

Enhancing Water Quality Prediction using Adaptive Neuro-Fuzzy Inference System with Neural Network Integration for Operational Safety

Ramya S¹ and Srinath S²

^{1,2}Department of Computer Science & Engineering, JSS Science and Technology University, Mysuru, India
Email: ramya.shivanagu@gmail.com, srinath@sjce.ac.in

Abstract—Accurately predicting water quality parameters plays a crucial role in managing wastewater treatment systems, helping to safeguard public health and ensure a safer working environment for plant operators. In this study, we introduce a hybrid model that merges the strengths of Adaptive Neuro-Fuzzy Inference Systems (ANFIS) with Artificial Neural Networks (ANN) to enhance the prediction of Biochemical Oxygen Demand (BOD)—a key indicator of water quality. By combining ANFIS's ability to handle fuzzy, imprecise data with the learning power of ANN, the model is well-equipped to deal with the inherent uncertainties and nonlinear behavior of wastewater data. The model was developed and tested using daily data from an industrial sewage treatment plant (STP), capturing the dynamic and often unpredictable nature of real-world operations. Its performance was assessed using a range of evaluation metrics including MSE, RMSE, MAE, SSE, and the R² score. Results were impressive, with R² values reaching 0.9854 for training, 0.9850 for validation, and 0.9718 for testing. These outcomes underscore the model's strong predictive capability and highlight its potential to support more efficient, reliable, and cost-effective wastewater management practices.

Keyword— Water Quality, sewage treatment plant, Adaptive Neuro-Fuzzy Inference System, Artificial Neural Network.

I. INTRODUCTION

Wastewater treatment plants (WWTPs) play a vital role in safeguarding both the environment and public health. By cleaning wastewater before it's released, these facilities help prevent the contamination of rivers, lakes, and other natural water bodies. Without them, harmful pollutants—ranging from organic waste and toxic chemicals to disease-causing microorganisms—would flow unchecked into ecosystems, damaging aquatic life and endangering nearby communities [1]. The absence of proper treatment infrastructure not only increases environmental risks but also worsens water scarcity and affects overall quality of life, especially in densely populated or resource-limited areas. This research centers on predicting Biochemical Oxygen Demand (BOD), a critical indicator used to evaluate water quality. BOD reflects the amount of oxygen needed by microorganisms to break down organic material in the water [2]. Although lab-based methods for measuring BOD are accurate, they're also slow—often taking up to five days—delaying important operational decisions. Real-time sensors offer a faster solution, but their high cost can be prohibitive for many plants. These challenges underscore the urgent need for faster, more affordable, and reliable approaches to estimating BOD levels.

Machine learning (ML) offers a smart, data-driven way to predict Biochemical Oxygen Demand (BOD), especially when paired with historical records. These methods dramatically cut down on both cost and time, without sacrificing accuracy. Unlike traditional lab testing or high-end sensor systems, ML models can deliver fast and reliable predictions—making them a practical solution for water treatment facilities. By providing near real-time insights, they empower operators to make quicker, more informed decisions that improve safety and efficiency [3]. Over the years, a wide range of ML algorithms have been successfully applied to water quality prediction. Techniques like Random Forest, Support Vector Regression (SVR), Gradient Boosting, and Artificial Neural Networks (ANN) have shown strong performance in handling the complex and often nonlinear patterns found in environmental datasets. Past studies have used these models to accurately estimate parameters such as pH, Chemical Oxygen Demand (COD), and Total Suspended Solids (TSS), further proving their value in this field.

In this study, we take things a step further by developing a hybrid model that combines the Adaptive Neuro-Fuzzy Inference System (ANFIS) with Neural Networks to enhance the accuracy of BOD predictions. This fusion leverages ANFIS’s ability to manage fuzzy, imprecise data along with the deep learning strengths of Neural Networks [4]. The result is a robust model capable of navigating the uncertainties and nonlinear dynamics typical of wastewater data. By delivering high accuracy at a lower cost, this approach represents a significant advancement in real-time water quality monitoring—offering WWTPs a dependable and scalable solution for ensuring operational safety and efficiency.

II. METHOD AND METHODOLOGY

A. Data Collection

The dataset is collected daily from an industrial STP, using well-equipped laboratory instruments for accurate and reliable measurements. This data reflects the dynamic nature of the wastewater treatment process and captures daily variations in wastewater quality and treatment performance. Key parameters measured include various influent and effluent characteristics, along with key operational and biological indicators (Table 1).

TABLE I: STATISTICS OF COLLECTED DATA FROM STP

Feature	Mean	Std	Min	25%	50%	75%	Max
PHIN	9.00	2.28	5.34	6.44	9.11	10.95	12.8
CODIN	1360	638	448	997	1232	1440	3840
BODIN	841	313	398	564	798	1117	1407
ConductivityIN	2412	2624	957	1229	1484	1933	14211
TSSIN	425	353	110	233	330	462	2012
AmoniaNitrogenIn	82.75	57.94	22	44	56	95	207
PHAR	6.92	0.22	6.39	6.77	6.95	7.1	7.45
MLSSAR	1510	366	796	1216	1500	1808	2320
MLVSSAR	1405	234	990	1239	1371	1614	1832
SVAR	597	153	250	480	550	750	880
DOAR	2.65	1.82	0.18	0.82	2.84	3.95	7.38
PHCO	7.48	0.36	6.68	7.29	7.46	7.64	8.98
BODCO	12	6	1	7	11	18	25
CODCO	49	32	16	32	40	56	216
TSSCO	33	6	21	30	32	35	56
PHUFI	7.52	0.26	6.98	7.38	7.53	7.67	8.08
CODUFI	36	16	16	24	32	40	88
TSSUFI	21	3	12	19	21	24	32
TurbidityUFI	5.18	2.86	0.84	3.53	4.5	5.8	15.9
BODOUT	2.04	2.33	0.2	0.4	0.94	3	8

The table showcases descriptive statistics for essential wastewater treatment parameters, emphasizing their role in evaluating treatment efficiency. Parameters like PHIN (average: 9.00), CODIN (average: 1360.68 mg/L), and BODIN (average: 841.62 mg/L) capture the chemical and biological composition of incoming wastewater, which is critical for designing effective treatment strategies. ConductivityIN (average: 2412.38 $\mu\text{S/cm}$) reflects the level of dissolved ions, potentially indicating industrial or saline influences, while TSSIN (average: 425.89 mg/L) highlights the suspended solids requiring filtration. AmoniaNitrogenIn (average: 82.76 mg/L) reveals nitrogen contamination, demanding close monitoring to prevent ecological issues like eutrophication. Post-treatment pH levels (PHAR, average: 6.93) confirm the effluent's stability and suitability for discharge. Monitoring these parameters is vital for optimizing treatment processes, meeting environmental regulations, and protecting aquatic ecosystems, ensuring sustainable wastewater management practices.

B. Adaptive Neuro-Fuzzy Inference System with ANN

The hybrid ANFIS-ANN model brings together the best of both worlds—fuzzy logic and neural networks—to tackle the complexity and uncertainty often found in real-world systems. ANFIS contributes by applying fuzzy if-then rules that can manage vague or imprecise information, offering a more human-like way of reasoning through uncertainty [5]. Meanwhile, the ANN works behind the scenes to recognize patterns and make data-driven predictions using its layered learning structure. In this approach, both ANFIS and ANN receive the same input data but process it independently: ANFIS evaluates fuzzy membership values based on its rule base, while the ANN learns from the data directly. Their individual outputs are then averaged to produce the final prediction. By combining these two methods, the model benefits from the adaptability and learning power of neural networks while preserving the transparency and interpretability offered by fuzzy logic [6]. This synergy not only improves prediction accuracy but also provides insights into the underlying relationships in the data (Figure 1).




Figure 1: Model Architecture for hybrid model ANFIS-ANN

The ANFIS architecture is built as a layered system, where each layer plays a distinct role in transforming raw input data into a meaningful prediction. It all begins with the Input Layer, which simply receives the features from the dataset—each input node corresponds to one of the variables being analyzed. Next, the Membership Function Layer takes over, converting these crisp input values into fuzzy sets using predefined membership functions. This step essentially adds a layer of human-like reasoning by capturing degrees of belonging, rather than strict classifications [7]. For an input variable x_i and a specific membership function j , the output is expressed as $\mu_{ij}(x_i)$, which reflects how well the input fits into a particular fuzzy category. The Rules Layer follows, where fuzzy if-then rules are applied, combining the fuzzy inputs to generate rule activations. These activations are then passed to the Normalization Layer, which ensures that the influence of each rule is proportionally weighted, making the overall system more stable and interpretable. After that, the Consequent Layer calculates the outcomes of each rule based on its activation strength. Finally, in the Output Layer, all these results are aggregated into a single, crisp output—completing the prediction process. This structured, step-by-step approach allows ANFIS to capture complex, nonlinear relationships while retaining a clear and logical flow from input to output.

$$\mu_{ij}(x_i) = \exp\left(-\frac{(x_i - c_{ij})^2}{2\sigma_{ij}^2}\right) \quad (1)$$

Where c_{ij} and σ_{ij} are the center and width of the Gaussian membership function, respectively. The rule layer represents fuzzy rules, with each node corresponding to a rule that computes the product of the membership values.

$$w_k = \prod_{i=1}^n \mu_{ij}(x_i) \quad (2)$$

Where w_k is the firing strength of rule k . The normalization layer then normalizes the firing strengths of these rules to ensure they sum up to one.

$$w_k^- = \frac{w_k}{\sum_{k=1}^R w_k} \quad (3)$$

Where w_k^- is the normalized firing strength of rule k and R is the number of rules. The model calculates the output of each rule with a linear combination of the input variables and consequent factors in consequent layer.

$$y_k = w_k^- \cdot \left(a_{k0} + \sum_{i=1}^n a_{ki} x_i \right) \quad (4)$$

In equation 4 consequent parameters are a_{k0} and a_{ki} . Lastly, the output layer(y) combined these rule outputs to generate the final prediction.

$$y = \sum_{k=1}^R y_k \quad (5)$$

The final output is the average of the ANFIS and ANN outputs [8]. This hybrid model brings a powerful blend of strengths to the table. By merging the transparency and rule-based reasoning of ANFIS with the adaptive learning power of neural networks, it becomes especially effective in managing systems where both uncertainty and complexity are at play.

$$Output = \frac{ANFIS_Output + ANN_Output}{2} \quad (6)$$

One of its biggest advantages is its ability to model nonlinear relationships with precision—an essential feature when predicting dynamic parameters like water quality. Another key benefit is its built-in ability to fine-tune membership functions automatically, eliminating much of the manual effort typically required and leading to faster, more accurate results. That said, this enhanced performance does come at a cost [9]. Because the model involves training both ANFIS and ANN components, it demands a higher level of computational power. This can be a drawback in scenarios where real-time processing is needed or where computing resources are limited. Still, for many applications, the trade-off is worth it—offering a smart, flexible solution that balances accuracy with interpretability.

III. EXPERIMENT SETUP

A. Hyperparameter Tuning

The data preprocessing phase lays the groundwork for building a reliable and efficient model. It starts with cleaning and preparing the dataset to ensure all values are accurate, consistent, and ready for analysis. To bring all features (X) and the target variable (y) onto a common scale, MinMaxScaler is applied, normalizing the data to fall within a range of 0 to 1.

This step is crucial—it eliminates differences in scale between variables, which help the model train more smoothly and efficiently. Once the data is normalized, it's split into three distinct sets: 70% for training, 15% for validation, and 15% for testing [10]. The validation set plays a key role during training—it allows for fine-tuning of hyperparameters and provides early warnings against overfitting.

Meanwhile, the test set is kept completely separate until the very end, offering an unbiased evaluation of how well the model performs on new, unseen data. To make the data compatible with the PyTorch deep learning framework, the final step involves converting it into PyTorch tensors.

This transformation ensures seamless integration with the neural network architecture, setting the stage for efficient training, validation, and testing workflows.

At the heart of the hybrid model lies the ANFIS component, which leverages fuzzy logic and Gaussian membership functions to make sense of the non-linear relationships in the input data. Each input feature is linked to a set of fuzzy rules—10 in total per feature in this implementation. These rules use Gaussian functions to calculate membership values, with each function shaped by a mean (μ) and standard deviation (σ). By taking the product of these membership values across all features, the model generates normalized activations that reflect the collective influence of all rules. This step captures subtle patterns and relationships that might be missed by more rigid systems [11]. Working alongside this is the ANN component, which brings the learning power of deep neural networks into the equation. It consists of two hidden layers: the first with 64 neurons, and the second with 32. Both layers use the ReLU activation function to introduce non-linearity, allowing the model to learn from complex, layered data structures. The network ends with a single-node output layer, dedicated to making the final prediction. Together, this combination of fuzzy reasoning and deep learning strikes a strong balance—maintaining interpretability and structure while also offering the flexibility and precision needed to capture complex dynamics in the data.

To harness the strengths of both systems, the outputs from the ANFIS and ANN components are averaged, creating a balanced prediction that blends the clarity of fuzzy logic with the adaptive power of deep learning. During training, the model fine-tunes its parameters—this includes the mean (μ) and standard deviation (σ) of the Gaussian membership functions in the ANFIS, as well as the weights of the neural network [12]. This optimization is carried out using backpropagation, with the Adam optimizer (set to a learning rate of 0.01) driving the learning process for faster and more stable convergence. The model is trained over 2000 epochs, with a portion of the data set aside for validation. This validation data helps track the model's performance in real-time, ensuring it generalizes well and doesn't overfit. The Mean Squared Error (MSE) is used as the loss function, guiding the model to reduce prediction errors throughout the training cycle. This hybrid learning framework is especially well-suited for complex and noisy datasets. By combining interpretability with learning flexibility, it delivers reliable and accurate results—even when the data includes non-linear behaviors and unpredictable patterns.

B. Performance metrics

Evaluating models is also essential to evaluate models in measuring accuracy and effectiveness since it helps in checking their performance and guides in optimization efforts and reliability in decision-making processes [13]. MSE checks the performance of the model, guides optimization, and ensures reliability in decision making due to the difference between y_i and $\hat{y}_i$.

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (7)$$

Square root of the MSE (RMSE) represents the standard deviation of the prediction errors. It provides an easily interpretable measure of the average magnitude of prediction error in the same units as the original values.

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (8)$$

MAE is average absolute difference between actual and predicted values. Unlike MSE, it doesn't square the errors, making it less sensitive to outliers. It reflects the average prediction error in a model.

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (9)$$

Sum of Squared Errors (SSE) is sum of squared differences between predicted and actual outputs, and a lower SSE indicates a better fit.

$$SSE = \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (10)$$

Coefficient of Determination (R^2) represents the proportion of variance in the dependent variable that is predictable from the independent variables.

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (11)$$

The performance metrics mentioned (MSE, RMSE, MAE, SSE, and R^2) are commonly used because they provide a well-rounded assessment of model accuracy by addressing different aspects of error and model fit [14]. These metrics are chosen because they complement each other by covering different error perspectives (absolute vs. squared, sensitivity to outliers) and the model's explanatory power.

IV. RESULTS AND DISCUSSION

The Hybrid ANFIS-ANN model is developed using Python (version 3.8 or later), with essential libraries like PyTorch and NumPy forming the core of the implementation. It runs on a standard Windows operating system, utilizing a modern multi-core processor—such as an Intel i5—and at least 8 GB of RAM, ensuring smooth and efficient execution even when working with complex datasets. This setup allows the model to seamlessly integrate fuzzy logic rules with the learning capabilities of neural networks. Designed specifically for regression tasks, the hybrid approach strikes a thoughtful balance between high prediction accuracy and interpretability—making it a practical and insightful tool for real-world applications.

TABLE II: PERFORMANCE EVALUATION OF HYBRID MODEL ANFIS-ANN

Dataset	MSE	RMSE	MAE	SSE	R^2 Score
Training	0.0003	0.0159	0.0123	0.2762	0.9854
Validation	0.0003	0.0181	0.0127	0.0772	0.985
Testing	0.0005	0.0234	0.0128	0.1288	0.9718

The Hybrid ANFIS-ANN model delivers exceptional predictive accuracy across training, validation, and testing datasets. In the training phase, the model achieves a remarkably low MSE of 0.0003 and a high R^2 score of 0.9854, demonstrating a strong alignment between predicted and actual values. Similarly, the validation phase shows consistent performance with an identical MSE of 0.0003 and an R^2 score of 0.9850, highlighting the model's ability to generalize effectively to unseen data. In the testing phase, a slight increase in MSE to 0.0005 and a minor decrease in R^2 to 0.9718 indicate minimal performance drop when applied to new data. The consistently low RMSE and MAE across all phases further confirm the model's precision and reliability. These results underscore the Hybrid ANFIS-ANN model's capability to accurately capture complex, nonlinear relationships, making it a highly effective solution for predicting BODOUT values in wastewater treatment systems.




Figure 2: (a) Hybrid model actual vs predicted graph (b) loss curve over the epoch

The Actual vs Predicted BODOUT, illustrates the model's performance in figure 2, highlights the model's effectiveness, showing that predicted BODOUT values are closely aligned with actual values for both validation and test datasets. The clustering of points near the red dashed diagonal line signifies high prediction accuracy with minimal deviation. The accompanying plot, loss curve over epochs, depicts a rapid decrease in loss during the initial training phase, followed by a steady plateau at a low loss value, indicating successful convergence and efficient learning. These results collectively demonstrate the strong predictive capability and reliable training dynamics of the Hybrid ANFIS-ANN model.




Figure 3: Sample membership function for parameters

The membership functions for sample parameters: PHIN, CODIN, and BODIN, illustrating the fuzzy membership degrees associated with normalized feature values. Each curve represents a distinct fuzzy rule, defining the degree to which a parameter belongs to a specific fuzzy set is shown in figure 3. These membership functions play a critical role in addressing uncertainties and imprecision within the input data. By leveraging these functions, the model ensures accurate predictions and supports robust decision-making, particularly in the complex and dynamic context of wastewater treatment processes.




Figure 4: Major Significant deviations of BOD for 5 years data.

The residuals from a predictive model over a 5-year period are present in figure 4. The blue line represents the residual values, which measure the gap between actual and predicted data points. Significant deviations, where the residuals surpass a certain threshold, are marked with red dots. Using a hybrid model that combines ANFIS and ANN, these significant deviations in BOD were accurately predicted. This demonstrates the model's capability to identify anomalies and effectively handle complex non-linear relationships in the data.

V. CONCLUSION

This study showcases the successful application of a hybrid ANFIS-ANN model for predicting BOD in wastewater treatment processes. By combining the adaptive, fuzzy logic-based decision-making of ANFIS with the deep learning capabilities of ANN, the model is able to navigate the complexities and uncertainties inherent in water quality data. The results from our experiments are impressive, with the model achieving a high R^2 score of 0.9718 during testing, underscoring its accuracy and reliability. This hybrid model provides a practical, cost-

effective solution for wastewater treatment plants, enhancing operational safety, boosting efficiency, and ensuring compliance with environmental standards. However, one limitation of this study is its reliance on daily data from a single industrial STP, which may limit the model's applicability to other treatment systems. Additionally, the model's accuracy could be impacted when tested on datasets with different characteristics or lower-quality measurements. Despite these limitations, the findings of this research set a new benchmark for water quality monitoring and open the door for further innovations in predictive modeling for wastewater management.

REFERENCES

- [1] Pratap, B., Kumar, S., Nand, S., Azad, I., Bharagava, R. N., Ferreira, L. F. R., & Dutta, V. (2023). Wastewater generation and treatment by various eco-friendly technologies: Possible health hazards and further reuse for environmental safety. *Chemosphere*, 313, 137547.
- [2] Aguilar-Torrejón, J. A., Balderas-Hernández, P., Roa-Morales, G., Barrera-Díaz, C. E., Rodríguez-Torres, I., & Torres-Blancas, T. (2023). Relationship, importance, and development of analytical techniques: COD, BOD, and, TOC in water—An overview through time. *SN Applied Sciences*, 5(4), 118.
- [3] Inbar, O., Shahar, M., & Avisar, D. (2024). Predictive modeling of BOD throughout wastewater treatment: A generalizable machine learning approach for improved effluent quality. *Environmental Science: Water Research & Technology*, 10(10), 2577-2588.
- [4] Varalakshmi, V., Saravanan, M., Abhilash, P., Ramadasu, T. L., Saritha, P., & Vempada, S. R. (2024). Optimization of Wastewater Treatment Processes Using AI-Driven Machine Learning Algorithms for Enhanced Biological Degradation Efficiency. *Library Progress International*, 44(3), 17485-17492.
- [5] Ibrahim, M., Haider, A., Lim, J. W., Mainali, B., Aslam, M., Kumar, M., & Shahid, M. K. (2024). Artificial neural network modeling for the prediction, estimation, and treatment of diverse wastewaters: a comprehensive review and future perspective. *Chemosphere*, 142860.
- [6] Bressane, A., Garcia, A. J. D. S., Castro, M. V. D., Xerfan, S. D., Ruas, G., & Negri, R. G. (2024). Fuzzy Machine Learning Applications in Environmental Engineering: Does the Ability to Deal with Uncertainty Really Matter?. *Sustainability*, 16(11), 4525.
- [7] Ayed, N., & Bougatef, K. (2023). Performance Assessment of Logistic Regression (LR), Artificial Neural Network (ANN), Fuzzy Inference System (FIS) and Adaptive Neuro-Fuzzy System (ANFIS) in Predicting Default Probability: The Case of a Tunisian Islamic Bank. *Computational Economics*, 1-33.
- [8] Kanwal, S., & Jiriwibhakorn, S. (2024). Advanced Fault Detection, Classification, and Localization in Transmission Lines: A Comparative Study of ANFIS, Neural Networks, and Hybrid Methods. *IEEE Access*.
- [9] Küçüktopcu, E., Cemek, E., Cemek, B., & Simsek, H. (2023). Hybrid statistical and machine learning methods for daily evapotranspiration modeling. *Sustainability*, 15(7), 5689.
- [10] Nimmy, P., Nagaraja, K. V., Srilatha, P., Karthik, K., Sowmya, G., Kumar, R. V. & Ali, M. R. (2023). Implication of radiation on the thermal behavior of a partially wetted dovetail fin using an artificial neural network. *Case Studies in Thermal Engineering*, 51, 103552.
- [11] Leski, J., & Henzel, N. (2024). A neuro-fuzzy system based on logical interpretation of if-then rules. In *Fuzzy Learning and Applications* (pp. 359-388). CRC Press.
- [12] Kaveh, A. (2024). Applications of artificial neural networks and machine learning in civil engineering. *Studies in computational intelligence*, 1168, 472.
- [13] Hassan, E., Shams, M. Y., Hikal, N. A., & Elmougy, S. (2023). The effect of choosing optimizer algorithms to improve computer vision tasks: a comparative study. *Multimedia Tools and Applications*, 82(11), 16591-16633.
- [14] Keddouda, A., Ihaddadene, R., Boukhari, A., Atia, A., Arıcı, M., Lebbihiat, N., & Ihaddadene, N. (2024). Photovoltaic module temperature prediction using various machine learning algorithms: Performance evaluation. *Applied Energy*, 363, 123064.