

Beyond Tides and Time: Machine Learning's Triumph in Water Quality Forecasting

Yinpu Li^{1,#}, Siqi Mao^{1,*,#}, Yaping Yuan^{2,#}, Ziren Wang^{3,#}, Yixin Kang¹, Yuanxin Yao¹

¹Department of Statistics, Florida State University, Tallahassee, USA

²Department of Mathematics and Statistics, University of Massachusetts Amherst, Amherst, USA

³Department of Statistics, Rice University, Houston, USA

*Corresponding author: mden17g@gmail.com

#These authors contributed equally to this work

Received September 26, 2023; Revised October 27, 2023; Accepted November 02, 2023

Abstract Water resources are essential for sustaining human livelihoods and environmental well-being. Accurate water quality prediction plays a pivotal role in effective resource management and pollution mitigation. In this study, we assess the effectiveness of five distinct predictive models—linear regression, Random Forest, XGBoost, LightGBM, and MLP neural network—in forecasting pH values within the geographical context of Georgia, USA. Notably, LightGBM emerges as the top-performing model, achieving the highest average precision. Our analysis underscores the supremacy of tree-based models in addressing regression challenges, while revealing the sensitivity of MLP neural networks to feature scaling. Intriguingly, our findings shed light on a counter-intuitive discovery: machine learning models, which do not explicitly account for time dependencies and spatial considerations, outperform spatial-temporal models. This unexpected superiority of machine learning models challenges conventional assumptions and highlights their potential for practical applications in water quality prediction. Our research aims to establish a robust predictive pipeline accessible to both data science experts and those without domain-specific knowledge. In essence, we present a novel perspective on achieving high prediction accuracy and interpretability in data science methodologies. Through this study, we redefine the boundaries of water quality forecasting, emphasizing the significance of data-driven approaches over traditional spatial-temporal models. Our findings offer valuable insights into the evolving landscape of water resource management and environmental protection.

Keywords: water quality prediction, linear-regression, random forest, xgboost, lightgbm, mlp neural network

Cite This Article: Yinpu Li, Siqi Mao, Yaping Yuan, Ziren Wang, Yixin Kang, and Yuanxin Yao, "Beyond Tides and Time: Machine Learning's Triumph in Water Quality Forecasting." American Journal of Applied Mathematics and Statistics, vol. 11, no. 3 (2023): 89-97. doi: 10.12691/ajams-11-3-2.

1. Introduction

In today's era, the significance of our water resources' quality is unparalleled, influencing every aspect of human existence and the natural world. Water, a fundamental necessity for survival, holds implications that extend well beyond quenching mere thirst. The presence of polluted water sources and deteriorated aquatic ecosystems has given rise to a multitude of problems encompassing public health, ecological balance, economic factors, and societal fairness. The proactive monitoring and projection of water quality play a pivotal role in protecting aquatic ecosystems such as rivers, lakes, and oceans. Deteriorated water quality can lead to the devastation of habitats, adverse effects on aquatic organisms, and disruption of the overall ecological equilibrium.

Among the critical parameters for evaluating water quality, the water's pH stands out as one of the most critical factors. It quantifies the level of acidity or

alkalinity within the water. Water possessing a pH value of 11 or higher has the potential to induce irritation in the eyes, skin, and mucous membranes, underscoring the importance of pH assessment in water quality investigations. As highlighted by Geetha et al. (2016) [1], the Internet of Things (IoT) is also contributing to advancements in water quality monitoring. Consequently, the prediction of water quality, particularly in relation to pH levels, has become increasingly imperative in recent times.

Recognizing the significance of maintaining water quality, the need extends beyond mere monitoring to encompass proactive prediction. This proactive approach guarantees timely public alerts regarding potential contamination, subsequently averting the associated health risks and economic losses. There are many traditional water quality prediction methods, such as multiple linear regression [2] and auto-regressive integrated moving average (ARIMA) [3]. However, the linear nature of multiple linear regression poses a limitation in detecting nonlinear relationships among water quality parameters [4]. Similarly, the primary drawback of ARIMA lies in its

underlying assumption of linearity [5]. During the process of model identification, the time series data must undergo scrutiny to determine their stationarity, a crucial aspect in constructing the ARIMA model. Notably, conventional methodologies struggle to effectively capture the non-linear [6] and non-stationary [7] characteristics inherent in water quality due to their intricate and complex nature.

In recent years, machine learning approaches have been widely applied to multiple domains and achieved gratifying results (e.g., [8-10, 28-30]). When it comes to estimating water quality using machine learning, Lu et al. [11] applied two hybrid decision tree-based water quality machine learning models: extreme gradient boosting (XGBoost) and random forest (RF), proposed to obtain more accurate short-term water quality prediction results, by using the water resources of Gales Creek site in Tualatin River. Huang et al. (2019) [12] established a prediction system for urban estuary water quality and used the gradient boosting machine model to fill and predict the flow. Wang et al. (2022) [13] applied several machine learning models—multiple linear regression, artificial neural networks, random forest, and extreme gradient boosting (XGBoost)—were developed to predict NH_4^+ -N in the Xiaoqing River estuary, China. The shapely additive explanations method [14] was used to interpret the XGBoost model and discover the influence of the upper reaches of the river on the estuary. In the research [15], Li et al. (2022) evaluated five tree-based models, namely classification tree, random forest, CatBoost, XGBoost, and LightGBM, and employed a state-of-the-art explanation method SHAP to explain the models.

As the volume of data continues its relentless expansion, traditional approaches are proving inadequate to cope with the demands of researchers. With the advent of increased computing power, data-driven models like artificial neural networks (ANNs) have undergone substantial improvements. These models excel at capturing the inherent functional relationships inherent within water quality data, as evidenced by examples in Zhang et al. (1998) [16]. Even in situations where articulating intricate data relationships proves challenging, ANNs have proven their effectiveness. Moreover, ANNs require fewer initial assumptions [17] while delivering heightened precision [18] compared to established techniques. Furthermore, Singh et al. (2009) [19] utilized the ANN model to predict the water quality of the Gomti River in India, showcasing the versatility of these models. García-Alba et al., 2019 [20] employed an ANN-based model to predict estuary bathing water quality, integrating laboratory analysis, machine learning, and numerical simulation for real-time water quality management. Peng et al. (2019) [21] proposed a framework for real-time prediction of daily water quality, successfully applying it to Lake Chaohu in China, thereby improving predictions for parameters such as dissolved oxygen and total phosphorus.

The primary aim of this inquiry is to offer enlightenment to individuals who possess a strong background in data science but may be less acquainted with the realm of environmental research. We present a comprehensive framework for the application of data science knowledge and methodologies, facilitating their conversion into tangible applications across diverse research domains. This encompasses areas such as water

quality prediction, enabling the practical utilization of these skills in various contexts.

In this study, we test the proposed framework by applying the dataset used in [22]. This input data consists of daily water quality samples from 37 sites, providing measurements related to pH values in Georgia, USA. The input features consist of 11 common indices including the volume of dissolved oxygen, temperature, and specific conductance. The proposed framework is examined by forecasting water quality in terms of the “power of hydrogen (pH)” value based on the input data.

The remaining sections of the paper are structured as follows: In Section 2, we delve into the acquired data and the prediction methodology. This section extensively elucidates the various models employed in our study, encompassing Linear Regression, XGBoost, LightGBM, Random Forest, and a Multiple-layer Perceptron Neural Network. It also encompasses a thorough description of the evaluation metrics employed and provides a comprehensive overview of the entire implementation process. Moreover, this section entails an in-depth performance comparison across different models, followed by a detailed analysis. Additionally, we include a SHAP (SHapley Additive exPlanations) analysis within this section. Section 3 succinctly elucidates why our machine learning models outperform the original model, which accounts for both time dependency and spatial factors. Lastly, Section 4 encapsulates the main conclusions drawn from the study and outlines potential avenues for future research.

Table 1. Data Description

Variable Name	Size	Description
features	1×11	A list of water indices to measure.
location_ids	37×1	IDs of the water stations.
X_{te}	1×282	Test set input data: water indices for 282 contiguous dates until 2018-01-01; each element is a 37 × 11 matrix: 37 spatial locations by 11 features.
X_{tr}	1×423	Training set input data: water indices for 423 contiguous dates from 2016-01-28; each element is a 37 × 11 matrix: 37 spatial locations by 11 features.
Y_{te}	37×282	Test set output data: water quality for 37 locations in 282 contiguous dates until 2018-01-01.
Y_{tr}	37×423	Training set output data: water quality for 37 locations in 423 contiguous dates from 2016-01-28.
location_group	1×3	The groups of water stations, each group forms a connected spatial network (i.e., water system).

2. Material and Methods

2.1. Data

The daily monitoring data for 37 sites in Georgia, USA from 2016 to 2018 was collected by the United States Geological Survey¹. The dataset presented in the original

¹ USGS: <https://www.usgs.gov/>. Accessed Feb, 2018

Spatial-Autoregressive Dependency Learning II (SADL-II)

A spatio-temporal model (SADL-II) was proposed in the original paper [22]. On the basis of spatial topological restrictions, SADL-II was suggested to automatically learn the conditional independence pattern as well as quantify the numerical values of the spatial dependency. It is considered as an alternative benchmarking model in this paper.

Linear Regression with Elastic Net Regularization

Multiple Linear Regression is a statistical technique used to model the relationship between multiple independent variables and a dependent variable. It extends the concept of simple linear regression by considering multiple predictors. The general form of Multiple Linear Regression is as follows:

$$Y = \beta_0 + \beta_1 \times X_1 + \dots + \beta_n \times X_n + \epsilon,$$

where $\beta_0, \beta_1, \dots, \beta_n$ are n coefficients and ϵ is a random error, Y is the target variable, and X_1, X_2, \dots, X_n are predictors.

Elastic Net is a linear regression technique that combines the features of both L1 (Lasso) [23] and L2 (Ridge) [24] regularization methods. It is a regularization and variable selection technique used in machine learning and statistics to prevent overfitting and improve the stability and generalization of linear regression models. The loss function is as follows:

$$\mathcal{L} = \frac{\sum_{i=1}^n (y_i - x_i' \beta)^2}{2n} + \lambda \frac{1-\alpha}{2} \sum_{j=1}^k \beta_j^2 + \lambda \alpha |\beta_j|.$$

In this loss function, the new parameter α is a "mixing" parameter that balances the two approaches. If α is one, this is Lasso regression, and if α is zero, then this is Ridge regression.

Lasso can improve prediction accuracy and model interpretability by performing both feature selection and regularization. Lasso's feature selection property provides a clear interpretation of the most important predictors in the model. The non-zero coefficients indicate the features that strongly impact the prediction outcome [23].

XGBoost

XGBoost was first introduced in [25]. They provided an in-depth explanation of the algorithm, its optimization techniques, and empirical evaluation of its performance on various datasets. XGBoost scales beyond billions of examples using far fewer resources than existing systems. The objective function (loss function and regularization) at iteration t that we need to minimize is the following:

$$\mathcal{L}^{(t)} = \sum_{i=1}^n l \left(y_i, y_i + f_t(x_i) \right) + \Omega(f_t),$$

where l is the loss function, f_t is the t -th tree output, and Ω is the regularization term to control the complexity of the model and prevent the model from overfitting.

LightGBM

LightGBM, an open-source machine learning framework created by Microsoft, is an optimized tool

tailored for gradient boosting. Gradient boosting is a widely-used machine learning approach that constructs predictive models by combining the outcomes of numerous weak models, often in the form of decision trees. One of the primary distinctions that sets LightGBM apart from the traditional gradient boosting tree decision technique is the use of a method called GOSS (Gradient-based One-Side Sampling).

During the training process, GOSS retains all the data points with significant gradients while randomly subsampling the data with lower gradients. This strategic approach effectively reduces the search space, enabling GOSS to converge more swiftly.

LightGBM has gained acclaim for its remarkable speed and efficiency, rendering it a preferred choice across a range of machine learning tasks, including classification, regression, and ranking.

Random Forest

In 2001, Leo Breiman [26] presented the Random Forest algorithm and discussed its principles and advantages. Random Forest is a powerful ensemble learning method in machine learning, primarily used for classification and regression tasks. It is an ensemble of decision trees, where multiple decision trees are trained independently, and their outputs are combined to make predictions.

Multiple-layer Perceptron (MLP) neural network

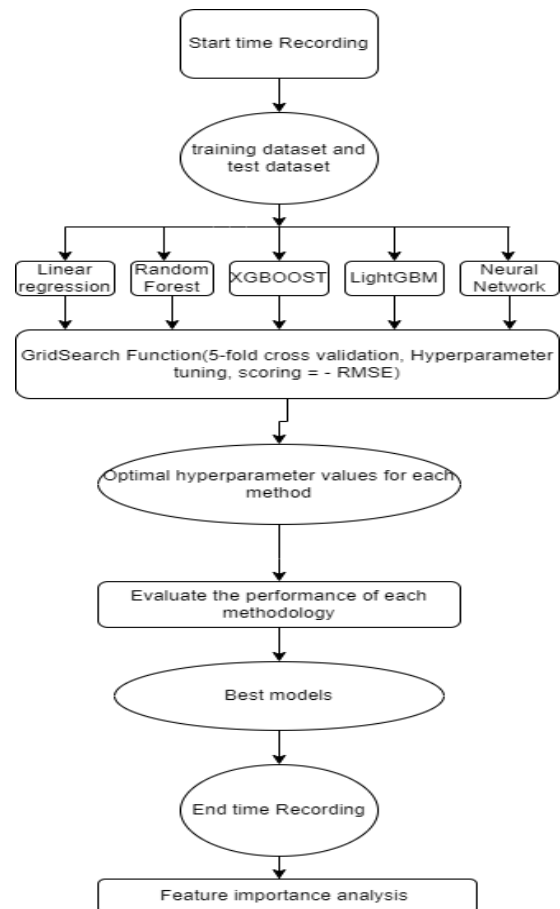


Figure 2. Flowchart of The Proposed Water Quality Prediction Framework.

Multiple-layer Perceptron (MLP) neural network represents an artificial neural network model that leverages the backpropagation technique to iteratively

refine the connections between neurons, thereby enhancing its predictive accuracy. This implementation integrates the Multi-Layer Perceptron (MLP) algorithm [27], harnessing backpropagation and stochastic gradient descent strategies for training and evaluating datasets. It offers a range of customizable parameters, allowing users to finely tune the model's performance by adjusting factors such as the number of hidden layers, activation functions, optimization solvers, and more. MLPRegressor stands as an effective solution for addressing regression tasks, demonstrating proficiency in capturing intricate non-linear relationships between input and output variables.

2.5. Implementation

In this section, we detail the proposed prediction framework through hyperparameter tuning, generating predictions, and assessing their effectiveness using tailored scoring metrics. Additionally, we capture the time taken for these tasks and endeavor to identify disparities in feature selection, as evidenced by variations in feature importance.

1. Choose the Scoring Function: For all candidate methods, we use the negative root of mean squared error as the scoring function for hyperparameter tuning and evaluation.
2. Define Hyperparameters: Set up the hyperparameter space for each method.
3. Start Time Recording: Begin recording the time taken for the entire procedure.
4. Initialize Models: Initialize the models for each candidate method.
5. Hyperparameter Tuning: Perform hyperparameter tuning using cross-validation to find the best hyperparameters for each method.
6. Final Predictions: Make final predictions on the testing data using the best-tuned models.
7. Score Calculation: Calculate the evaluation metrics on the test data for each method.
8. End Time Recording: Stop recording the time taken for the procedure.
9. Feature Importance: Obtain feature importance scores from the best-tuned models.
10. Save Results: Save the results of each method's performance and feature importance for model comparison.

2.6. Model Selection

2.6.1. Evaluation Metrics

Based on the nature of the problem, and to compare against the original paper, we have chosen the following metrics: root mean square error (RMSE), mean absolute percentage error (MAPE), weighted mean absolute percentage error (WMAPE), weighted under prediction (WUPRED) and weighted over prediction (WOPRED). These metrics are defined by the following equations:

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

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right|$$

$$WMAPE = \frac{\sum_{i=1}^n |y_i - \hat{y}_i|}{\sum_{i=1}^n |y_i|}$$

$$WUPRED = \frac{\sum_{i=1}^n I\{y_i > \hat{y}_i\} (y_i - \hat{y}_i)}{\sum_{i=1}^n y_i}$$

$$WOPRED = \frac{\sum_{i=1}^n I\{y_i < \hat{y}_i\} (\hat{y}_i - y_i)}{\sum_{i=1}^n y_i}$$

where y_i represents the observed PH value, \hat{y}_i is the value of prediction. In the context of all five metrics, our goal is to minimize them as much as possible. RMSE assesses the error magnitude in the same units as the predicted values. In contrast, MAPE highlights relative percentage errors, effectively mitigating the influence of outliers on evaluations. WMAPE builds upon MAPE by introducing weighted adjustments that account for the varying importance of different data points. Meanwhile, WUPRED and WOPRED quantify errors when forecasts fall short or exceed actual values, shedding light on potential model biases. Collectively, these metrics provide a comprehensive evaluation of predictive performance, considering both error direction and magnitude, and can be invaluable in guiding enhancements to predictive models.

2.6.2 Results

This section presents the prediction results and error analysis results of PH values. During the cross validation process, 5-fold cross validation is selected. To test if the spatio-temporal features are contributing to the predictions, three model implementation strategies are derived:

1. **Strategy 1:** We select eleven numerical features as input features to predict the target variable.
2. **Strategy 2:** We select the standardized eleven numerical features as input features to predict the target variable.
3. **Strategy 3:** The standardized numerical features and the one-hot encoded categorical features are selected to predict the target variable.

For strategy 1, the model performance summary is shown in Table 4. Based on the performance of

Benchmarking and other advanced models, we can conclude that employing advanced models is essential. In addition, the XGBoost beats all the other candidate models in terms of MAPE, WMAPE and WOFOREC. For metric WUFOREC, MLP has the best performance. Meanwhile, the lightGBM outperforms the other models under the metric RMSE. Tree-based methods have shown their superiority in regression problems. There are several candidate models outperform the SADL-II (Original Paper).

Table 4. Water Quality Prediction Results (Numerical Features Only)

Models	RMS E	MAP E	WMAPE E	WUPRE D	WOPRE D
SADL-II (Original Paper)	11.50	N/A	N/A	N/A	N/A
Benchmarking	29.52	32.21	32.35	14.55	17.80
LightGBM	10.88	9.34	9.48	5.20	4.28
Linear Regression	12.04	10.72	10.90	5.65	5.24
MLP	11.62	10.00	10.15	5.13	5.02
Random Forest	11.47	9.71	9.87	5.34	4.52
XGBoost	10.96	9.23	9.37	5.17	4.19

For strategy 2 (model performance summary detailed in Table 5), we have the same conclusion for the tree based models as what we had in strategy 1. In addition, the performance of XGBoost and LightGBM had a slight improvement after the standardization of numerical features. However, the performance of MLP got worse than the linear regression with elastic net regularization after the standardization.

Table 5. Water Quality Prediction Results (Numerical Features Only; Standardized)

Models	RMSE	MAPE	WMAPE	WUPRED	WOPRED
SADL-II (Original Paper)	11.50	N/A	N/A	N/A	N/A
Benchmarking	29.52	32.21	32.35	14.55	17.80
LightGBM	10.84	9.33	9.46	5.16	4.30
Linear Regression	12.04	10.74	10.92	5.66	5.25
MLP	14.20	12.92	13.16	3.83	9.33
Random Forest	11.47	9.71	9.87	5.35	4.51
XGBoost	10.96	9.22	9.36	5.17	4.19

In strategy 3 (model performance detailed in Table 6), the spatio-temporal features are involved. The lightGBM beats all the other models when it comes to the metric RMSE, MAPE and WMAPE. However, the general performance of models slightly degraded after the spatio-temporal features are inputted. Same as the results in strategy 1 and strategy 2, there are several candidate models outperform the SADL-II (Original Paper).

We also provided hyper-parameters tuning time for strategy 2 and strategy 3 in Table 7 and Table 8. In the tables, "Total Fits" stands for the number of fits in hyper-parameters tuning of each model. "Tuning Time" is for running time of entire hyper-parameters tuning process of each model. "Fitting Time (Best Model)" stands for the training time for the selected best hyper-parameters of each model. "Average Tuning" is for the average hyper-

parameters tuning time. The tuning time for strategy 1 is similar to the one for strategy 2, therefore, it is not provided in the paper. In addition, the benchmark model does not have a hyper-parameters tuning process. In conclusion, the lightGBM has a lowest average tuning time when the spatio-temporal features are included. Otherwise, the linear regression with elastic net regularization is the most efficient one.

Table 6. Water Quality Prediction Results (Numerical Features: Standardized; Categorical Features: One-Hot Encoding)

Models	RMSE	MAPE	WMAPE	WUPRED	WOPRED
SADL-II (Original Paper)	11.50	N/A	N/A	N/A	N/A
Benchmarking	29.52	32.21	32.35	14.55	17.80
LightGBM	11.04	9.22	9.36	5.15	4.21
Linear Regression	11.91	10.79	10.96	6.14	4.82
MLP	14.42	14.54	14.61	9.29	5.32
Random Forest	11.71	9.71	9.87	5.26	4.61
XGBoost	11.14	9.29	9.44	5.28	4.16

Table 7. Running Time Summary (Numerical Features: Standardized)

Models	Total Fits	Tuning Time	Average Tuning	Fitting Time (Best Model)
LightGBM	2000	159.92	0.08	0.19
Linear Regression	150	3.32	0.02	0.07
MLP	40	1948.63	48.72	116.05
Random Forest	1200	1934.51	1.61	64.58
XGBoost	5760	2429.88	0.42	4.30

Table 8. Running Time Summary in Seconds (Numerical Features: Standardized; Categorical Features: One-hot encoding)(in second)

Models	Total Fits	Tuning Time	Average Tuning	Fitting Time (Best Model)
LightGBM	2000	213.44	0.11	0.46
Linear Regression	150	44.39	0.30	0.82
MLP	40	2147.05	53.68	65.81
Random Forest	1200	3633.05	3.03	115.67
XGBoost	5760	13320.13	2.31	4.16

2.7. SHapley Additive exPlanations (SHAP)

For each prediction, the SHAP values associated with environmental variables serve to quantify their localized contributions to that specific prediction (Lundberg and Lee 2017). The mathematical definition of the SHAP value is provided below:

$$\begin{aligned} \Phi_i(F, x) &= \sum_{Z \subseteq X'} \left[\binom{M}{|Z|} \right]^{-1} [F(Z) - F(Z \setminus x_i)] \\ &= \sum_{Z \subseteq X'} \frac{(|Z|-1)!(M-|Z|)!}{M!} [F(Z) - F(Z \setminus x_i)]. \end{aligned}$$

In this equation, $\Phi_i(F, x)$ represents the SHAP value corresponding to feature x_i in the context of a model F constructed on a set of features X . M represents the total number of input features, X' denotes the set of all potential

feature combinations that include feature x_i , and $|Z|$ signifies the number of features within a particular feature combination $|Z|$. Additionally, $F(Z)$ and $F(Z \setminus x_i)$ represent distinct predictive models trained on $|Z|$ and $Z \setminus x_i$ (which is Z with feature x_i removed), respectively. Thus, the SHAP value is determined by aggregating the marginal contributions $F(Z) - F(Z \setminus x_i)$ from all feasible feature combinations (Z) through a weighted average.

SHAP values are computed using the Python implementation of SHAP, as described by Lundberg and

Lee in [14]. In this study, we calculated SHAP values for some candidate models for all strategies. They provided similar results crossing the strategies, so we displayed the visualization of SHAP values for strategy 2 for illustration purposes.

Figure 3 and Figure 4 are the illustration of SHAP values for LightGBM model. From the plots, the feature X6 (“Dissolved oxygen, water, unfiltered, milligrams per liter (Maximum)”) has the significant impact in predicting the target variable. Its impact is much higher than the other features. From Figure 5, XGBoost’s SHAP values generally agree with the SHAP values of LightGBM.




Figure 3. Average Impact (SHAP value) on model output for LightGBM




Figure 4. Impact (SHAP value) on model output for LightGBM




Figure 5. Average Impact (SHAP value) on model output for XGBoost

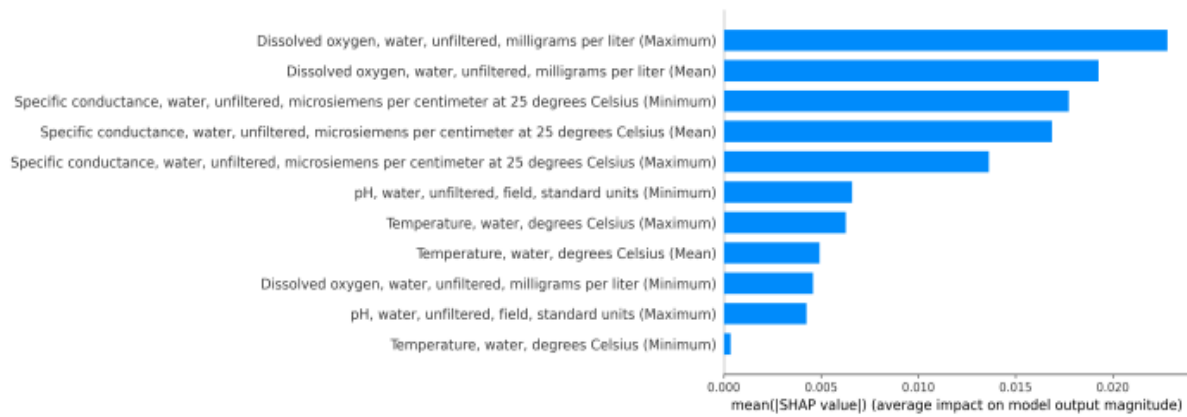


Figure 6. Average Impact (SHAP value) on model output for Linear Regression

3. Findings

We found some machine learning models are showing better performance than the spatial topological model described in the original paper. There could be several reasons why the machine learning models outperform the SADL-II model in the original paper that takes into account time dependency and spatial features. Here are a few possible explanations:

3.1. Data Quality and Quantity

The availability of a larger and higher-quality dataset can have a significant impact on model performance. The data provided by the original paper is extensive and contains no missing values, the machine learning models benefit from having more reliable information.

3.2. Feature

The efficacy of machine learning models is profoundly influenced by the quality and pertinence of the features employed for training. It is conceivable that the features provided by the original paper and we meticulously engineered have proven to be insightful and pertinent in predicting water pH values. Remarkably, the machine learning approaches leverage the automated feature selection during training time, which is done via the feature importance scores, aiding in feature selection and engineering.

3.3. Model Complexity and Flexibility

Machine learning models, especially ensemble methods like Random Forest and XGBoost, have the capability to capture complex relationships and patterns in the data. They can automatically learn interactions between features and nonlinear relationships that the original model might not have been able to capture effectively.

These ensemble models combine multiple individual models to make predictions. This aggregation reduces the risk of overfitting, as the errors of individual models tend to cancel out, leading to more robust predictions on unseen data. And they strike a balance between bias and variance. By averaging or combining predictions from multiple models, they reduce variance while maintaining a low bias, resulting in more accurate predictions.

3.4. Hyperparameter Tuning

We employ a systematic hyperparameter tuning process for all candidate models. This process aims to strike a balance among model complexity, interpretability, and predictability. It enables us to find the model that best fits the data, even without prior expert knowledge in the water quality field.

4. Discussion

In our investigation to predict water pH values, we adopted several Machine learning approaches and extended our analysis to include temporal and spatial features. Surprisingly, despite adding these features, the improvement in prediction accuracy was not as significant as expected. However, the results still demonstrated superior performance compared to the original temporal spatial model proposed by Chang et.al [7] in their study.

Our findings challenge the belief (from the original paper) that incorporating time-dependent and spatial features would invariably enhance the accuracy of predictive models. The success of machine learning models lies in the meticulous process of feature selection, engineering, and model optimization.

Suitable feature engineering helps the machine models find hidden patterns in the data that the original model didn't catch. Also, ensemble tree methods like LightGBM and XGBoost are good at finding complicated patterns in the data.

In comparison to the original model's approach of leveraging time dependencies and spatial dependencies, the machine learning models exhibited higher predictive accuracy. This discrepancy may be attributed to the more sophisticated modeling techniques and the capacity to harness the power of the ensemble approach, both of which are integral to modern machine learning.

Our study underscores the importance of a holistic approach to predictive modeling, wherein data preprocessing, feature engineering, model selection, and hyperparameter tuning collectively contribute to the overall performance. While time-dependent and spatial features can certainly enhance predictive capabilities, our results emphasize the critical role of model selection and optimization in achieving remarkable accuracy.

In summary, our exploration into predicting water pH values with machine learning has not only yielded superior results compared to the original model but also shed light on the multifaceted nature of predictive modeling. As the field of machine learning continues to evolve, the intricate interplay between features, algorithms, and optimization techniques emerges as a pivotal factor in determining the success of predictive endeavors. We will test this standardized framework on other spatial-temporal prediction tasks in future work.

ACKNOWLEDGMENTS

The source code for this project is available on GitHub at the following URL: <https://github.com/YinpuLi/water-quality-prediction>. The repository contains the implementation of the algorithms and methods comparison discussed in this paper.

References

- [1] Geetha, S., Gouthami, S.: Internet of things enabled real time water quality monitoring system. *Smart Water* 2(1). 1–19. 2016.
- [2] Rajae, T., Boroumand, A.: Forecasting of chlorophyll-a concentrations in south san francisco bay using five different models. *Applied Ocean Research* 53, 208–217. 2015.
- [3] Araghinejad, S.: Data-driven Modeling: Using MATLAB® in Water Resources and Environmental engineering. *Springer Science & Business Media*. Vol. 67. 2013.
- [4] Nourani, V., Alami, M.T., Vousoughi, F.D.: Self-organizing map clustering technique for ann-based spatiotemporal modeling of groundwater quality parameters. *Journal of Hydroinformatics* 18(2), 288–309. 2016.
- [5] Zare, A., Bayat, V., Daneshkare, A.: Forecasting nitrate concentration in ground-water using artificial neural network and linear regression models. *International agrophysics* 25(2). 2011.
- [6] Huo, S., He, Z., Su, J., Xi, B., Zhu, C.: Using artificial neural network models for eutrophication prediction. *Procedia Environmental Sciences* 18, 310–316. 2013.
- [7] Chang, F.-J., Chen, P.-A., Chang, L.-C., Tsai, Y.-H.: Estimating spatio-temporal dynamics of stream total phosphate concentration by soft computing techniques. *Science of the Total Environment* 562, 228–236. 2016.
- [8] Chen, D.Q., Mao, S.-Q., Niu, X.-F.: Tests and classification methods in adaptive designs with applications. *Journal of Applied Statistics* 50(6), 1334–1357. 2023.
- [9] Li, Y., Linero, A.R., Murray, J.: Adaptive conditional distribution estimation with bayesian decision tree ensembles. *Journal of the American Statistical Association*, 1–14. 2022.
- [10] Henrique, B.M., Sobreiro, V.A., Kimura, H.: Literature review: Machine learning techniques applied to financial market prediction. *Expert systems with applications* 124, 226–251. 2019.
- [11] Lu, H., Ma, X.: Hybrid decision tree-based machine learning models for short-term water quality prediction. *Chemosphere* 249, 126169. 2020.
- [12] Huang, P., Trayler, K., Wang, B., Saeed, A., Oldham, C.E., Busch, B., Hipsey, M.R.: An integrated modelling system for water quality forecasting in an urban eutrophic estuary: The swan-canning estuary virtual observatory. *Journal of Marine Systems* 199, 103218. 1995.
- [13] Wang, S., Peng, H., Liang, S.: Prediction of estuarine water quality using interpretable machine learning approach. *Journal of Hydrology* 605, 127320. 2022.
- [14] Lundberg, S.M., Lee, S.L.: A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol.30. 2017.
- [15] Li, L., Qiao, J., Yu, G., Wang, L., Li, H.-Y., Liao, C., Zhu, Z.: Interpretable tree-based ensemble model for predicting beach water quality. *Water Research* 211, 118078. 2022.
- [16] Zhang, G., Patuwo, B.E., Hu, M.Y.: Forecasting with artificial neural networks: The state of the art. *International journal of forecasting* 14(1), 35–62. 1998.
- [17] Anmala, J., Meier, O.W., Meier, A.J., Grubbs, S.: Gis and artificial neural network-based water quality model for a stream network in the upper green river basin, Kentucky, USA. *Journal of Environmental Engineering* 141(5), 04014082. 2015.
- [18] Li, L., Jiang, P., Xu, H., Lin, G., Guo, D., Wu, H.: Water quality prediction based on recurrent neural network and improved evidence theory: a case study of qiantang river, China. *Environmental Science and Pollution Research* 26, 19879–19896. 2019.
- [19] Singh, K.P., Basant, A., Malik, A., Jain, G.: Artificial neural network modeling of the river water quality—a case study. *Ecological modelling* 220(6), 888–895. 2009.
- [20] Garc'ia-Alba, J., B'arcena, J.F., Ugarteburu, C., Garc'ia, A.: Artificial neural networks as emulators of process-based models to analyse bathing water quality in estuaries. *Water research* 150, 283–295. 2019.
- [21] Peng, Z., Hu, W., Liu, G., Zhang, H., Gao, R., Wei, W.: Development and evaluation of a real-time forecasting framework for daily water quality forecasts for lake chaohu to lead time of six days. *Science of the total environment* 687, 218–231. 2019.
- [22] Zhao, L., Gkountouna, O., Pfoser, D.: Spatial auto-regressive dependency interpretable learning based on spatial topological constraints. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 5(3), 1–28. 2019.
- [23] Tibshirani, R.: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 58(1), 267–288 (1996).
- [24] Hoerl, A.E., Kennard, R.W.: Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12(1), 55–67. 1970.
- [25] Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining*, pp. 785–794. 2016.
- [26] Breiman, L.: Random forests. *Machine learning* 45, 5–32. 2001.
- [27] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11), 2278–2324. 1998.
- [28] Linero AR, Basak P, Li Y, Sinha D. Bayesian survival tree ensembles with submodel shrinkage. *Bayesian Analysis*. 2022 Sep;17(3):997-1020.
- [29] Li, Y., 2021. Bayesian Ensemble Tree Models for Nonparametric Problems (Doctoral dissertation, The Florida State University).
- [30] Mao, S., 2022. Time Series and Machine Learning Models for Financial Markets Forecast (Doctoral dissertation, The Florida State University).

