

# Optimizing Time Series Models for Water Demand Forecasting <sup>†</sup>

Gal Perelman <sup>1,\*</sup> , Yaniv Romano <sup>2</sup> and Avi Ostfeld <sup>1</sup> 

<sup>1</sup> Faculty of Civil and Environmental Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel; ostfeld@technion.ac.il

<sup>2</sup> Faculty of Electrical Engineering and Computer Science, Technion—Israel Institute of Technology, Haifa 32000, Israel; yromano@technion.ac.il

\* Correspondence: gal-p@campus.technion.ac.il

<sup>†</sup> Presented at the 3rd International Joint Conference on Water Distribution Systems Analysis & Computing and Control for the Water Industry (WDSA/CCWI 2024), Ferrara, Italy, 1–4 July 2024.

**Abstract:** This study focuses on optimizing time series forecasting models for water demand in a North Italian city as part of the Battle of the Water Demand Forecast (BWDF) challenge. It aims to accurately predict water demands across ten district-metered areas (DMAs) using historical data and weather information over a one-week horizon. The methodology encompasses data preprocessing, including missing data imputation, feature engineering, and novel normalization techniques, followed by the development and hyperparameter optimization of various data-driven models such as random forest, XGB, LSTM, and Prophet. Extensive cross-validation tests assess each model's performance, revealing that our refined approach markedly enhances forecast accuracy, demonstrating the importance of model and parameter selection for effective water demand forecasting.

**Keywords:** time series forecasting; water demand; data-driven models; data normalization

## 1. Introduction

In the realm of water distribution systems (WDS), ensuring a reliable and efficient water supply stands as a critical challenge, further amplified by global trends such as population growth, urbanization, and climate change. Estimating near-future water demands is crucial for improving WDS management, as it supports various applications such as pump control, maintenance scheduling, leakage identification, and even cyberattack protection [1]. Due to its significant importance, the domain of short-term water demand forecasting has been researched extensively in the past decades through various methods including statistics-based, machine learning, and deep learning approaches [2]. In recent years, advances in automatic meter reading (AMR) and advanced metering infrastructure (AMI) have revolutionized the processes of data collection, storage, and analysis. Consequently, more advanced forecasting methods have been introduced, which incorporate a wider range of data sources and place greater emphasis on accuracy, anomaly detection, and real-time robustness [3]. Despite the comprehensive research conducted on water demand forecasting, selecting the optimal models and fine-tuning their performance continue to present significant challenges. Factors such as local consumption patterns, climate variability, and specific WDS characteristics require a tailored approach to model selection and parameter optimization.

### *Battle of the Water Demand Forecast*

The Battle of the Water Demand Forecast (BWDF) poses the challenge of forecasting the weekly demands of ten district-metered areas (DMAs) in a city in North Italy [4]. The dataset provided includes historic demand data for each DMA, derived from mass balance calculated based on flow measurements, alongside four weather parameters (rainfall depth, air temperature, windspeed, and air humidity). Additionally, key characteristics of each



**Citation:** Perelman, G.; Romano, Y.; Ostfeld, A. Optimizing Time Series Models for Water Demand Forecasting. *Eng. Proc.* **2024**, *69*, 9. <https://doi.org/10.3390/engproc2024069009>

Academic Editors: Stefano Alvisi, Marco Franchini, Valentina Marsili and Filippo Mazzoni

Published: 29 August 2024



**Copyright:** © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

DMA are provided such as the consumption sector (residential, commercial, industrial) and the number of consumers. The battle unfolds over four test periods, with a new batch of data released after each period, such that teams have historical data up to the start of the next test period. The objective is to predict the most accurate forecasts according to three criteria:

- i1—Mean absolute error (MAE) of the first 24 h of the test period.
- i2—Max absolute error (MaxAE) of the first 24 h of the test period.
- i3—MAE of the period between the 24th hour and the rest of the test period.

## 2. Methodology

The proposed methodology follows a comprehensive data-driven workflow, encompassing data preprocessing, feature engineering, hyperparameter optimization, and prediction. The method includes several distinctive features aimed at improving forecast accuracy. The different DMAs vary in demand magnitudes, periodic patterns (daily only, combined daily and weekly), long-term trends, and seasonality; hence, a different model was adjusted for each DMA. Furthermore, acknowledging the distinct accuracy requirements between the initial 24 h period and the remainder of the week, we implemented two specialized forecasting models for each DMA. This approach is rooted in the understanding that models specifically tailored to different time horizons can significantly enhance overall forecasting accuracy. Accordingly, the proposed forecasting architecture was designed to tailor individual models per each DMA and forecasting horizon, resulting in 20 distinct models that were developed and optimized throughout the following stages.

### 2.1. Outlier Detection

The preprocessing stage incorporated outlier detection and filtering using one or more of the following methods: z-score, interquartile range (IQR), and rolling IQR. Additionally, the outlier filtering process included the removal of consecutive identical records, as this often indicates a malfunctioning or “stuck” sensor. Identified outliers were deleted and subsequently addressed during the next phase of missing data completion.

### 2.2. Missing Data Completion

The dataset provided contained approximately 5.5% missing values. To address this issue and ensure the integrity of our forecasting model, a k-NN imputation technique was employed. This method allowed us to accurately estimate and fill in missing data by leveraging the similarity between historical demands and weather observations, thus preserving the underlying structure and patterns in the dataset.

### 2.3. Feature Engineering

A thorough feature engineering process was undertaken, involving the meticulous selection and transformation of relevant variables that significantly influence water demands. The process included the construction of time variables such as hour of the day, day of the week, week of the year, and month. Next, the time variables were converted to periodic waveforms (sin and cos functions) to capture the cyclic nature of these features (hour 23 should be similar to hour 0). Two more time-related features that were introduced are “is special” and “is DST”, which indicate whether a data record is a holiday and whether it is daylight savings time or not. Lagged features were also constructed by shifting previous time steps of the demand and weather variables to capture temporal dependencies and trends. Lastly, moving statistics (moving average and standard deviation) and the decomposition of the data to its trend, seasonal, and noise components were examined, although not necessarily used for all models.

#### 2.4. Data Normalization

A cornerstone of the presented methodology is data normalization. Besides standard normalization methods like standardization and min–max, two temporal dynamic normalization techniques were employed and found effective. The first one was moving window standardization (MWS) and the second was periodic window standardization (PWS). These methods adopt time-dependent means and standard deviations (STDs) for data standardization instead of considering the entire dataset parameters. PWS uses the mean and STD of a corresponding fixed period, which in our case is one week. It ensures that the normalization process respects the temporal boundaries and characteristics unique to each week, enhancing model sensitivity to long-term seasonal and cyclical trends. The MWS dynamically adjusts the normalization parameters based on a sliding window, offering a tailored approach that adapts to recent changes and fluctuations in the data. This adaptability makes the model more responsive to short-term variations and anomalies, potentially improving forecasting accuracy in rapidly changing conditions.

#### 2.5. Models

Multiple data-driven models were developed, including SARIMA, random forest (RF), XGBoost (XGB), a long short-term memory (LSTM) network, and Prophet. Notably, an advanced version of the XGB model was specifically designed for multi-series forecasting. This approach allows the model to forecast for one DMA while incorporating information from other DMAs, leveraging cross-correlations between areas to enhance predictive accuracy. In the Multi-XGB model, each DMA is associated with a cluster of other DMAs to be used as supporting information. Accordingly, several DMA clustering alternatives were tested to find the optimal information base for each DMA.

#### 2.6. Optimization and Cross-Validation

Next, model configurations underwent optimization, including the selection of the prediction algorithm, normalization method, and hyperparameters. To this end, an extensive random search was employed, significantly enhanced by parallelizing numerous experiments, to explore the vast space of potential configurations efficiently. The process was repeated for each test period to fine-tune the configurations, ensuring they were optimally adjusted to the specific period conditions. Overall, a total of 13.8 million train-predict experiments were conducted with a total run time of 440 million seconds, meaning almost 14 years of computations. The best-performing models were identified based on the BWDF metrics. These selected models were then rigorously tested through cross-validation, ensuring their robustness and generalizability. The models that consistently demonstrated superior performance across a diverse set of validation periods were ultimately chosen.

### 3. Results

The optimal models yielded by the above process indicate the superiority of XGB and its multi-series version over the other algorithms that were tested (SARIMA, RF, LSTM, and Prophet). Notably, the temporal dynamic normalization methods, PWS and MWS, were found to improve prediction accuracy compared to standard normalization techniques in most cases. Among the 80 models developed (two for each DMA for each test period), 45 used PWS and 11 used MWS. To demonstrate the performance of the proposed approach, the selected models were validated against four validation periods, corresponding to the last week leading up to each real test period. The results are depicted in Figure 1 along with the averaged metrics of  $i1$ ,  $i2$ ,  $i3$ , and mean absolute percentage error (MAPE) over the four periods. Most DMAs have a MAPE of less than 5%. Additionally, the MAEs of the short and long terms (metrics  $i1$  and  $i3$ ) are in the same magnitude for all DMAs, which underscores the robustness of the proposed approach against accuracy degradation on the longer horizon.




**Figure 1.** Averaged model accuracy across validation periods of the last weeks before test periods.

#### 4. Summary

This study presents a method to predict weekly water demands within the competitive framework of the BWDF. The method employs a set of tailored models for each DMA and respective forecasting horizon. The results underscore the robustness and adaptability of the proposed approach, showcasing its capacity to deliver accurate predictions across diverse DMAs and varying conditions. The achieved accuracy levels meet the requirement for most WDS management applications, illustrating the practical viability of our method for improving WDS efficiency and reliability. While basic forecasting can nowadays be readily produced using numerous available time series analysis packages, the value of a systematic approach cannot be overstated. It refines accuracy beyond the baseline predictions and facilitates ongoing adjustments of models to adapt to dynamic conditions.

**Author Contributions:** Conceptualization, G.P., Y.R. and A.O.; methodology, G.P., Y.R. and A.O.; software, G.P.; validation, G.P.; formal analysis, G.P.; investigation, G.P., Y.R. and A.O.; resources, Y.R.; writing—original draft preparation, G.P.; writing—review and editing, G.P., Y.R. and A.O.; visualization, G.P.; supervision, Y.R. and A.O. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The original data used in the study are openly available at <https://wdsa-ccwi2024.it/battle-of-water-networks/> (accessed on 28 August 2024).

**Acknowledgments:** This research was supported by The Bernard M. Gordon Center for Systems Engineering at the Technion.

**Conflicts of Interest:** The authors declare no conflicts of interest.

#### References

1. Donkor, E.A.; Mazzuchi, T.A.; Soyer, R.; Alan Roberson, J. Urban Water Demand Forecasting: Review of Methods and Models. *J. Water Resour. Plan. Manag.* **2014**, *140*, 146–159. [[CrossRef](#)]
2. Pacchin, E.; Gagliardi, F.; Alvisi, S.; Franchini, M. A Comparison of Short-Term Water Demand Forecasting Models. *Water Resour. Manag.* **2019**, *33*, 1481–1497. [[CrossRef](#)]

3. Niknam, A.; Zare, H.K.; Hosseinasab, H.; Mostafaeipour, A.; Herrera, M. A Critical Review of Short-Term Water Demand Forecasting Tools—What Method Should I Use? *Sustainability* **2022**, *14*, 5412. [CrossRef]
4. Alvisi, S.; Franchini, M.; Marsili, V.; Mazzoni, F.; Salomons, E. Battle of Water Demand Forecasting (BWDF). Available online: [https://wdsa-ccwi2024.it/wp-content/uploads/2024/01/BWDF\\_Instructions\\_rev4.pdf](https://wdsa-ccwi2024.it/wp-content/uploads/2024/01/BWDF_Instructions_rev4.pdf) (accessed on 12 March 2024).

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.