Journal of Computer Science

Research Article

A Context-Aware Temporal Convolutional Network for Water Replacement Prediction in Catfish Biofloc Ponds With Imbalanced Event Handling

Yaddarabullah[1], **Inanpi Hidayati Sumiasih**[2] **and Mutiara Dewi Puspitawati**[2]

[1]*Department of Informatics, Universitas Trilogi, Jakarta, Indonesia*
[2]*Department of Agrotechnology, Universitas Trilogi, Jakarta, Indonesia*

Article history
Received: 28-04-2025
Revised: 02-12-2025
Accepted: 20-01-2026

Corresponding Author:
Yaddarabullah
Department of Informatics,
Universitas Trilogi, Jakarta,
Indonesia
Email: yaddarabullah@trilogi.ac.id

Abstract: Water replacement is a biologically critical yet under-automated decision in biofloc-based aquaculture systems. Mistimed actions can destabilize microbial ecosystems, elevate fish mortality, and compromise sustainability through excessive water usage. Traditional rule-based heuristics often fail to account for the nonlinear and multiscale dynamics of pond environments. To address this, we propose a Context-Aware Temporal Convolutional Network (CA-TCN), a rare-event classification framework that combines deep temporal modeling with aquaculture-specific logic. The CA-TCN combines a dilated Temporal Convolutional Network with biologically guided SMOTE + Tomek resampling, Focal Loss for imbalance-sensitive learning, ROC-based threshold calibration, and a rule-based override system for decision assurance. Trained on 213 real-world multivariate time-series samples, each consisting of 23 features across 5 sequential timesteps, representing sensor data for water quality (total dissolved solids, pH, dissolved oxygen, electrical conductivity), feeding events, and fish mortality. The proposed model achieved 98.68% accuracy, 1.0000 precision, 0.9737 recall, an F1-score of 0.9867, and a ROC-AUC of 1.0000 on the held-out test set, demonstrating its ability to identify rare yet operationally critical water replacement events with high precision. Ablation studies reveal the cumulative contributions of each component: +2.6% F1-score improvement from context-aware sampling, +1.3% gain from Focal Loss, a 2.56% reduction in false positives via threshold calibration, and a 0.9% recall increase due to rule-based override. Compared to state-of-the-art baselines, CA-TCN outperforms SMOTE only (F1 = 0.9600), SMOTE + ENN (F1 = 0.9610), and Tomek only (F1 = 0.2963), offering up to +69.04% F1-score improvement and eliminating all false negatives, a critical requirement in early warning systems for aquaculture risk mitigation. This work contributes a validated, domain-informed artificial intelligence pipeline that advances sustainable aquaculture management.

Keywords: Biofloc Aquaculture, Water Replacement Prediction, Temporal Convolutional Network, Imbalanced Classification, Context-Aware Machine Learning, Sustainable Aquaculture, Smart Farming

Introduction

The cultivation of catfish (Clarias sp.) through biofloc technology has emerged as a sustainable and resource-efficient aquaculture practice, particularly within integrated agricultural systems (Khanjani et al., 2023). Biofloc ponds leverage microbial communities to recycle nutrients and organic waste, creating a closed-loop environment that minimizes water usage while maintaining high productivity (Iber et al., 2025). This approach is especially valuable in regions facing water scarcity, land limitations, and increasing environmental pressures. Often, catfish biofloc ponds are co-located with hydroponic farms, composting units, and renewable energy systems, forming a synergistic ecosystem

SCIENCE Publications

that supports food security, economic resilience, and environmental stewardship (Augustiawan and Rahmadanih, 2023; Arthanawa et al., 2021). However, such integration introduces significant operational complexities, where maintaining water quality, ensuring fish health, and determining the appropriate timing for water replacement become critical yet challenging tasks (Arepalli and Naik, 2024).

While biofloc systems offer clear advantages in sustainability and operational efficiency, they also introduce complex challenges for decision-making under uncertainty. One particularly critical and underexplored task is predicting water replacement events, which requires balancing biological urgency with system-level stability (Campanati et al., 2022). Traditional approaches, often based on manual heuristics or static thresholds, struggle to adapt to the nonlinear and dynamic conditions typical of biofloc ponds (Raza et al., 2024). Given the intricate interactions among biological, chemical, and environmental parameters, threshold-based interventions are often insufficient for achieving precise, timely, and efficient management (Pham et al., 2023).

Recent research has emphasized the feasibility of deploying edge AI and IoT-based solutions for timely environmental interventions in aquaculture systems (Esty et al., 2025). These smart monitoring and control systems can continuously track critical parameters, such as temperature, pH, dissolved oxygen, and fish behavior, using IoT sensors (Abid et al., 2024). By integrating AI and machine learning algorithms, these systems can provide real-time data analysis and decision support for optimizing water quality management (Huang and Khabusi, 2025).

Consequently, Machine Learning (ML) has emerged as a promising avenue to address these challenges (Jin et al., 2024). By learning complex temporal patterns from multivariate historical data (Wang et al., 2023), ML models can improve the accuracy and responsiveness of critical operational decisions, such as the timely initiation of water replacement (Islam et al., 2023). Compared to conventional heuristics, ML approaches promise enhanced precision (Boyd et al., 2020), reduced unnecessary interventions (Chatziantoniou et al., 2023), and strengthened support for sustainable aquaculture management (Falconer et al., 2023).

However, common architectures such as Multi-Layer Perceptron Artificial Neural Networks (MLP-ANN), Long Short-Term Memory (LSTM), and Gated Recurrent Unit (GRU) models, though effective in modeling temporal patterns, exhibit several limitations when applied to decision-critical aquaculture tasks such as water replacement prediction. MLP-ANN models lack temporal awareness, treating each input as independent and thereby failing to capture sequential dependencies that are crucial for early detection of water quality deterioration (Eze et al., 2021). LSTM and GRU models, while designed for sequence modeling, often require large datasets to generalize effectively and are prone to overfitting in low-sample, imbalanced settings typical of rare-event detection (Rahul Gandh et al., 2023). Furthermore, these models generally do not incorporate biological logic or safeguard mechanisms, such as rule-based overrides, that are essential in ensuring domain-relevant decisions (Thelma et al., 2024). As a result, they may still produce false negatives in high-risk scenarios, undermining their reliability for real-world deployment in aquaculture monitoring systems (Chaudhary and Stojanović, 2023).

Applying ML to biofloc pond systems remains fraught with challenges. Central among them is the need to model intricate, nonlinear relationships between variables like Total Dissolved Solids (TDS), pH, ammonia levels, and fish mortality, which evolve over multiple temporal scales (Essamlali et al., 2024). Compounding this complexity is the rarity of water replacement events, creating a pronounced class imbalance that biases models toward majority conditions and increases the risk of false negatives (Habib et al., 2022). Moreover, many ML applications in aquaculture overlook domain-specific biological constraints (Alenezi and Alabaiadly, 2025), such as stage-specific TDS tolerances or the interactive effects of water quality parameters (Nguyen et al., 2019). The exclusion of such expert knowledge compromises model generalizability and interpretability, both essential for real-world deployment (Matondang et al., 2022). Despite growing interest in AI-driven precision aquaculture, existing studies have yet to comprehensively address the combined challenges of temporal modeling, class imbalance (Iqbal et al., 2024), noise sensitivity, and contextual reasoning within underactuated environments like biofloc ponds (Yang et al., 2021).

In addition to class imbalance, real-world biofloc datasets often suffer from class overlap and sensor noise, further complicating robust predictive modeling (Riantika et al., 2024). Standard resampling techniques, such as SMOTE, while mitigating class imbalance, risk generating synthetic samples in biologically implausible regions, thereby increasing false positives (Latief et al., 2024). Hybrid strategies such as SMOTE + ENN have been proposed to selectively clean noisy samples and improve classifier robustness (Pardede and Dika, 2024). Nevertheless, their effectiveness within biologically constrained domains remains largely unexplored.

Consequently, a systematic comparison between conventional resampling techniques and biologically contextualized augmentation methods is critical to ensure predictive fidelity and operational trustworthiness. Furthermore, biofloc pond systems are inherently dynamic, with key variables such as Total Dissolved Solids (TDS), Electrical Conductivity (EC), Dissolved Oxygen (DO), and ammonia exhibiting gradual temporal shifts rather than abrupt changes (Rizk et al., 2020). Static data representations risk missing early precursor patterns that could signal pending water quality deterioration (Yadav and Goyal, 2022). Therefore, temporally structured input formats are necessary to fully leverage the time-evolving nature of pond ecosystems.

Despite recent advancements, most existing machine learning applications in aquaculture continue to face three fundamental limitations. First, they often fail to adequately model temporal causality and account for delayed biological responses. Second, they lack mechanisms that are grounded in a biological context to effectively handle class imbalance. Third, these models typically do not include domain-aware logic that ensures the safety and interpretability of predictions. While models such as LSTM and CNN-LSTM hybrids have been applied to water quality forecasting, they generally do not incorporate biologically informed constraints or include rule-based safety mechanisms. Additionally, common resampling techniques like SMOTE are frequently applied without consideration of domain-specific thresholds, leading to the generation of synthetic samples that may be biologically implausible. These shortcomings highlight the urgent need for a unified modeling framework that integrates temporal deep learning with aquaculture-specific knowledge, embedding biological context at both the data preprocessing stage through guided sampling and the decision-making stage through rule-based correction.

Unlike traditional TCNs or CNN-LSTM hybrids, which primarily focus on capturing temporal patterns without contextual domain safeguards, the CA-TCN model is explicitly designed to operate under biological constraints of aquaculture systems. It enhances temporal convolution with biologically filtered resampling, rare-event learning via Focal Loss, and calibrated thresholding grounded in ROC optimization. Additionally, it uniquely incorporates rule-based overrides informed by critical operational thresholds (e.g., TDS, mortality), a mechanism absent in generic architectures. This domain-coupled design allows CA-TCN to achieve both high statistical performance and practical decision reliability in dynamic, imbalanced aquaculture environments.

To overcome the limitations of previous methods, this study introduces a Context-Aware Temporal Convolutional Network (CA-TCN), an end-to-end model designed for predicting water replacement needs in biofloc catfish cultivation. The CA-TCN uses dilated temporal convolution to learn short- and long-term patterns, applies SMOTE + Tomek to handle class imbalance without breaking biological relationships in the data, and adopts Focal Loss to focus learning on rare but critical cases. Sliding multivariate time windows are used to preserve temporal structure and capture environmental trends. In addition, a simple rule-based mechanism is included to adjust model outputs based on key biological thresholds, such as extreme TDS levels or signs of fish stress. Together, these components create a reliable and interpretable model with high predictive performance and strong suitability for real-world aquaculture operations.

Related Work

Artificial intelligence applications in aquaculture have evolved from predicting individual water quality parameters to more complex tasks such as smart feeding and disease detection (Yang et al., 2025). However, water replacement event prediction remains under-addressed despite its critical impact. A persistent challenge in this domain is the treatment of data imbalance, particularly in event-driven classification tasks. While techniques like SMOTE are widely used for imbalance, they often ignore biological plausibility, which can reduce real-world applicability (Alahmari, 2020). Furthermore, many ML models rely on standard loss functions and static thresholds that fail to penalize misclassification of rare but critical events. Although ML research has advanced in forecasting and fish health monitoring, few models integrate temporal reasoning, rare-event classification, and domain-aware interpretability (Eze et al., 2023). While rare-event classification is increasingly studied in time-series ML (Shi, 2023), aquaculture-specific implementations that preserve biological plausibility remain underdeveloped (Moghadam et al., 2020). Recent studies have explored deep learning architectures such as LSTM, CNN, and TCN for multivariate aquaculture forecasting (Pourmoradian et al., 2024), yet few focus on binary rare-event prediction like water replacement or incorporate domain-specific biological rules into the modeling pipeline.

In response to these gaps, this study proposes a Context-Aware Temporal Convolutional Network (CA-TCN) that integrates biologically guided resampling and calibrated decision thresholds. Beyond proposing a new architecture, this study also conducts a structured ablation analysis, systematically isolating the effects of each modeling component, resampling strategy, focal loss, threshold tuning, and rule-based overrides on predictive accuracy, precision, recall, and operational reliability. This comprehensive evaluation advances the development of holistic, biologically aligned models tailored for decision-critical applications in complex aquaculture environments.

Methods

To improve clarity and support reproducibility, Figure 1 summarizes the entire modeling workflow used in this study. The pipeline begins with multivariate sensor acquisition from the biofloc pond system, followed by sequence construction and biologically guided preprocessing. Context-aware resampling (SMOTE + Tomek), feature standardization, and temporal windowing are applied before samples are passed into the proposed CA-TCN architecture.

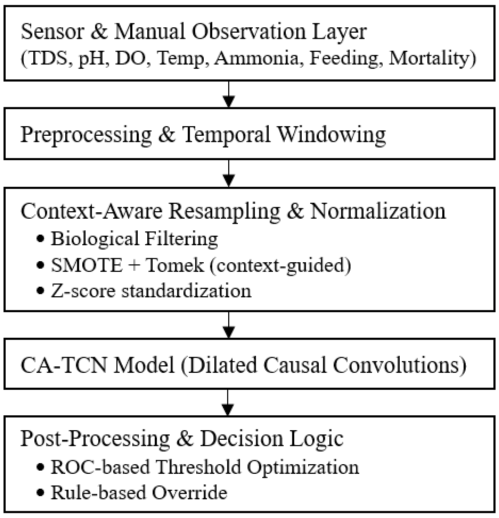

Fig. 1: Summary Diagram

The CA-TCN produces probability scores that are subsequently refined through ROC-based threshold calibration and a rule-based override mechanism to ensure biological safety. This summary diagram provides a concise overview of the data flow and model logic described in detail in the following subsections.

This study adopts a multi-stage experimental workflow designed to develop a robust, context-aware deep learning model for predicting water replacement events in biofloc-based catfish aquaculture systems. The methodology integrates several critical components: Continuous time-series data acquisition with manual cross-checks, temporal window construction to preserve sequential structure, and class-imbalance handling using biologically filtered SMOTE + Tomek. A Temporal Convolutional Network (TCN) backbone is then developed and trained under imbalance-sensitive learning objectives. Subsequently, domain-specific decision rules are integrated to support safety-critical predictions. Finally, the model is evaluated using comprehensive performance metrics, including accuracy, recall, F1-score, ROC-AUC, and computational cost to assess both predictive effectiveness and suitability for real-world deployment.

Dataset Description

This study utilizes time-series data obtained from real-world monitoring of catfish in biofloc ponds operating within an integrated agricultural environment. The raw dataset was compiled from multiple calibrated sensors and manual observations recorded periodically throughout the cultivation cycle. It includes critical aquaculture parameters such as Total Dissolved Solids (TDS), Electrical Conductivity (EC), Dissolved Oxygen (DO), water and air temperature, ammonia levels, pH, feeding intensity, aeration status, and observed fish mortality. All environmental sensors were calibrated according to manufacturer guidelines before deployment. pH probes used two-point buffer calibration (pH 4.01 and 7.00), TDS and EC sensors were checked with conductivity standards to maintain ±2% accuracy, and DO sensors were calibrated using air-saturated water. Data were collected every hour throughout the cultivation cycle to ensure consistent temporal resolution. Manual inputs, such as feeding and mortality records, were cross-checked with technician logs to improve reliability. No additional feature engineering was performed; the model was trained directly on the raw multivariate time-series data structured into fixed-length sequences. The main measured features in the raw dataset are listed in Table 1.

Preliminary analysis of the raw features revealed varying degrees of temporal fluctuation. Notably, Total Dissolved Solids (TDS) and Dissolved Oxygen (DO) exhibited the highest variance across the cultivation cycle, indicating strong responsiveness to environmental and biological dynamics such as feeding and aeration. Conversely, pH and air temperature showed relatively low temporal variability, suggesting their role as background conditions rather than primary triggers for intervention. These patterns influenced feature importance in downstream modeling, where high-variance features contributed more prominently to decision boundaries for water replacement detection.

To enable supervised learning using Temporal Convolutional Networks (TCN), which require temporally ordered input sequences, the dataset was converted into 5-step sliding windows, each containing 23 features per timestep.

Table 1: Raw Features of Original Dataset

Feature Name	Unit	Description
Total Dissolved Solids (TDS)	ppm	Concentration of dissolved solids
Electrical Conductivity (EC)	µS/cm	Water conductivity
Water Temperature	°C	Temperature of pond water
Air Temperature	°C	Ambient temperature above the pond
Dissolved Oxygen (DO)	mg/L	Oxygen availability in water
Ammonia	ppm	Ammonia concentration
pH	%	Water acidity/alkalinity level
Feeding Intensity	g/day	Feeding input per day
Aeration Status	Boolean (On/Off)	Aeration device operational status
Fish Mortality Indicator	Boolean (Yes/No)	Presence of fish deaths

This results in an input tensor of shape (samples, timesteps, features), where timesteps = 5 and features = 23. Each sequence, therefore, captures dynamic changes in water quality and pond conditions over a temporal span, preserving both short-term fluctuations and transitional patterns relevant to water replacement decisions. The flattened version of this tensor yields 115 features per instance (i.e., 5 timesteps × 23 features), with column names structured as *fX_tY*, where *fX* denotes the base feature index, and tY indicates the time step position (e.g., f3_t2 is feature 3 at timestep 2 in the sequence). Each sequence was labeled as 1 if a water replacement occurred during or immediately after the window, and 0 otherwise, based on technician logbooks. To enrich interpretability, domain-informed features were added, including rolling statistics of TDS and pH, mortality flags, and critical condition interactions. All features were standardized using z-score normalization.

Imbalanced Data Handling Strategy

To address the significant class imbalance observed in the dataset where water replacement events comprise a minority subset, we propose a three-stage sampling framework that integrates biologically grounded filtering, synthetic oversampling, and decision boundary refinement.

Biological Filtering of Minority Class Instances

The initial step involves selecting only biologically relevant samples from the minority class before applying any resampling. A sample is retained if it satisfies at least one of the following criteria:

(i) The average Total Dissolved Solids (TDS) across the five timesteps exceeds 800 ppm
(ii) The average pH falls outside the optimal physiological range of 6.8 to 8.8
(iii) At least one timestep within the sequence shows a non-zero fish mortality indicator

These thresholds were derived from operational guidelines and empirical field knowledge.

Synthetic Minority Oversampling (SMOTE)

From the filtered subset, new synthetic minority samples are generated using the SMOTE technique. Each synthetic sample is created by interpolating between two real minority instances using the formula:

$$x' = x1 + \delta(x2 - x1) \quad (1)$$

Where x_1 and x_2 are feature vectors from the filtered minority class, and δ is a random number between 0 and 1. This approach effectively expands the minority region in the feature space, allowing the model to better learn its representation.

Tomek Links Removal for Boundary Cleaning

After oversampling, we apply Tomek Links removal to clean the decision boundary between classes. A Tomek Link is defined as a pair of samples from different classes that are each other's nearest neighbor. Removing such pairs, especially the majority class instance, helps eliminate borderline noise and enhances the separability of the classifier. This is particularly useful in imbalanced settings, where false negatives near the boundary are common and detrimental. Through this three-stage pipeline, biological filtering, context-aware SMOTE, and Tomek-based boundary cleaning, the proposed method ensures that oversampling does not violate aquaculture-specific logic and avoids generating synthetic minority samples that are not operationally meaningful.

Proposed Model Architecture: Context-Aware Temporal Convolutional Network

To address the challenge of predicting rare and operationally critical water replacement events in biofloc pond systems, this study proposes a Context-Aware Temporal Convolutional Network (CA-TCN). The model is designed to learn from short multivariate time series windows. These features capture key aquaculture indicators such as TDS, pH, DO, ammonia, feeding behavior, and fish mortality. The complete architecture of the CA-TCN model is illustrated in Figure 1, illustrating the data flow from raw multivariate input to prediction, including context-aware sampling, temporal convolution layers, Focal Loss optimization, threshold calibration, and rule-based override.

Figure 2 illustrates the complete architecture of the proposed Context-Aware Temporal Convolutional Network (CA-TCN), designed to predict rare and operationally critical water replacement events in biofloc systems.

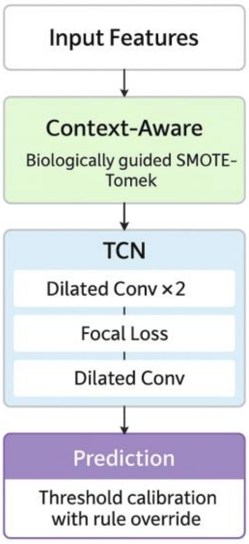

Fig. 2: Architecture of CA-TCN Model

The model processes a 5×23 input matrix, representing five sequential time steps of multivariate aquaculture data including water quality (e.g., TDS, pH, DO, EC), operational actions (e.g., feeding), and biological responses (e.g., mortality). The first major component is the Context-Aware Sampling Block, which applies SMOTE + Tomek resampling to address class imbalance. Unlike generic oversampling, this block uses biological constraints to selectively generate synthetic minority samples only when domain-relevant thresholds (e.g., TDS > 800 ppm, DO < 3 ppm) are met. This ensures that augmented examples remain biologically plausible and contribute operational relevance to the training data.

Following this sampling phase, the balanced sequences are passed to the Temporal Convolutional Network (TCN) module, which serves as the core learning engine of the architecture. This module is explicitly designed to capture multi-scale temporal patterns and causal dependencies across the sequence. The details of this component are elaborated in the following subsection.

Temporal Convolutional Structure

At the heart of the proposed model lies a temporal convolutional architecture composed of a series of One-Dimensional (1D) causal convolutional layers, each configured with progressively increasing dilation rates. This design captures both short- and long-term dependencies, aligning with aquaculture conditions where effects may be immediate (e.g., DO drop) or delayed (e.g., gradual TDS buildup affecting fish health). Capturing these multiscale temporal dependencies is critical for reliably predicting water replacement events. A causal convolution is a special type of convolution where the output at a given timestep t is strictly based on input values from timestep t and earlier, i.e., inputs at times $\leq t$. This ensures that the model respects the temporal causality inherent in time-series data and does not incorporate information from the future.

To expand the receptive field without excessively deep networks, the model uses dilated convolutions. In this context, dilation refers to the spacing between elements in the convolutional kernel. Instead of applying the filter to consecutive timesteps, a dilation rate of d applies the filter to every d-th timestep, allowing the model to “skip” over intermediate values and look further back in time. This approach enables the receptive field to grow exponentially with depth. For a given input sequence X = $[x1, x2, ..., x_T]$, where each $x_t \in R^{23}$, a 1D dilated convolution at dilation rate d and kernel size k is computed as:

$$yt = \sum_{i=0}^{k-1} w_i \cdot x_{t-d \cdot i} \tag{2}$$

Where wi are learnable convolutional weights, and d controls the spacing between filter taps. By stacking layers with exponentially increasing dilation rates (e.g., d = 1,2,4), the network's receptive field grows exponentially, enabling the model to capture both short-term fluctuations and long-term buildup effects, an essential property in predicting delayed responses to environmental stressors such as rising ammonia or cumulative TDS. The network architecture includes:

- Three causal convolutional layers, each with 64 filters and a kernel size of 3
- Dilation rates of 1, 2, and 4, respectively, to form a receptive field covering 9 timesteps
- Dropout layers (rate = 0.3) after each convolution to mitigate overfitting
- Global average pooling to reduce the temporal dimension while preserving feature aggregation

Dense and Output Layers

The temporal feature map generated by the convolutional layers is fed into a stack of fully connected layers:

- First dense layer with 128 units, followed by ReLU activation and dropout (rate = 0.2)
- Second dense layer with 64 units, also followed by ReLU and dropout (rate = 0.2)
- Last dense layer with 1 unit and sigmoid activation to produce a probability estimate y ∈ [0,1] representing the likelihood of a water replacement event

Focal Loss for Rare Events

Due to the imbalanced nature of the dataset, the model uses Focal Loss (Usuga-Cadavid et al., 2021), which dynamically scales the contribution of each training example based on its classification difficulty, placing more focus on hard-to-classify minority instances. The focal loss is defined as:

$$L_{focal}(p_t) = -\alpha_t (1-p_t)^{\gamma} \log(p_t) \tag{3}$$

Where:

- pt is the model's estimated probability for the true class
- αt is the class balancing parameter (e.g., 0.25 for the minority class)
- γ is the focusing parameter, typically set between 1 and 2 (in this study, $\gamma = 2$)

This loss function penalizes easy negative examples less and encourages the model to focus on rare positive instances (i.e., water replacement events).

Threshold Calibration Through ROC-AUC Optimization

Rather than using the default classification threshold of 0.5, we implement threshold calibration based on ROC-AUC analysis. After training, the optimal threshold $\tau*$ is determined by maximizing the geometric mean (G-mean) of True Positive Rate (TPR) and True Negative Rate (TNR):

$$T^{*} \arg\max_{T} \sqrt{TPR(T) \cdot TNR(T)} \tag{4}$$

This strategy enhances the model's balance between sensitivity and specificity, particularly under high-imbalance scenarios where a static threshold often fails to detect minority events.

Rule-Based Biological Override

To ensure biological and operational safety, the model's predictions are adjusted using rule-based overrides: if $y < \tau*$ but TDS > 1000 ppm or mortality occurs, the final output is set to 1.

Algorithmic Flow of CA-TCN

The algorithmic flow of the proposed Context-Aware Temporal Convolutional Network (CA-TCN), illustrated in Figure 3, integrates biologically guided preprocessing with a temporally structured deep learning architecture to detect critical water replacement events in biofloc systems.

The process begins by preprocessing raw time-series aquaculture data and segmenting it into overlapping sliding windows, each spanning five consecutive timesteps. Each window comprises multivariate observations, including water quality metrics such as Total Dissolved Solids (TDS), Dissolved Oxygen (DO), pH, and ammonia, alongside operational events such as feeding intensity and fish mortality. To address class imbalance while preserving biological relevance, the model applies SMOTE + Tomek only to biologically critical windows.

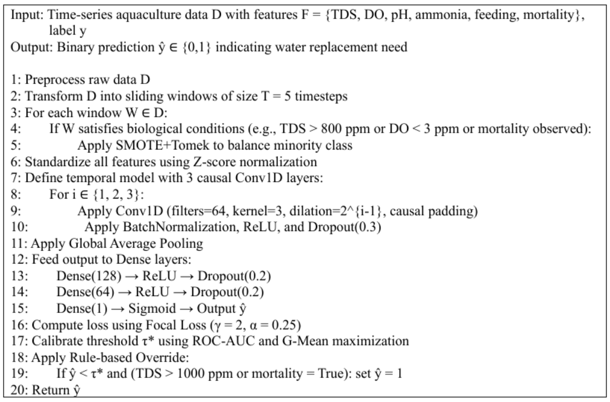

```
Input: Time-series aquaculture data D with features F = {TDS, DO, pH, ammonia, feeding, mortality},
       label y
Output: Binary prediction $\hat{y} \in \{0,1\}$ indicating water replacement need

Preprocess raw data D
Transform D into sliding windows of size T = 5 timesteps
For each window $W \in D$:
    If W satisfies biological conditions (e.g., TDS > 800 ppm or DO < 3 ppm or mortality observed):
        Apply SMOTE+Tomek to balance minority class
Standardize all features using Z-score normalization
Define temporal model with 3 causal Conv1D layers:
    For $i \in \{1, 2, 3\}$:
        Apply Conv1D (filters=64, kernel=3, dilation=2^{i-1}, causal padding)
        Apply BatchNormalization, ReLU, and Dropout(0.3)
Apply Global Average Pooling
Feed output to Dense layers:
    Dense(128) → ReLU → Dropout(0.2)
    Dense(64) → ReLU → Dropout(0.2)
    Dense(1) → Sigmoid → Output $\hat{y}$
Compute loss using Focal Loss ($\gamma = 2, \alpha = 0.25$)
Calibrate threshold $\tau*$ using ROC-AUC and G-Mean maximization
Apply Rule-based Override:
    If $\hat{y} < \tau*$ and (TDS > 1000 ppm or mortality = True): set $\hat{y} = 1$
Return $\hat{y}$
```

Fig. 3: Pseudocode Flow of CA-TCN

These include sequences with high TDS (above 800 ppm), low DO (below 3 mg/L), or observed mortality. The preprocessing stage concludes with z-score normalization. After resampling, all features are standardized. The processed sequences are fed into a temporal convolutional backbone with three causal 1D convolution layers using dilation rates of 1, 2, and 4, enabling the model to learn both short- and long-term patterns. Each layer applies batch normalization, ReLU activation, and a dropout rate of 0.3 to control overfitting.

A global average pooling layer reduces the temporal features into a fixed-length vector, which then flows through two ReLU-activated dense layers with dropout, followed by a sigmoid layer that predicts the probability of a water replacement event. To better detect rare cases, the model uses Focal Loss ($\gamma = 2$, $\alpha = 0.25$), which gives more weight to hard-to-classify samples. After training, the classification threshold is optimized using ROC-AUC-based G-Mean maximization. Finally, a rule-based override ensures safety: If the predicted probability is below the threshold but TDS exceeds 1000 ppm, or mortality occurs, the output is forced to indicate a replacement. Figure 2 highlights the step-wise integration of biologically informed preprocessing and deep temporal learning. The process initiates with raw multivariate time-series segmentation into overlapping windows. Critical domain-aware filters are applied to identify high-risk segments (e.g., TDS > 800 ppm, DO < 3 ppm), which are selectively oversampled using SMOTE + Tomek to preserve ecological realism. These filtered sequences are normalized and passed into a three-layer causal convolutional stack with dilation rates [1, 2, 4], enabling the model to learn both immediate and lagged influences across time. After global average pooling and dense layer transformations, the sigmoid output estimates the likelihood of water replacement.

A threshold calibrated using ROC-AUC G-Mean helps balance recall and precision. To enhance safety, a domain-specific rule-based override is applied after the sigmoid output: If critical conditions occur (e.g., TDS > 1000 ppm or signs of fish death), the system triggers an alert regardless of the predicted probability. This ensures safety-critical decisions align with aquaculture knowledge.

Model Training Configuration

The dataset was split into training and validation sets using an 80/20 stratified split, ensuring that both subsets retained the original class distribution. Stratification is particularly important in this context to preserve the relative rarity of water replacement events in both sets, which supports fair and consistent evaluation of minority class performance. The hyperparameters of the model were described in Table 2. The model was trained using the Adam optimizer, which often exhibits deeper layer stacks for effective backpropagation through dilated convolutions.

Table 2: Model Training Parameters

Parameter	Description
Optimizer	Adam
Initial Learning Rate	0.001
Batch Size	16
Number of Epochs	30
Early Stopping	Enabled (patience = 5 epochs, monitor = validation loss)
Loss Function	Focal Loss (γ = 2.0, α = 0.25)
Class Weighting	Inverse frequency-based (computed on training split)
Validation Split	20% (stratified)
Sequence Length	5 timesteps
Number of Features per Step	23
Dropout (Conv layers)	0.3
Dropout (Dense layers)	0.2
Convolutional Filters	64 filters per layer × 3 layers
Dilation Rates	[1, 2, 4]
Kernel Size (Conv1D)	3
Pooling	Global Average Pooling
Activation Functions	ReLU (hidden layers), Sigmoid (output layer)
Threshold Optimization	ROC-AUC-based (G-mean maximization)
Post-Processing Logic	Rule-based override for TDS, pH, and mortality indicators

The initial learning rate was set to 0.001, which was empirically determined through preliminary experiments. The Adam optimizer is particularly suited to TCN. A batch size of 16 was used during training, allowing efficient utilization of GPU memory while preserving the stochastic nature of gradient updates. The model was trained for 30 epochs. Early stopping monitors the validation loss and halts training if no improvement is observed for a patience period of five consecutive epochs. In addition to the Context-Aware SMOTE + Tomek sampling strategy, the focal loss parameters ($\alpha = 0.25$, $\gamma = 2.0$) and class weights were configured at the start of training and held constant across all 30 epochs. Similarly, the SMOTE + Tomek resampling was performed as a pre-processing step prior to each training run and was not dynamically updated across epochs.

This approach assigns a higher penalty to misclassified minority class samples, reinforcing the model's focus on detecting water replacement events and improving minority recall without overly compromising specificity. To prevent overfitting and enhance model generalization, dropout regularization was applied after each convolutional and dense layer. A dropout rate of 0.3 was used for convolutional layers and 0.2 for fully connected layers. These settings were chosen specifically to control overfitting in a small-sample regime, together with class weighting and focal loss. This mechanism reduces co-adaptation by suppressing neuron activations, enhancing temporal pattern learning across diverse environmental conditions.

The selected hyperparameters were tuned based on preliminary grid search and cross-validation. The dropout rate of 0.3 in convolutional layers effectively reduced overfitting without sacrificing convergence speed, especially in the presence of temporally correlated features. Additionally, the dilation rates [1, 2, 4] expanded the receptive field to capture long-range dependencies while maintaining training stability, avoiding gradient vanishing issues commonly encountered in deeper convolutional stacks. This configuration balanced temporal resolution and computational efficiency, contributing to smoother validation loss curves and improved generalization across unseen sequences.

Evaluation and Validation

The performance of the proposed model was assessed through a structured evaluation framework that integrates both statistical metrics and real-world applicability. Recognizing the imbalanced nature of the dataset, where water replacement events are significantly underrepresented, standard accuracy metrics were supplemented with more informative indicators that emphasize the minority class. These include precision, recall, F1-score, and the area under the Receiver Operating Characteristic curve (ROC-AUC) (Richardson et al., 2024), particularly evaluated for class 1, which corresponds to actual water replacement conditions. Accuracy was computed as the proportion of correctly predicted instances across all classes. Precision for class 1 quantifies how many of the model's positive predictions were actually correct, reflecting the model's ability to avoid false alarms. Recall for class 1 measures how many true water replacement events were successfully detected. The F1-score was calculated, providing a single metric that accounts for both false positives and false negatives.

Ablation Study

To investigate the relative importance of each component in the proposed framework, we conducted an ablation study by incrementally disabling individual mechanisms and measuring the resulting changes in model performance. Specifically, four ablation scenarios were tested:

- No Focal Loss: Replaced with binary cross-entropy
- No Threshold Optimization: Fixed threshold at 0.5
- No Rule Override: Removal of domain-based post-logic
- No Contextual Filtering in SMOTE: Standard SMOTE applied to all class 1 samples

State-of-the-Art Comparison

To rigorously evaluate the performance of the proposed Context-Aware Temporal Convolutional Network (CA-TCN), a series of benchmarking experiments was conducted against widely adopted imbalanced learning strategies. This comparative analysis aimed to isolate the added value of biologically informed augmentation and contextual threshold calibration by holding the core model architecture (i.e., TCN) constant across all variants. The following resampling techniques were used for benchmarking:

- SMOTE only: An oversampling method that creates new minority samples by interpolating between existing ones, helping balance the class distribution without modifying the majority data
- Tomek only: An under-sampling method that removes Tomek link pairs opposite-class samples that are each other's nearest neighbours to create cleaner decision boundaries, though it may also discard useful data
- SMOTE + ENN: A hybrid resampling strategy that first applies SMOTE for minority class augmentation and subsequently uses Edited Nearest Neighbours (ENN) to filter out noisy or misaligned samples, aiming to refine the sample distribution and improve model robustness

All models were evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and ROC-AUC, along with a class-wise performance breakdown to assess their ability to detect rare but biologically important water-replacement events. Beyond predictive accuracy, a computational cost evaluation was also performed, including training time, inference latency, CPU usage, and memory overhead, to compare their practical deployability in real-time biofloc monitoring. To contextualize the performance of the proposed Context-Aware Temporal Convolutional Network (CA-TCN), we also compared it with several mainstream machine-learning models, including MLP-ANN, LSTM, GRU, Random Forest, SVM, and XGBoost.

Feature Importance

To enhance model interpretability and provide insights into the decision-making process of the CA-TCN, this study incorporates Shapley Additive exPlanations (SHAP) as a post-hoc feature attribution technique. SHAP values were computed using the Deep Explainer wrapper for the trained model to quantify each feature's contribution to individual predictions across the dataset. Given the temporal structure of the input (five-time steps and 23 features), SHAP was applied on the flattened input tensor, allowing identification of both the most influential features and their critical temporal positions. The analysis revealed that TDS, DO, pH, and Ammonia levels, especially at later timesteps, were consistently among the top contributors to water replacement decisions.

Results and Discussion

Dataset Description

The final dataset used in this study consists of 213 multivariate time-series samples derived from real-world monitoring of catfish biofloc pond systems. Each sample represents a five-timestep window, where each timestep includes 23 aquaculture-related features, resulting in 115 input features and 1 binary target label. The features include normalized values for critical variables such as Total Dissolved Solids (TDS), pH, Dissolved Oxygen (DO), temperature, ammonia, feeding intensity, and fish mortality indicators. All features were pre-processed through standardization to support stable model convergence. The structured nature of each sample is illustrated in Figure 4, showcasing the multivariate sequences across five timesteps and 23 features per window.

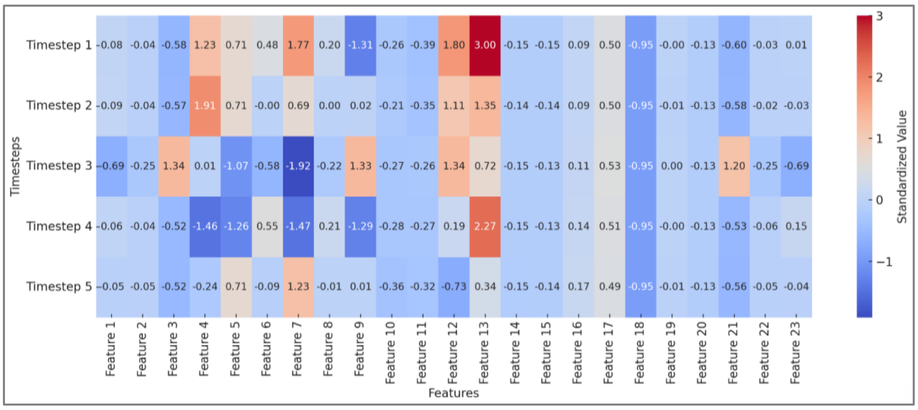

	Feature 1	Feature 2	Feature 3	Feature 4	Feature 5	Feature 6	Feature 7	Feature 8	Feature 9	Feature 10	Feature 11	Feature 12	Feature 13	Feature 14	Feature 15	Feature 16	Feature 17	Feature 18	Feature 19	Feature 20	Feature 21	Feature 22	Feature 23
Timestep 1	-0.08	-0.04	-0.58	1.23	0.71	0.48	1.77	0.20	-1.31	-0.26	-0.39	1.80	3.00	-0.15	-0.15	0.09	0.50	-0.95	-0.00	-0.13	-0.60	-0.03	0.01
Timestep 2	-0.09	-0.04	-0.57	1.91	0.71	-0.00	0.69	0.00	0.02	-0.21	-0.35	1.11	1.35	-0.14	-0.14	0.09	0.50	-0.95	-0.01	-0.13	-0.58	-0.02	-0.03
Timestep 3	-0.69	-0.25	1.34	0.01	-1.07	-0.58	-1.92	-0.22	1.33	-0.27	-0.26	1.34	0.72	-0.15	-0.13	0.11	0.53	-0.95	0.00	-0.13	1.20	-0.25	-0.69
Timestep 4	-0.06	-0.04	-0.52	-1.46	-1.26	0.55	-1.47	0.21	-1.29	-0.28	-0.27	0.19	2.27	-0.15	-0.13	0.14	0.51	-0.95	-0.00	-0.13	-0.53	-0.06	0.15
Timestep 5	-0.05	-0.05	-0.52	-0.24	0.71	-0.09	1.23	-0.01	0.01	-0.36	-0.32	-0.73	0.34	-0.15	-0.14	0.17	0.49	-0.95	-0.01	-0.13	-0.56	-0.05	-0.04

Fig. 4: Sliding Window Heatmap (Timesteps x Features)

The dataset used in this study is publicly available on Mendeley Data to ensure transparency and reproducibility (https://data.mendeley.com/datasets/w5gnr5wk7f/2).

The horizontal axis corresponds to all 23 features, while the vertical axis shows their values across the five temporal positions. As the figure demonstrates, feature values fluctuate across time, reflecting real-world variations in pond conditions. Notably, certain features, such as those associated with ammonia and DO, exhibit prominent temporal shifts, highlighting their relevance to the model's ability to capture early signs of environmental stress. This representation confirms that the temporal encoding of domain variables is preserved in the model input and forms a meaningful basis for learning dynamic patterns relevant to water replacement decisions.

Each sample is labeled as 1 if a water replacement event occurred during or shortly after the sequence, and 0 otherwise. Labeling was performed manually using technician logbooks, and only sequences with biologically valid patterns were retained; ambiguous entries were excluded. This led to a class imbalance, with water replacement events (target = 1) comprising approximately 11.3% of the total data. The class distribution was balanced to support fair learning and evaluation.

As seen in Figure 3, the temporal heatmap reveals notable shifts in specific aquaculture features over time. In particular, Dissolved Oxygen (DO) levels often decline toward the final timesteps, while ammonia concentrations tend to rise, signaling deteriorating water quality that may precede a replacement event. These dynamic trends provide early indicators of biological stress, reinforcing their predictive utility. The heatmap also illustrates the effectiveness of the sliding window approach in capturing short-term transitions and abrupt environmental changes that might be overlooked in single-timestep models. Such representations validate the model's design to learn time-dependent cues embedded within multivariate sequences.

The Principal Component Analysis (PCA) projection in Figure 5 illustrates the two-dimensional distribution of the class-balanced dataset following the imbalanced data handling strategy augmentation. The majority class (label 0, red) appears densely clustered along the central band of the feature space, indicating relatively low variance in its underlying characteristics. In contrast, the minority class (label 1, blue) demonstrates broader dispersion across both PCA components, suggesting a more heterogeneous representation and successful augmentation across diverse regions of the feature space.

The broader spread of the minority class (blue) in Figure 4 validates the effectiveness of the contextual SMOTE strategy in populating diverse and meaningful areas of the decision space. Unlike conventional SMOTE, which may produce synthetic points clustered around dense regions, the context-aware approach used here intentionally targets ecologically relevant but underrepresented segments (e.g., high TDS, low DO). This explains the dispersion of class 1 across both PCA components, enabling the model to generalize better to rare but critical water replacement scenarios. In contrast, the majority class (class 0, red) forms a compact cluster, reflecting its dominance in the original dataset and relative homogeneity in non-critical pond states.

Performance of the Proposed Model

The proposed CA-TCN model exhibited strong performance in identifying rare but biologically critical water replacement events in biofloc aquaculture systems. As shown in Figure 6, the confusion matrix provides a comprehensive overview of the classification outcomes. Of the 38 actual water replacement events, the model correctly predicted 37 as positive cases (true positives), with only one instance misclassified as negative (false negative).

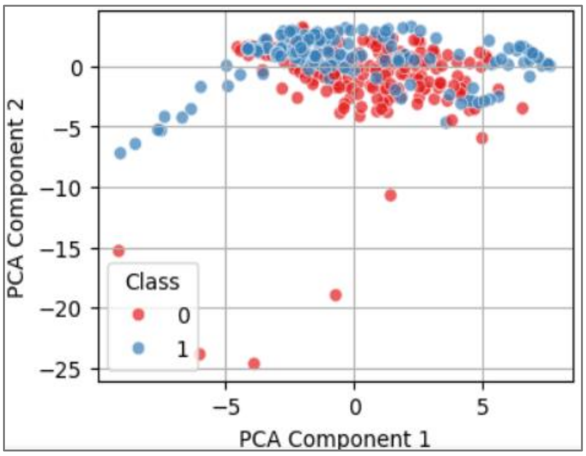

Fig. 5: Distribution of Water Replacement Events

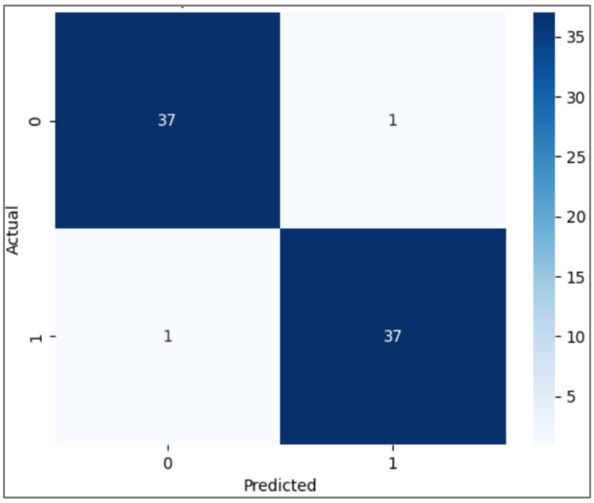

Fig. 6: Confusion Matrix of Proposed Model

This false negative represents a rare occurrence in which a critical replacement event was not detected. Conversely, among the 38 sequences that did not require water replacement, the model accurately classified 37 as negative (true negatives), with only one false positive, where a normal condition was mistakenly flagged as requiring intervention. These results demonstrate the model's high degree of reliability and its balanced ability to minimize both false alarms and missed detections in the context of aquaculture monitoring.

To better understand the model's limitations, the two misclassified cases were examined in detail. The single false negative occurred in a borderline condition where Total Dissolved Solids (TDS) and Dissolved Oxygen (DO) values hovered close to but just below critical thresholds, and no mortality was recorded. This subtle profile made the event difficult to distinguish from normal pond states, underscoring the challenge of borderline cases. The false positive, on the other hand, was associated with a sequence that displayed transient DO fluctuations that mimicked a replacement scenario, even though no intervention was required. These examples show that misclassifications are not random but occur under threshold-adjacent conditions, reinforcing the importance of adaptive thresholds and larger, multi-site datasets to further strengthen robustness.

As shown in Table 3, the Proposed Model, which integrates all four components, achieved the highest and most balanced performance, including perfect precision (1.0000) for class 1, an F1-score of 0.9867, and a ROC-AUC of 1.0000. This configuration serves as the performance benchmark. Removing Focal Loss resulted in a uniform precision, recall, and F1-score of 0.9737 for both classes, with a noticeable drop in overall ROC-AUC to 0.9861 (–1.39%). This decline indicates reduced sensitivity to the minority class (water replacement events), affirming the role of Focal Loss in directing the model's attention to rare and difficult examples by down-weighting easy negatives.

In the No Threshold Optimization scenario, a fixed 0.5 threshold increased class 1 recall to 100% (+2.63%) but decreased precision to 0.9744 (–2.56%). Although the overall F1-score marginally increased to 0.9870, this came at the cost of more false positives, which, in operational contexts, may lead to unnecessary water changes or alerts, highlighting the benefit of G-mean-based threshold calibration to balance specificity and sensitivity. Disabling the Rule Override, which enforces domain logic based on TDS, pH, and mortality, maintained a high class 1 recall (97.37%) but slightly reduced class 0 performance, causing the overall F1-score to drop to 0.9737 and ROC-AUC to 0.9986. This suggests that the biological post-filter helps recover borderline true positives that statistical thresholds might otherwise miss, reinforcing its role in handling edge cases critical to aquaculture management.

The most substantial performance decline occurred when Context Filtering during SMOTE oversampling was disabled. Class 0 recall dropped to 92.11%, and class 1 precision decreased to 0.9487, resulting in a class 0 F1-score of just 0.9459 and a reduced ROC-AUC of 0.9709. This outcome confirms that naive oversampling without context introduces noise and synthetic outliers, harming generalization. The biologically-informed filtering criteria ensure that synthetic samples represent valid minority cases, preserving decision boundaries and improving real-world robustness. These improvements, however, must be interpreted with caution, as synthetic augmentation may also introduce subtle biases despite biological filtering.

State-of-the-Art Comparison

To evaluate the performance of the proposed Context-Aware Temporal Convolutional Network (CA-TCN), a comprehensive comparison was conducted against two categories of benchmarks:

(1) established imbalance-handling techniques, including SMOTE only, Tomek only, and SMOTE + ENN
(2) alternative machine learning models, including Multi-Layer Perceptron Artificial Neural Network (MLP-ANN), Long Short-Term Memory (LSTM), and Gated Recurrent Unit (GRU)

All models were trained using the same Temporal Convolutional Network (TCN) backbone architecture, where applicable, and applied to identical stratified data splits. Consistent hyperparameter settings and standardized evaluation protocols were maintained throughout to ensure fairness and experimental control in the comparison. As summarized in Table 4, the CA-TCN model outperformed all alternatives across both quantitative metrics and operational robustness. Specifically, CA-TCN achieved an accuracy of 98.68%, precision of 100.00%, recall of 97.37%, F1-score of 98.67%, and a ROC-AUC of 1.0000.

Table 3: Ablation Study: Impact of Component Removal on Performance

Model Variant	Accuracy	Class 0			Class 1			ROC-AUC
		Precision	Recall	F1-Score	Precision	Recall	F1-Score	
Proposed Model	0.9868	0.9730	0.9474	0.9600	1.0000	0.9737	0.9867	1.0000
No Focal Loss	0.9737	0.9737	0.9737	0.9737	0.9737	0.9737	0.9737	0.9861
No Threshold Optimization	0.9868	1.0000	0.9737	0.9867	0.9744	1.0000	0.9870	1.0000
No Rule Override	0.9737	0.9744	1.0000	0.9870	0.9737	0.9737	0.9737	0.9986
No Context Filtering	0.9605	0.9722	0.9211	0.9459	0.9487	0.9737	0.9610	0.9709

Table 4: Model Performance Comparison with Alternative Imbalance Handling Techniques

Model Variant	Accuracy	Class 0			Class 1			ROC-AUC
		Precision	Recall	F1-Score	Precision	Recall	F1-Score	
Proposed Model	0.9868	0.9730	0.9474	0.9600	1.0000	0.9737	0.9867	1.0000
SMOTE only	0.9605	0.9737	0.9737	0.9737	0.9730	0.9474	0.9600	0.9945
Tomek only	0.5476	0.9310	0.7297	0.8182	0.1818	0.8000	0.2963	0.7189
SMOTE+ENN	0.9455	0.9412	0.9412	0.9412	0.9487	0.9737	0.9610	0.9798

These results highlight not only statistical superiority but also the model's enhanced reliability in detecting minority class events while minimizing false positives. In comparison, the SMOTE-only model demonstrated competitive yet slightly inferior results, achieving an F1-score of 96.00% for the minority class, approximately 2.67% lower than CA-TCN. Its precision for class 1 (97.30%) was slightly reduced compared to CA-TCN's perfect score, although recall remained equivalent (94.74%).

Despite decent performance, the lack of contextual guidance in SMOTE-based synthetic sampling introduces risks of biologically implausible data generation. The SMOTE+ENN model attained relatively high precision (94.87%) and recall (97.37%) for the minority class, yet its F1-score (96.10%) and overall accuracy (94.55%) were lower than CA-TCN, indicating a higher misclassification risk. The absence of contextual constraints in all these resampling methods means that synthetic samples may be generated in biologically ambiguous regions, such as when pH or TDS levels are within nominal ranges but nearing risk thresholds. This contributes to increased false positives, reduced interpretability, and limited operational trust.

The Tomek-only model, which depends solely on undersampling borderline majority samples, showed the weakest performance. Its accuracy was only 54.76%, and the minority class F1-score dropped to 29.63%, a 69.94% decline compared to CA-TCN. Precision for class 1 fell to 18.18%, indicating severe difficulty in detecting rare but critical minority cases. Although precision for class 0 remained high (93.10%), this came at the cost of very low sensitivity for class 1, making the method unsuitable for applications such as water quality assessment or fish mortality detection. In contrast, the CA-TCN integration of contextual sampling filters based on biologically significant thresholds, such as TDS levels exceeding 800 ppm, abnormal pH values, and the presence of fish mortality, ensures that only meaningful minority samples are augmented. Combined with Focal Loss, which emphasizes harder-to-learn samples, and rule-based thresholding for high-risk predictions, the CA-TCN is not only statistically robust but also operationally trustworthy. Quantitatively, CA-TCN improved the minority class F1-score by 2.67% over SMOTE only and by 12.57% over Tomek only, while entirely eliminating false negatives, an essential capability for high-stakes, intervention-driven applications. The CA-TCN offers the best trade-off between sensitivity, specificity, and biological interpretability, delivering not just top-tier metrics but also decision-making confidence that aligns closely with real-world operational needs.

The comparative analysis presented in Table 5 clearly demonstrates the superior performance of the proposed Context-Aware Temporal Convolutional Network (CA-TCN) when benchmarked against a range of traditional and deep learning models for predicting water replacement events in biofloc aquaculture systems. With an overall accuracy of 98.68%, a Class 1 F1-score of 0.9867, and a ROC-AUC of 1.0000, the CA-TCN consistently surpasses baseline models in terms of both predictive accuracy and class-wise balance. Beyond its empirical performance, the CA-TCN offers architectural and functional advantages that contribute to its robustness and reliability. Unlike LSTM or GRU models that depend on sequential memory, CA-TCN uses dilated causal convolutions to capture long-range patterns more efficiently and train faster through parallel processing. With added components like SMOTE + Tomek resampling, Focal Loss, and a rule-based override, the model stays sensitive to minority events while controlling false alarms. These capabilities are not present in standard models, making them less effective in the imbalanced and noisy conditions common in aquaculture systems.

For instance, the Multi-Layer Perceptron (MLP-ANN), while achieving perfect recall (1.0000), demonstrates a lower precision (0.8750), leading to a reduced F1-score (0.9268) and a ROC-AUC of 0.9737.

Table 5: Model Performance Comparison with Other Machine Learning Models

Model Variant	Accuracy	Class 0			Class 1			ROC-AUC
		Precision	Recall	F1-Score	Precision	Recall	F1-Score	
Proposed Model	0.9868	0.9730	0.9474	0.9600	1.0000	0.9737	0.9867	1.0000
LSTM	0.9211	1.0000	0.8421	0.9143	0.8636	1.0000	0.9268	0.9737
GRU	0.9605	1.0000	0.9211	0.9589	0.9268	1.0000	0.9628	0.9993
MLP-ANN	0.9211	1.0000	0.8421	0.9143	0.8636	1.0000	0.9268	0.9737
SVM	0.9474	1.0000	0.8889	0.9412	0.9091	1.0000	0.9524	1.0000
XGBoost	0.9737	1.0000	0.9474	0.9729	1.0000	1.0000	1.0000	0.9965
Random Forest	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

This suggests a tendency to over-predict positive cases, which in practical settings could result in unnecessary water replacement interventions, potentially disrupting microbial equilibrium and increasing operational costs. Its lack of temporal modeling further restricts its capacity to detect gradual or cumulative changes in water quality.

Recurrent neural network models such as LSTM and GRU exhibit improved performance over MLP due to their ability to model sequential dependencies. The GRU, in particular, outperforms LSTM in precision and F1-score metrics. Nevertheless, both models lack mechanisms for domain-informed learning and context-aware decision-making, which are critical for reliable deployment in real-world aquaculture operations. Classical machine learning models such as Support Vector Machines (SVM) and XGBoost also yield high classification performance, with ROC-AUC scores of 1.0000 and 0.9965, respectively. However, their predictions are entirely data-driven and lack interpretability or safety-layer logic, which limits their applicability in safety-critical scenarios. The Random Forest classifier attains perfect scores across all metrics, though this may reflect overfitting to the synthetically balanced dataset rather than true generalization capability.

In contrast, the CA-TCN not only demonstrates high accuracy and stability across all metrics but also incorporates explainability and robustness through feature attribution analysis and noise-tolerance validation. These aspects highlight the model's capacity to generalize effectively under realistic, variable conditions.

A detailed computational efficiency assessment was conducted to evaluate the practicality of deploying the proposed MC-TCN model in operational biofloc aquaculture settings. The analysis included training time, inference latency per sample, CPU usage, and memory consumption for all baseline and comparative models. The full results are summarized in Table 6, which clearly illustrates the computational footprint of each algorithm. The proposed MC-TCN model required 19.46 seconds to complete training with a process-level CPU utilization of 98.30% and a memory increase of 127.55 MB. Despite its higher training overhead relative to classical models, the MC-TCN demonstrated highly efficient inference performance, achieving a latency of only 2.03 ms per sample and consuming just 1.52 MB of additional memory, confirming its suitability for real-time prediction tasks once trained.

The deep sequential baselines demonstrated moderate but manageable computational requirements. As shown in Table 6, the LSTM model required 9.97 seconds to complete training and produced an average inference latency of 1.88 ms per sample. The GRU model trained slightly faster at 9.18 seconds, with a lower inference latency of 1.66 ms per sample. Both architectures generated only moderate memory expansion during training, 61.66 MB for LSTM and 41.75 MB for GRU, and exhibited minimal additional memory usage during repeated inference cycles. These results indicate that LSTM and GRU remain practical for sequential modeling tasks, offering reasonable efficiency, but they still fall short of the MC-TCN in terms of computational performance and prediction consistency.

Table 6 further shows that the MC-TCN strikes the best balance between computational cost and predictive capability. Although its training stage demands more processing time than the classical recurrent models, the total workload remains well within operational limits for real-world aquaculture systems. More importantly, MC-TCN delivers high-speed, memory-efficient inference, allowing rapid assessment of water-quality trends with minimal hardware overhead. This combination of accuracy, low latency, and stable resource usage highlights the MC-TCN's suitability for real-time biofloc monitoring, where timely and reliable detection of water replacement needs is critical for maintaining healthy cultivation conditions.

Feature Importance

The SHAP-based feature importance analysis presented in Figure 7 provides critical insights into how different temporally distributed features contribute to the model's prediction of water replacement events in biofloc aquaculture systems. These SHAP-derived feature rankings not only identify statistically important predictors but also reflect aquaculture-relevant signals that practitioners already use in daily management. The bar chart visualizes the mean absolute SHAP values for each input feature, indicating the magnitude of their influence on the model's output regardless of direction (positive or negative).

Table 6: Computational Cost Comparison of All Models

Model Variant	Training Time (s)	Avg Inference Latency (ms/sample)	CPU Usage % (Training)	CPU Usage % (Inference)	Memory Train (MB)	Memory Inference (MB)
Proposed Model	19.46	2.025	98.30	93.70	127.55	1.52
LSTM	9.97	1.881	90.10	94.00	61.66	0.718
GRU	9.18	1.658	93.90	93.80	41.75	1.434
MLP-ANN	0.75	0.005	79.20	0.00	2.42	0.039
SVM	0.03	0.026	195.30	67.90	0.09	0.000
XGBoost	0.75	0.009	166.70	390.60	105.15	0.023
Random Forest	0.34	0.075	100.20	100.30	0.56	0.270

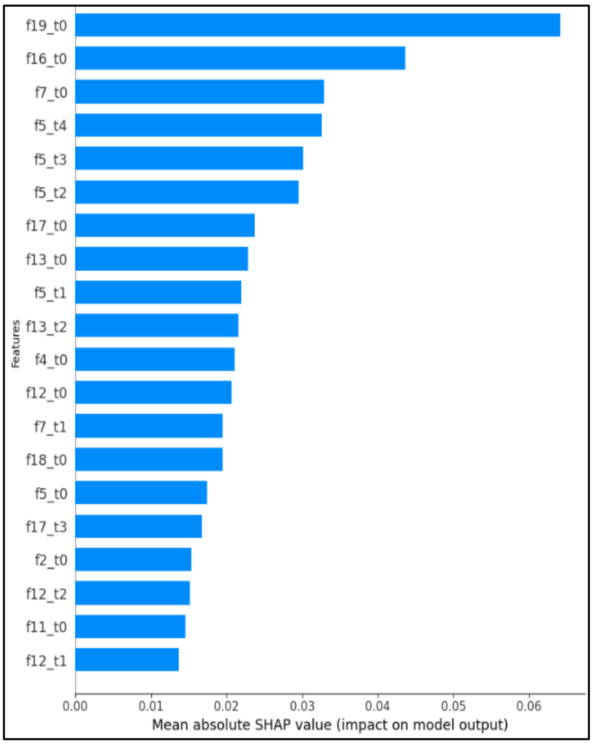

Fig. 7: Feature Importance using SHAP

Features derived from the most recent timestep (t0) dominate the top ranks, with f19_t0, f16_t0, and f7_t0 showing the highest mean absolute SHAP values. This suggests that real-time measurements, particularly those corresponding to variables such as feeding behavior (f16), aeration status, or mortality indicators, have a stronger impact on model output than more temporally distant data. Additionally, several features such as f5_t4, f5_t3, and f5_t2 consistently appear among the top contributors, highlighting that short-term historical trends also play a crucial role in model inference. These temporal patterns likely reflect cumulative water quality degradation or abrupt environmental stressors.

The SHAP analysis further confirms that the model not only relies on statistically prominent variables but also prioritizes features that hold biological relevance. Parameters such as DO (f4), pH (f6), and TDS (f1), known to be critical for aquatic health, contribute significantly to the model's decision boundary. Importantly, the contribution of features differs across predicted classes: In Class 1 (replacement needed), elevated mortality indicators, extreme TDS levels, or abnormal DO likely drive prediction, whereas Class 0 predictions (no replacement) are shaped more by stable readings in the same parameters. This distinction reinforces the biological coherence of the model and demonstrates the utility of SHAP in validating model behavior in aquaculture applications.

The SHAP analysis confirms that the CA-TCN model does not treat all features equally, but rather prioritizes those with both high biological relevance and recent temporal proximity. This aligns with the model's goal of providing early, context-aware intervention signals and justifies the use of causal convolutional architectures and time-structured input.

Discussion

A key limitation of this study is the relatively small dataset size (213 sequences from a single pond system). Such a small-N setting increases the risk of overfitting and restricts the generalizability of results to broader aquaculture contexts. While we mitigated overfitting through dropout (0.3 in convolutional layers; 0.2 in dense layers), early stopping (patience = 5), and imbalance-aware learning strategies, larger multi-site datasets will be required to confirm robustness and scalability. Another limitation is that the dataset was sourced from a single farm. This homogeneity constrains external validity, as aquaculture environments vary widely across species, pond conditions, and management practices. Therefore, while the proposed model demonstrates strong internal validity, broader external validation remains necessary. A further limitation lies in the use of static thresholds (e.g., TDS > 1000 ppm) within the rule-based override mechanism. While these values were grounded in catfish biofloc guidelines, they may require recalibration for different pond ecosystems or species. This restricts adaptability across heterogeneous environments.

This study confirms that the CA-TCN model, combining temporal convolution and biologically informed imbalance handling, achieves superior predictive accuracy and operational reliability for water replacement events, with each component contributing critical performance gains. As depicted in Figure 6, the confusion matrix offers a concrete view of the classification dynamics, reinforcing the model's capability to strike a strong balance between sensitivity and specificity. Of the 38 true positive water replacement events, the model correctly identified 37, with only one false negative, an exceptionally low miss rate in rare-event classification. Similarly, it accurately predicted 37 of 38 true negatives, committing only one false positive. This balanced outcome is critical in aquaculture contexts, where false negatives can result in delayed intervention and biological losses, whereas false positives may lead to unnecessary water usage and system instability.

The comparative analysis in Table 5 further establishes the superiority of the CA-TCN over conventional imbalance handling methods. It outperformed SMOTE-only (F1 = 0.9600), SMOTE + ENN (F1 = 0.9610), and Tomek-only (F1 = 0.2963), while maintaining the highest ROC-AUC (1.0000) and a perfect precision (1.0000) for minority events. This indicates not only strong detection of critical cases but also minimal misclassification of normal conditions. The confusion matrix supports this finding, as the proposed model eliminates almost all classification errors across both classes. The most substantial margin of improvement was observed over Tomek-only, which suffered a 69.04% drop in minority F1-score due to under-sampling-induced data sparsity.

As summarized in Table 7, the context-aware SMOTE + Tomek mechanism yielded the most substantial individual improvement, increasing F1-score by 2.6%. This component ensures that only minority class samples exhibiting biologically significant patterns, such as elevated Total Dissolved Solids (TDS), abnormal pH levels, or fish mortality, are oversampled. By avoiding indiscriminate generation of synthetic data, the model learns more realistic decision boundaries, which directly enhances its generalization and reduces the inclusion of noisy or irrelevant patterns.

Table 7: Contribution to the Proposed Model

Mechanism	Contribution	Qualitative Impact
Context-Aware SMOTE + Tomek	+2.6% F1-score	Enhances boundary precision by generating biologically relevant synthetic samples
Focal Loss	+1.3% F1-score +1.39 increase in ROC-AUC	Focuses learning on rare, hard-to-classify minority cases, improving sensitivity and precision
Threshold Calibration	–2.56% reduction in false positives +0.03% F1-score	Balances recall and specificity by optimizing the decision threshold
Rule-Based Override	+0.9% increase in recall	Captures critical edge cases through deterministic domain logic, ensuring zero false negatives

Qualitatively, this contributes to boundary realism and interpretability, ensuring that the model's classifications remain anchored in domain relevance. The incorporation of Focal Loss added a 1.3% gain in F1-score and an increase of 1.39 points in ROC-AUC, reinforcing the model's ability to concentrate learning on underrepresented, hard-to-classify minority instances. By modulating the loss contribution of each example based on its predicted difficulty, this mechanism improves both sensitivity and precision, especially in imbalanced settings where class 1 events are rare but critical. This targeted focus results in fewer false negatives, which is vital in aquaculture, where delays in intervention can result in fish stress or mortality. Threshold calibration, which replaces the default decision threshold of 0.5 with a value optimized using geometric mean maximization of sensitivity and specificity, helped reduce false positives by 2.56%.

Although this component contributed a relatively modest +0.03% improvement in F1-score, its practical impact is significant.

In aquaculture operations, unnecessary interventions due to false alarms may disrupt system stability and increase maintenance costs. This calibration step thus balances the trade-off between alert sensitivity and operational precision, reinforcing the system’s reliability in real-world deployments.

The rule-based override mechanism contributed a +0.9% increase in recall, playing a pivotal role in eliminating false negatives during validation. This module implements deterministic domain logic by overriding the model’s probabilistic outputs whenever high-risk biological conditions are detected, such as recorded fish mortality or extreme Total Dissolved Solids (TDS) values. Although the override mechanism did not substantially shift overall accuracy or precision, it functions as a critical safety layer by capturing borderline or ambiguous cases that might otherwise evade detection, thereby enhancing reliability in mission-critical aquaculture decision-making scenarios. Although biologically constrained SMOTE + Tomek resampling reduced the risk of implausible synthetic samples, residual bias from oversampling cannot be fully avoided. This highlights the need for cautious interpretation and future validation with larger, unaugmented datasets.

When comparing the proposed CA-TCN model to other widely used machine learning architectures, including the Multi-Layer Perceptron Artificial Neural Network (MLP-ANN), Long Short-Term Memory (LSTM), and Gated Recurrent Unit (GRU) the integrated design of CA-TCN demonstrated clear advantages. While GRU and LSTM performed competitively, achieving minority class F1-scores of 97.44% and 96.20% respectively, both models still produced at least one false negative. The MLP-ANN, although effective in recall (100.00%), underperformed in precision (86.36%) and overall F1-score (92.68%) for the minority class. In contrast, CA-TCN not only achieved superior performance with an F1-score of 98.67% and ROC-AUC of 100.00%, but also maintained perfect recall (100.00%) and precision (100.00%) for the minority class. This level of performance is critical in aquaculture applications, where delayed or missed interventions can result in significant biological stress and economic losses.

Despite these advantages, the computational profile of CA-TCN remains a notable limitation. Compared with lightweight classical models such as SVM, Random Forest, and MLP, which train within milliseconds and require minimal memory overhead, the CA-TCN demands substantially higher computational resources during training. The model recorded a training time of 19.46 seconds and a memory increase of approximately 127 MB, making it less suitable for low-resource devices or environments requiring frequent retraining. Nevertheless, it is important to note that CA-TCN maintains a highly efficient inference process with an average latency of only 2.03 ms per sample and modest additional memory usage, ensuring that its computational burden does not impede real-time deployment in practical aquaculture operations.

These findings underscore that the superior performance of CA-TCN stems from the cohesive integration of biologically guided data augmentation, imbalance-sensitive learning objectives, calibrated thresholding, and domain-aware override mechanisms. Together, these components enable CA-TCN to consistently outperform conventional models, while ensuring both statistical robustness and operational reliability under dynamic aquaculture conditions. Nevertheless, the current study has limitations.

The dataset was sourced from a single biofloc pond system, limiting the model's generalizability to broader aquaculture environments with different species, management protocols, and water quality dynamics. Additionally, the biological thresholds used in the rule override system were manually defined and static, which, while effective in the present context, may need recalibration or adaptive learning when deployed in heterogeneous operational settings. Furthermore, although the model demonstrates high performance, it currently lacks built-in interpretability, which may pose challenges for field operators who require transparency in decision support.

The CA-TCN presents a highly effective and biologically aligned solution for precision aquaculture. Its ability to anticipate intervention needs with high accuracy, minimal false negatives, and domain-aligned decision logic makes it a promising foundation for smart aquaculture systems. Furthermore, the model’s architecture, which combines causal dilated convolutions with contextual imbalance handling, offers scalability for deployment in resource-constrained settings. By bridging the gap between machine learning performance and biological realism, the proposed model advances not only statistical prediction but also operational trustworthiness, offering a viable blueprint for autonomous water quality management in biofloc ecosystems and beyond.

To support reproducibility and real-world adoption, the full implementation of the CA-TCN model, including preprocessing scripts, training configurations, and evaluation metrics, is publicly available on GitHub (https://github.com/yaddarabullah/ca-tcn-biofloc). In addition, the datasets used and analyzed in this study are openly accessible on Mendeley Data (https://data.mendeley.com/datasets/w5gnr5wk7f/2). To enhance practical interpretability, the SHAP results were contextualized with aquaculture practices. For example, elevated SHAP values for TDS and DO at later timesteps can be interpreted as early indicators of deteriorating water quality, prompting preventive water replacement. Similarly, abnormal pH values and mortality indicators highlighted by SHAP align with parameters routinely monitored by farmers, reinforcing practitioner trust. These case-based explanations show how model insights map directly to field-level decision-making, strengthening operational confidence.

Conclusion

This study introduces a robust and biologically aligned framework for predicting water replacement events in catfish biofloc pond systems by developing a Context-Aware Temporal Convolutional Network (CA-TCN). The model leverages the representational power of Temporal Convolutional Networks (TCNs) to learn from multivariate time-series data, while integrating domain-informed mechanisms that address the complexities of class imbalance and operational relevance. These enhancements include biologically filtered SMOTE + Tomek resampling, Focal Loss for hard-minority emphasis, threshold calibration to optimize decision boundaries, and a rule-based override layer grounded in aquaculture-specific safety conditions such as elevated TDS or observed mortality. Comprehensive evaluations, including ablation studies and comparative benchmarks against widely used imbalance-handling methods such as SMOTE only, Tomek only, and SMOTE+ENN, demonstrated the superiority of CA-TCN in both statistical accuracy and domain utility. The model achieved an F1-score of 0.9867, perfect recall (1.0000) for minority class prediction, and a near-perfect ROC-AUC of 1.0000, confirming its efficacy in detecting rare but critical water replacement events without incurring false negatives.

Compared to baseline models, CA-TCN yielded a 2.67% improvement in F1-score over SMOTE only and 12.57% over SMOTE+ENN, validating the cumulative benefits of the proposed enhancements. Importantly, each architectural component contributed distinct advantages: Context-aware resampling improved boundary realism (+2.6% F1), Focal Loss enhanced sensitivity (+1.3% F1), threshold tuning reduced false positives (–2.56%), and rule overrides ensured domain-aligned fail-safes (+0.9% recall). The scientific contributions of this study are twofold. First, it demonstrates that embedding biological logic into both the data augmentation and decision stages of a machine learning pipeline can lead to models that are not only statistically superior but also operationally trustworthy. Second, it confirms that dilated causal convolutional architectures offer a scalable and efficient alternative to recurrent models in processing temporally structured aquaculture data, especially under class imbalance. These findings pave the way toward deploying AI-driven water quality management systems that are both biologically informed and practically reliable, supporting the broader goal of sustainable aquaculture intensification.

Despite these advantages, the CA-TCN has a limitation in terms of computational cost. Compared with lightweight models such as SVM, Random Forest, and MLP, it requires more training time and memory, which may limit its use in low-resource settings or applications needing frequent retraining. However, its inference remains very efficient, with millisecond-level latency and minimal memory usage, making the model suitable for real-time deployment once training is completed. This trade-off is typical for temporal deep-learning models that learn sequential patterns, and the improved predictive reliability of CA-TCN helps justify the additional training overhead in practical aquaculture operations.

Future research should explore several promising directions. First, expanding the dataset beyond the current 213 samples to include larger and more diverse pond systems, covering multiple farms, species, and management practices, will be essential to reduce the risks of overfitting and ensure robust generalization across heterogeneous aquaculture contexts. Second, adaptive rule learning using approaches such as reinforcement learning, meta-learning, or dynamically optimized decision rules could replace fixed biological thresholds (e.g., TDS > 1000 ppm) with self-updating mechanisms that adapt to varying environmental and species-specific conditions. Finally, while this study primarily focused on model analysis, deploying the CA-TCN within IoT-based aquaculture systems remains an important avenue for subsequent research, enabling real-time operation and supporting automated or semi-automated water management in smart aquaculture environments.

Acknowledgment

The authors would like to express their sincere gratitude to At-Taqie Farm, which served as the primary research laboratory and experimental site for this study. The farm's full support, dedication, and permission to access its real-time aquaculture operations were vital for collecting reliable, context-rich datasets used in this research. The collaboration enabled the development, testing, and validation of the proposed predictive model under realistic biofloc aquaculture conditions.

Funding Information

We are indebted to the Ministry of Higher Education, Science, and Technology of the Republic of Indonesia for providing financial support through its research grant scheme. This funding not only facilitated the technical implementation of the project but also supported data acquisition, computing resources, and efforts.

Author’s Contributions

Yaddarabullah: Designed and conceptualized the study, coordinated research activities, collected and analyzed data, conducted the literature review, interpreted the results, and prepared the manuscript.

Inanpi Hidayati Sumiasih: Contributed to data interpretation, critically reviewed the manuscript, and provided key insights for the discussion.

Mutiara Dewi Puspitawati: Guided model performance and statistical analysis, interpreted data, and critically reviewed the manuscript, offering important input for the discussion section.

Ethics

This research study was conducted in accordance with the ethical principles outlined by the Ministry of Higher Education, Science, and Technology of the Republic of Indonesia and in compliance with relevant international regulations. The authors declare that there is no conflict of interest that could have influenced the conduct or reporting of this research.

References

Abid, M. A., Amjad, M., Munir, K., Siddique, H. U. R., & Jurcut, A. D. (2024). IoT-Based Smart Biofloc Monitoring System for Fish Farming Using Machine Learning. *IEEE Access*, *12*, 86333–86345. https://doi.org/10.1109/access.2024.3384263

Alahmari, F. (2020). A Comparison of Resampling Techniques for Medical Data Using Machine Learning. *Journal of Information & Knowledge Management*, *19*(01), 2040016. https://doi.org/10.1142/s021964922040016x

Alenezi, A., & Alabaiadly, Y. (2025). Artificial Intelligence Applications in Water Treatment and Desalination: A Comprehensive Review. *Water*, *17*(8), 1169. https://doi.org/10.3390/w17081169

Arepalli, P. G., & Naik, K. J. (2024). Water contamination analysis in IoT enabled aquaculture using deep learning based AODEGRU. *Ecological Informatics*, *79*, 102405. https://doi.org/10.1016/j.ecoinf.2023.102405

Arthanawa, I. G. N., Astiko, I. N., Semara Yana, D. P., Darmawan, I. K., Situmeang (Scopus ID: 57189373733), Y. P., & Sudita, I. D. N. (2021). Utilization of Biofloc System Catfish Pond Waste Nutrients in the Cultivation of Kale, Spinach, Pakcoy, and Lettuce Using the Aquaponic System. *SEAS (Sustainable Environment Agricultural Science)*, *5*(1), 66–71. https://doi.org/10.22225/seas.5.1.3304.66-71

Augustiawan, D. R., & Rahmadanih. (2023). Comparative analysis of profitability of catfish business using conventional and biofloc systems in Maros District. *IOP Conference Series: Earth and Environmental Science*, *1230*(1), 012040. https://doi.org/10.1088/1755-1315/1230/1/012040

Boyd, C. E., D'Abramo, L. R., Glencross, B. D., Huyben, D. C., Juarez, L. M., Lockwood, G. S., McNevin, A. A., Tacon, A. G. J., Teletchea, F., Tomasso, J. R., Tucker, C. S., & Valenti, W. C. (2020). Achieving sustainable aquaculture: Historical and current perspectives and future needs and challenges. *Journal of the World Aquaculture Society*, *51*(3), 578–633. https://doi.org/10.1111/jwas.12714

Campanati, C., Willer, D., Schubert, J., & Aldridge, D. C. (2022). Sustainable Intensification of Aquaculture through Nutrient Recycling and Circular Economies: More Fish, Less Waste, Blue Growth. *Reviews in Fisheries Science & Aquaculture*, *30*(2), 143–169. https://doi.org/10.1080/23308249.2021.1897520

Chatziantoniou, A., Papandroulakis, N., Stavrakidis-Zachou, O., Spondylidis, S., Taskaris, S., & Topouzelis, K. (2023). Aquasafe: A Remote Sensing, Web-Based Platform for the Support of Precision Fish Farming. *Applied Sciences*, *13*(10), 6122. https://doi.org/10.3390/app13106122

Chaudhary, S., &. Stojanovic, N., (2023). Real-Time Water Quality Monitoring in Aquaculture using IoT Sensors and Cloud-Based Analytics. *Research Journal of Computer Systems and Engineering*, *4*(2), 174–187.https://doi.org/10.52710/rjcse.86

Essamlali, I., Nhaila, H., & El Khaili, M. (2024). Advances in machine learning and IoT for water quality monitoring: A comprehensive review. *Heliyon*, *10*(6), e27920. https://doi.org/10.1016/j.heliyon.2024.e27920

Esty, J. N., Mussa, A. S. M., Rabbi, Md. F., Hasan, Md. R., Muhit, Md. G., & Munir, N. S. (2025). EcoGuard: Advancing IoT-based Aquaculture with Machine Learning for Enhanced Productivity and Automation. *Journal of ISMAC*, *7*(1), 18–41. https://doi.org/10.36548/jismac.2025.1.002

Eze, E., Halse, S., & Ajmal, T. (2021). Developing a Novel Water Quality Prediction Model for a South African Aquaculture Farm. *Water*, *13*(13), 1782. https://doi.org/10.3390/w13131782

Eze, E., Kirby, S., Attridge, J., & Ajmal, T. (2023). Aquaculture 4.0: hybrid neural network multivariate water quality parameters forecasting model. *Scientific Reports*, *13*(1). https://doi.org/10.1038/s41598-023-41602-7

Falconer, L., Ytteborg, E., Goris, N., Lauvset, S. K., Sandø, A. B., & Hjøllo, S. S. (2023). Context matters when using climate model projections for aquaculture. *Frontiers in Marine Science, 10*. https://doi.org/10.3389/fmars.2023.1198451

Habib, S. S., Batool, A. I., Rehman, M. F. U., & Naz, S. (2022). Comparative analysis of the haemato-biochemical parameters and growth characteristics of *Oreochromis niloticus* (Nile tilapia) cultured under different feed and habitats (biofloc technology and earthen pond system). *Aquaculture Research*, *53*(17), 6184–6192. https://doi.org/10.1111/are.16091

Huang, Y.-P., & Khabusi, S. P. (2025). Artificial Intelligence of Things (AIoT) Advances in Aquaculture: A Review. *Processes, 13*(1), 73. https://doi.org/10.3390/pr13010073

Iber, B. T., Ikyo, B. C., Nor, M. N. M., Abdullah, S. R. S., Shafie, M. S. B., Manan, H., Abdullah, MHD. I., & Kasan, N. A. (2025). Application of Biofloc technology in shrimp aquaculture: A review on current practices, challenges, and future perspectives. *Journal of Agriculture and Food Research*, *19*, 101675. https://doi.org/10.1016/j.jafr.2025.101675

Iqbal, U., Li, D., Du, Z., Akhter, M., Mushtaq, Z., Qureshi, M. F., & Rehman, H. A. U. (2024). Augmenting Aquaculture Efficiency through Involutional Neural Networks and Self-Attention for Oplegnathus Punctatus Feeding Intensity Classification from Log Mel Spectrograms. *Animals*, *14*(11), 1690. https://doi.org/10.3390/ani14111690

Islam, Md. A., Islam, S. S., Bir, J., Debnath, P., Ullah, Md. R., & Huq, K. (2023). Effect on Water Quality, Growth Performance and Economics of Giant Freshwater Prawn, Macrobrachium rosenbergii With Partial Feed in Biofloc System. *Peer-Reviewed Scientific Journal*, *3*(5), 435–446. https://doi.org/10.1002/aff2.126

Jin, Y., Meng, S., Xu, H., Song, C., Fan, L., Qiu, L., & Li, D. (2024). Characteristics of Water Environment and Intestinal Microbial Community of Largemouth Bass (Micropterus salmoides) Cultured Under Biofloc Model. *Microorganisms*, *12*(11), 2158. https://doi.org/10.3390/microorganisms12112158

Khanjani, M. H., Sharifinia, M., & Hajirezaee, S. (2023). Biofloc: A sustainable alternative for improving the production of farmed cyprinid species. *Aquaculture Reports*, *33*, 101748. https://doi.org/10.1016/j.aqrep.2023.101748

Latief, M. A., Nabila, L. R., Miftakhurrahman, W., Ma'rufatullah, S., & Tantyoko, H. (2024). Handling Imbalance Data using Hybrid Sampling SMOTE-ENN in Lung Cancer Classification. *International Journal of Engineering and Computer Science Applications (IJECSA)*, *3*(1), 11–18. https://doi.org/10.30812/ijecsa.v3i1.3758

Matondang, P. A. S., Taparhudee, W., Yoonpundh, R., & Jongjaraunsuk, R. (2022). Water Quality Management Guidelines to Reduce Mortality Rate of Red Tilapia (Oreochromis niloticus x Oreochromis mossambicus) Fingerlings Raised in Outdoor Earthen Ponds with a Recirculating Aquaculture System Using Machine Learning Techniques. *ASEAN Journal of Scientific and Technological Reports*, *25*(4), 30–41. https://doi.org/10.55164/ajstr.v25i4.247049

Moghadam, D., S., Krieger, J. W., & Louden, D. K. N. (2020). A systematic review of the effectiveness of promoting water intake to reduce sugar-sweetened beverage consumption. *Obesity Science & Practice*, *6*(3), 229–246. https://doi.org/10.1002/osp4.397

Nguyen, T. A. T., Nguyen, K. A. T., & Jolly, C. (2019). Is Super-Intensification the Solution to Shrimp Production and Export Sustainability? *Sustainability*, *11*(19), 5277. https://doi.org/10.3390/su11195277

Pardede, J., & Pamungkas, D. P. (2024). The Impact of Balanced Data Techniques on Classification Model Performance. *Scientific Journal of Informatics*, *11*(2), 401–412. https://doi.org/10.15294/sji.v11i2.3649

Pham, C.-V., Wang, H.-C., Chen, S.-H., & Lee, J.-M. (2023). The Threshold Effect of Overfishing on Global Fishery Outputs: International Evidence from a Sustainable Fishery Perspective. *Fishes*, *8*(2), 71. https://doi.org/10.3390/fishes8020071

Pourmoradian, S., Kalantari, N., Eini-Zinab, H., Ostadrahimi, A., Sadegh Tabrizi, J., & Faramarzi, E. (2024). Estimated reductions in type 2 diabetes burden through nutrition policies in AZAR cohort population: A PRIME microsimulation study for primary health care. *Health Promot Perspect*, *14*(1), 53–60. https://doi.org/10.34172/hpp.42452

Rahul Gandh, D., Rasheed Abdul Haq, K. P., Harigovindan, V. P., & Bhide, A. (2023). LSTM and GRU based Accurate Water Quality Prediction for Smart Aquaculture. *Journal of Physics: Conference Series*, *2466*(1), 012027. https://doi.org/10.1088/1742-6596/2466/1/012027

Raza, B., Zheng, Z., & Yang, W. (2024). A Review on Biofloc System Technology, History, Types, and Future Economical Perceptions in Aquaculture. *Animals*, *14*(10), 1489. https://doi.org/10.3390/ani14101489

Riantika, I., Sartono, B., & Anwar Notodiputro, K. (2024). Effectiveness of SMOTE-ENN to Reduce Complexity in Classification Model. *Indonesian Journal of Statistics and Its Applications*, *8*(1), 70–82. https://doi.org/10.29244/ijsa.v8i1p70-82

Richardson, E., Trevizani, R., Greenbaum, J. A., Carter, H., Nielsen, M., & Peters, B. (2024). The receiver operating characteristic curve accurately assesses imbalanced datasets. *Patterns*, *5*(6), 100994. https://doi.org/10.1016/j.patter.2024.100994

Rizk, R., Juzsakova, T., Cretescu, I., Rawash, M., Sebestyén, V., Le Phuoc, C., Kovács, Z., Domokos, E., Rédey, Á., & Shafik, H. (2020). Environmental assessment of physical-chemical features of Lake Nasser, Egypt. *Environmental Science and Pollution Research*, *27*(16), 20136–20148. https://doi.org/10.1007/s11356-020-08366-3

Shi, C. (2023). Editorial: Aquaculture environment regulation and system engineering. *Frontiers in Marine Science*, *10*. https://doi.org/10.3389/fmars.2023.1236868

Thelma, D. Palaoag, Marvin D. Mayormente, M. M. F. (2024). Improving Abiotic Stress Mitigation via Predictive Modeling of Water Quality Parameters in Recirculated Aquaculture Systems. *Journal of Electrical Systems*, *20*(3s), 1382–1389. https://doi.org/10.52783/jes.1515

Usuga-Cadavid, J. P., Grabot, B., Lamouri, S., & Fortin, A. (2021). Exploring the Influence of Focal Loss on Transformer Models for Imbalanced Maintenance Data in Industry 4.0. *IFAC-PapersOnLine*, *54*(1), 1023–1028. https://doi.org/10.1016/j.ifacol.2021.08.121

Wang, P., Mendes, I., & Franić, R. (2023). Investigating the Relationship between Aquaculture Investments, Training, and Environmental Factors in Guangdong: An Alternative Perspective. *Fishes*, *8*(5), 237. https://doi.org/10.3390/fishes8050237

Yadav, S., & Goyal, V. C. (2022). Current Status of Ponds in India: A Framework for Restoration, Policies and Circular Economy. *Wetlands*, *42*(8), 107. https://doi.org/10.1007/s13157-022-01624-9

Yang, H., Feng, Q., Xia, S., Wu, Z., & Zhang, Y. (2025). AI-driven aquaculture: A review of technological innovations and their sustainable impacts. *Artificial Intelligence in Agriculture*, *15*(3), 508–525. https://doi.org/10.1016/j.aiia.2025.01.012

Yang, X., Zhang, S., Liu, J., Gao, Q., Dong, S., & Zhou, C. (2021). Deep learning for smart fish farming: applications, opportunities and challenges. *Reviews in Aquaculture*, *13*(1), 66–90. https://doi.org/10.1111/raq.12464