Short-Term Forecasting of Household Water Demand in the UK Using an Interpretable Machine-Learning Approach

Xenochristou, Maria; Hutton, Chris; Hofman, Jan; Kapelan, Zoran

DOI
10.1061/(ASCE)WR.1943-5452.0001325

Publication date
2021

Document Version
Final published version

Published in
Journal of Water Resources Planning and Management

Short-Term Forecasting of Household Water Demand in the UK Using an Interpretable Machine Learning Approach

Maria Xenochristou, Ph.D.[1]; Chris Hutton, Ph.D.[2]; Jan Hofman, Ph.D.[3]; and Zoran Kapelan, Ph.D.[4]

Abstract: This study utilizes a rich UK data set of smart demand metering data, household characteristics, and weather data to develop a demand forecasting methodology that combines the high accuracy of machine learning models with the interpretability of statistical methods. For this reason, a random forest model is used to predict daily demands 1 day ahead for groups of properties (mean of 3.8 households/group) with homogenous characteristics. A variety of interpretable machine learning techniques [variable permutation, accumulated local effects (ALE) plots, and individual conditional expectation (ICE) curves] are used to quantify the influence of these predictors (temporal, weather, and household characteristics) on water consumption. Results show that when past consumption data are available, they are the most important explanatory factor. However, when they are not, a combination of household and temporal characteristics can be used to produce a credible model with similar forecasting accuracy. Weather input has overall a mild to no effect on the model's output, although this effect can become significant under certain conditions. **DOI: 10.1061/(ASCE)WR.1943-5452.0001325.**

Author keywords: Water demand forecasting; Smart demand metering; Random forest.

Introduction

Ensuring water availability for the future is a matter of increasing concern, especially in the context of a rapidly changing world. Understanding water consumption, as well as the drivers behind it, is the first step toward developing accurate demand forecasts and effective water demand management strategies. However, this is a difficult task because household water use reflects many time- and space-dependent factors, and research is often limited by data availability (Parker and Wilby 2013) and privacy concerns.

Jorgensen et al. (2009) reviewed several studies that used social variables to model and predict water consumption and concluded that most of them found different variables to be the most important explanatory factors of consumption. In addition, the explanatory potential of these models was limited, with R^2 (coefficient of determination) values reaching a maximum of ~30% (Jorgensen et al. 2009). This inability of the models to accurately represent consumption might be the reason for the high deviations between them.

Williamson et al. (2002) used a number of property characteristics (e.g., number of residents, appliance ownership, and property type) to predict monthly individual household consumption using statistical regression. This method could distinguish between a large number of households and explained 44% of the variance ($R^2 = 44\%$) in water demand. The rest was attributed to factors that were not included in the model, such as the garden size. However, aggregating consumption at the monthly scale means that temporal variables such as the day of the week cannot be used as explanatory factors. This might limit the amount of variance explained by the model, as well as the opportunity to understand how these variables influence consumption. In addition, for certain applications (e.g., operational requirements for water distribution systems), predictions with higher temporal resolution might be required.

Jorgensen et al. (2014) used a latent growth curve to predict consumption for single-person households over four quarters in 2009 and 2010. In this case, the maximum variance explained (R^2) in the rate of change of water consumption was 31%. This was achieved using three predictors, income, type of irrigation system, and beliefs relating to own consumption. However, accuracy could be improved if more variables were included in the analysis.

Duerr et al. (2018) also developed a water demand forecasting model using property (e.g., land and building value as well as green space), temporal (e.g., month and year), and weather (e.g., temperature and precipitation) characteristics. Several methods were compared for their ability to forecast monthly individual household consumption, including statistical and machine learning models. The one that performed best was the time-series model, with a minimum root-mean square error (RMSE) of 1,246 gal./month (equivalent of an average of 155 L/day), for predictions 1 month ahead. Similarly to previous studies, the level of accuracy is problematic, even when consumption is aggregated at the monthly scale.

Overview and Aim

The benefit of explanatory variables depends on the model's capability to capture the complicated relationships between them and water consumption. In most cases, even when explanatory variables (e.g., household and climatic variables) are utilized to produce water

[1]Ph.D. Student, Centre for Water Systems, Univ. of Exeter, North Park Rd., Exeter EX4 4QF, UK; presently, Postdoctoral Research Fellow, Dept. of Anesthesiology, Perioperative and Pain Medicine, Dept. of Biomedical Data Science, Dept. of Pediatrics, Stanford Univ., Stanford, CA 94305 (corresponding author). ORCID: https://orcid.org/0000-0001-5064-0813. Email: mxenoc@stanford.edu; maria.xenoc@gmail.com

[2]Water Resources Planning Manager, Wessex Water, Claverton Down Rd., Bath BA2 7WW, UK.

[3]Professor, Water Innovation and Research Centre, Dept. of Chemical Engineering, Univ. of Bath, Claverton Down Rd., Bath BA2 7AY, UK; Principal Scientist, KWR Water Research Institute, P.O. Box 1072, Nieuwegein 3430BB, Netherlands. ORCID: https://orcid.org/0000-0002-5982-603X

[4]Professor, Dept. of Water Management, Delft Univ. of Technology, Stevinweg 1, Delft 2628CN, Netherlands; Professor, Centre for Water Systems, Univ. of Exeter, North Park Rd., Exeter EX4 4QF, UK.

Note. This manuscript was submitted on July 5, 2019; approved on September 6, 2020; published online on January 20, 2021. Discussion period open until June 20, 2021; separate discussions must be submitted for individual papers. This paper is part of the ***Journal of Water Resources Planning and Management***, © ASCE, ISSN 0733-9496.

demand forecasts, this is often done using linear regression analysis or geodemographic profiling based on census data (Parker and Wilby 2013). These techniques have traditionally been used because they are simple and able to capture the relationships between the predictors and water demand in a transparent way (Goodchild 2003; Wong et al. 2010). However, their ability to model the complicated relationships between a set of predictors and water consumption may be limited. At the same time, the non-linear and non-univariate effect of some weather variables on water demand, as well as their interactions with other variables that were observed in previous studies (Parker and Wilby 2013; Parker 2014; Xenochristou et al. 2018, 2019, 2020b), require further attention.

Machine learning models are able to provide accurate water demand forecasts (Herrera et al. 2010; Anele et al. 2017; Chen et al. 2017; Zubaidi et al. 2018), but they have been traditionally considered as black box. This means that they are not easy to interpret, and sometimes even their structure and functionality is not well understood. The interpretability of machine learning models is a topic with increasing popularity as more methods are developed (Doshi-Velez and Kim 2017; Adadi and Berrada 2018; Carvalho et al. 2019; Molnar 2019) and find use in different fields, particularly in medical applications (Berk et al. 2016; Choi 2018; Cremona et al. 2018; Carmichael et al. 2019; Culos et al. 2020; Stanley et al. 2020). However, machine learning interpretability methods have not been applied and tested in the field of water demand forecasting. As a result, the ability to use these models to provide guidance to water utilities has been limited.

The aforementioned gaps in knowledge are addressed by developing and presenting a novel approach for water demand forecasting that combines the high accuracy of machine learning models with the interpretability of simpler methods. Combining both accuracy and interpretability is essential in order to produce accurate forecasts and provide water utilities with the knowledge to improve network operations and secure water for the future. Water demand modeling that reconstructs detailed household, temporal, and weather variables would enable planners to predict small area demands and test new tariffs (Clarke et al. 1997). In addition, these variables can enhance the understanding of water-use behaviors and thus support improved demand management practices (Duerr et al. 2018). This is particularly important when the distribution of customer demand is highly skewed, particularly on peak demand days, when a small number of customers are responsible for a high percentage of the total water use. Results of this study would allow demand management strategies to target particular household types (i.e., the types that use the most water) in order to reduce peak demands, which can be valuable during drought periods, as well as improve the understanding of the complicated relationships between weather and water consumption.

In order to achieve this, a machine learning model based on random forests is implemented to predict daily demands for small household groups with homogenous characteristics, with and without past consumption data. Next, three interpretability techniques [variable permutation, accumulated local effects (ALE) plots, and individual conditional expectation (ICE) curves] are used to assess the influence of a variety of household, temporal, and weather variables, as well as their interactions, on the model's predictions.

Data

The data set comprises of water demand data and household characteristics from the southwest of England collected by Wessex Water, one of the UK water companies, as well as weather data provided by the Meteorological Office of the UK (Met Office). A detailed description of each data type is available in this section.

Consumption Data

Water demand data were collected at the household level by the water company using smart meters, recording consumption every 15–30 min over a 3-year period (October 2014–September 2017). The raw data were carefully cleaned and processed before used in any further analysis. A process was implemented consisting of logical rules that aimed to exclude inconsistent or false data while maintaining the natural variability of water demand. More details about this process have been given by Xenochristou et al. (2020b). After the preprocessing of the data, 1,793 properties are included in the data set. Recordings for each property correspond to a maximum duration of 1,019 days, although this number is reduced for most properties due to gaps in the data.

Household Characteristics

The water company also collected household data relating to property and customer characteristics (garden size, rateable value, metering status, council tax band, acorn groups, and occupancy rate), available at the household level. Information about garden sizes and occupancy rates were collected by questionnaires that customers fill in when they want to switch to a smart water meter. The rest of the household properties were collected by their respective agencies.

In order to limit the processing time and reduce complexity, the properties in the data set are grouped in two to three segmentation categories for each household characteristic (Fig. 1). Garden sizes were divided into small (<60 m^2), medium (61–165 m^2), and large (>165 m^2) by the water company. Properties that are classed as unmetered are a representative sample of all unmetered customers in the study area and are not charged based on their meter readings. The water bill of unmetered properties in the UK is adjusted according to the property's rateable value, which is indicative of its rental value and was last updated in the 1970s (Dresner and Ekins 2006). The cutting points for the categories of the rateable value are chosen in order to acquire relatively equal groups that are at the same time distinct enough to identify any differences in their water consumption. The top and bottom 30% of the rateable values are classified as high and low, respectively, whereas the rest are classified as medium.

Acorn is a geodemographic segmentation of the UK's population based on social factors and population behavior (CACI 2014). According to the acorn guide, Consumer groups A, B, and C are classified as Affluent Achievers and Groups D and E as Rising Prosperity (CACI 2014). All Groups A–E are classified as Affluent in the following. Groups F–J are classified as Comfortable Communities, whereas Groups K–Q are Financially Stretched (similar to the same guide). Occupancy rate groups are divided into 1, 2, and 3+, based on the corresponding number of occupants living in each household. The council tax bands are divided into three classes containing Bands A–C, D–E, and F–H, with Class A being the lowest and Class H the highest paying council tax band.

The cutting points of the new categories for the acorn status, occupancy rate, and council tax band are selected based on a z-statistic, according to the following process. Each type of household (e.g., households in Tax band C) is associated with a certain water consumption distribution among all days in the data. A z-statistic is used in order to assess the similarity between the consumption distributions for different types of households (e.g., households in Council tax bands A, B, C, and so on). Similar consumption distributions that are also in close proximity in terms of the physical meaning of their characteristic (e.g., similarly

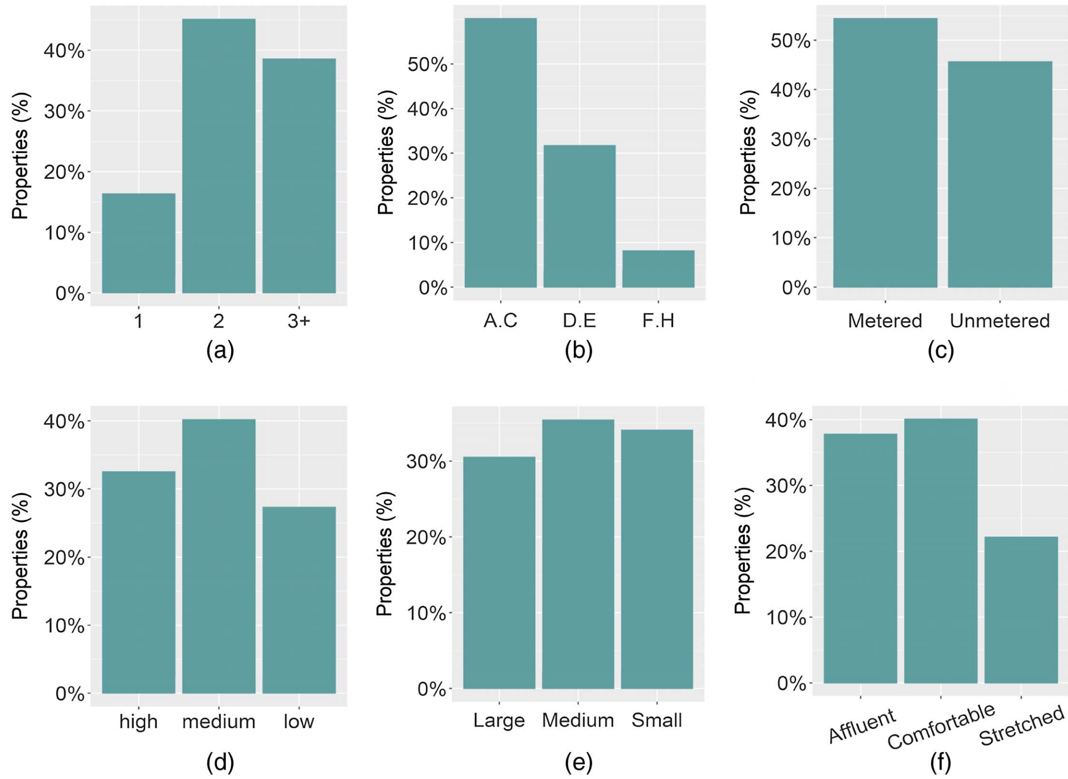

Fig. 1. Percentage of properties in each segmentation category of the six household characteristics: (a) occupancy rate; (b) council tax band; (c) metering status; (d) rateable value; (e) garden size; and (f) acorn group.

paying council tax bands) are grouped together into a larger category (e.g., Council tax bands A–C).

Fig. 1 demonstrates the percentage of properties in each segmentation category among all properties in the study area for each one of the six household characteristics.

Weather Data

The weather data set includes Met Office data on air and soil temperature at 10 cm depth, humidity, sunshine duration, and rainfall. These data are recorded at the hourly or daily scale over the same period (October 2014–September 2017) from hundreds of weather stations across the study area as part of the Met Office Integrated Data Archive System (MIDAS) Land and Marine Surface Stations Data (Met Office 2006a, b, c, d, e). The number of preceding consecutive days without rain is also calculated based on the rainfall data.

Out of the hundreds of weather stations in the study area, only 56 are included in the analysis, based on their proximity to the properties in the data set. Because the properties are scattered over a relatively large area, daily and hourly information from multiple weather stations is used to calculate one daily value for each weather variable as a weighted average of all 56. In order to do this, a weight is assigned to each weather station based on the number of properties that are the closest to it geographically (each property is closest to one of the weather stations). For example, if Weather station A is the nearest weather station to 100 properties and Weather station B is the nearest weather station to 160 properties, Weather station B is assigned a higher weight. Weather stations that have no properties in the nearest proximity are assigned a zero weight. This methodology is adopted in order to account for the location of the weather stations. Instead of calculating a mean value among all stations in the area, the proximity of the stations to the properties in the data set is taken into account. This is likely to result in more accurate estimates of weather values, especially for the weather variables that demonstrate a higher spatial variability.

Methodology

This section outlines the steps of the methodology adopted here in terms of the model variables, the household grouping, the modeling technique (random forests), the model and variable assessment methodologies, and finally the model's technical implementation.

Model Input Variables

The first step toward model building is to define the pool of variables that will be included in the analysis. All available variables are investigated for their influence on the model's results, for forecasts 1 day into the future, grouped into the following four types:

- Past consumption: a 7-day window of past consumption is used to capture the repetitive nature of water use over a calendar week. Past consumption consists of seven values, reflecting mean daily consumption for each one of the 7 days prior to the prediction day. Fig. 2 demonstrates an example of how water consumption, averaged across all properties in the data set, follows a weekly pattern over 2 consecutive weeks, from April 18 until May 2, 2016. In Fig. 2, May 2, which is a Monday, corresponds to unusually high consumption, which is typically characteristic of weekends. This is due to the fact that this day is also a Bank Holiday in the UK;
- Temporal variables: these refer to the season, month, day of the week, and type of day (working day or weekend/holiday) that consumption relates to. They are used as a proxy for time-varying behavioral and weather patterns;

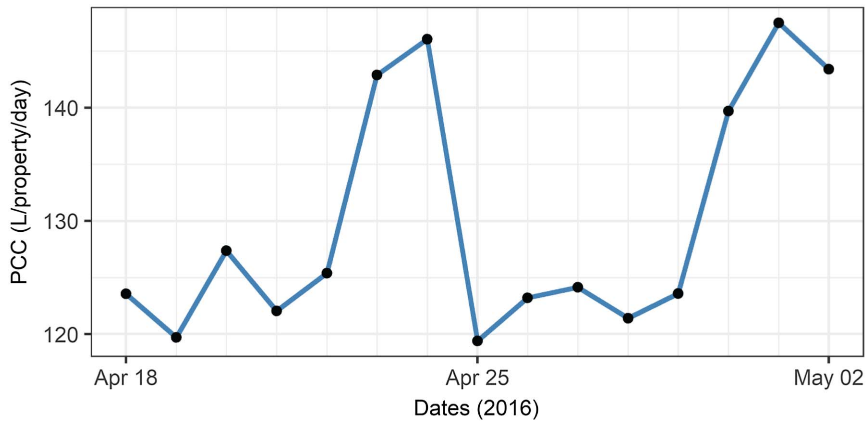

Fig. 2. Mean per capita consumption (PCC) among all properties in the study area for 2 consecutive weeks, between April 18 and May 2, 2016.

- Household characteristics: the six variables collected by the water company, namely garden size, rateable value, metering status, occupancy rate, council tax band, and acorn group, are regularly suspected to influence demand; and
- Weather variables: six variables relating to daily air and soil temperature at 10 cm depth, relative humidity, total sunshine hours and rainfall amount as well as the total number of preceding days without rain are used to account for the weather-induced variance in water consumption.

Household Grouping

In order to maintain the heterogeneity of the original data set, six household characteristics are used in order to create homogenous groups of properties. For example, one group comprises of properties with large gardens, high rateable value, metered consumption, affluent residents, Tax bands A–C and occupancy rate 3+. Because each household characteristic has three to four categories, this results in 3,072 household groups

$$\mathrm{HG}\,(3{,}072) = \mathrm{GS}\,(4) \times \mathrm{RV}\,(4) \times \mathrm{MS}\,(3) \times \mathrm{Acorn}\,(4) \times \mathrm{CT}\,(4) \times \mathrm{OR}\,(4)$$

where HG = household groups; GS = garden size; RV = rateable value; MS = metering status; CT = council tax band; and OR = occupancy rate.

Even though the theoretical number of groups is 3,072, some of the aforementioned household characteristics combinations contain no houses for all or part of the days in the data set (1,019 days in total), whereas others contain only one household. For this analysis, the minimum amount of households allowed in each group is set to two. Each data point represents consumption for a given group and a given day, resulting in 56,020 data points, containing 2–24 households each, or a mean of 3.8 households.

This grouping is adopted in order to reduce the number of data points and the noise in the consumption signal. Instead of having multiple individual households with identical characteristics and high variance in consumption, these are replaced by one representative household, with consumption equal to the mean among all properties in the group. Due to the small size of the final groups and the high variation in their characteristics, daily water consumption varies significantly among days and groups, from ~45 to ~390 L/capita/day, with a mean consumption of 127.4 L/capita/day.

Random Forests

A random forest (RF) model is an ensemble of decision trees that can be used for regression or classification purposes (Breiman 2001). The RF regression used here works by taking a set of input variables, which are then passed onto each of the decision trees in the forest. The uniqueness of a RF model lies in the fact that it implements randomness in the modeling process because at each node, the variable for splitting is chosen among a randomly selected sample of the independent variables (Herrera et al. 2010). Each tree gives a prediction, and the mean of these values is the prediction of the RF.

Hyperparameters in machine learning models are parameters whose values are fixed before the learning process begins. RFs' performance depends on three key hyperparameters, the number of features tested for splitting (mtry), the number of trees that comprise the forest (ntrees), and the tree depth, which can also be specified by the number of end points at each node (nodesize). The maximum number of mtry is equal to the total number of input variables. Reducing the mtry increases the randomness of the trees and reduces processing time whereas reducing the nodesize cause the trees to grow deeper, with the danger of overfitting.

It is commonly believed that default values of these hyperparameters (e.g., mtry = number of variables/3 in regression) can produce good results, although there is no theoretical framework that supports this assumption (Scornet 2017). A search for the optimum set of hyperparameters (mtry, nodesize, and ntrees) confirmed the belief that RFs are fairly robust to changes in hyperparameter values, at least when these are varied within reasonable limits. Thus, the hyperparameter nodesize for the models is set to 200 and the number of trees at 300, although all models are tuned for the optimum value of the mtry parameter.

RFs are chosen because they have been consistently found to outperform most other models in the literature (Chen et al. 2017), but at the same time, they are underrepresented in water demand forecasting (Herrera et al. 2010; Chen et al. 2017; Duerr et al. 2018). In addition, these models are quick to train because the trees are built in parallel and they have limited number of parameters that require tuning.

Model Performance Assessment

The forecasting accuracy of the models is assessed using the following three performance metrics: the mean square error (MSE), the mean absolute percentage error (MAPE), and the R^2 coefficient of determination. These metrics provide a range of information; the

MSE is sensitive to outliers; the MAPE is weighted more toward smaller values and is independent of units and therefore system capacity (Xenochristou 2019); the R^2 indicates the agreement between observed and predicted values.

Each one of these metrics is calculated as follows:

$$\mathrm{MSE} = \frac{1}{n}\sum_{i=1}^{n}(O_i - P_i)^2$$

$$\mathrm{MAPE} = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{O_i - P_i}{O_i}\right|$$

$$R^2 = \left[\frac{\sum_{i=1}^{n}(O_i - \hat{O})(P_i - \hat{P})}{\sqrt{\sum_{i=1}^{n}(O_i - \hat{O})^2 \sum_{i=1}^{n}(P_i - \hat{P})^2}}\right]^2$$

where n = total number of values; O = observed values; P = predicted values; O_i and P_i = ith observed and predicted value, respectively; and $\hat{O}$ and $\hat{P}$ = observed and predicted means, respectively (Xenochristou 2019).

The variable importance is calculated by assessing by how much accuracy drops when a variable is permutated (i.e., rearranged). Permutating a variable means shuffling its values and thus destroying the link between the predictor and the target; therefore, destroying its predictive capability. For example, shuffling the temperature variable would rearrange the temperature values by randomly assigning each one of them to a day in the data set. The MSE of the model is calculated before and after the permutation occurs; the higher the increase in MSE, the higher the importance of the variable that was permutated. The shuffling is repeated several times in order to achieve more accurate results. This process is affected by variable interactions for two reasons. First, correlated predictors mask each other's effect because they provide overlapping information to the model. At the same time, shuffling a variable that is strongly correlated with another one could create unrealistic data points (Molnar 2019). For example, assuming two correlated predictors, air and soil temperature, shuffling the air temperature values could create a day with soil temperature of 4°C and air temperature of 28°C.

The model predictors are evaluated for their impact on the dependent variable, i.e., the water demand, based on two types of interpretable machine learning methods, ALE plots (Apley and Zhu 2020) and ICE curves (Goldstein et al. 2015). In order to explain these methods, it is easier to explain the simpler concept of partial dependence plots (PDPs) first. PDPs work simply by forcing a predictor to take the whole range of its values for each point in the data (each data instance) and calculating the mean response of the model for each value of the predictor. The same happens for categorical predictors, except in this case, the variable is forced to take each one of its potential categories instead of a range of values. PDPs assume noncorrelated variables because in a different scenario, this process could create unrealistic data instances, as explained previously.

ALE plots also describe how a variable affects the prediction on average by calculating the variation in the model's result when varying the values of the predictor within a small window. ALE plots are centered at zero, so the value at each point is the difference to the mean prediction. Apley and Zhu (2020) first introduced ALE plots as a faster and nonbiased alternative to PDPs. ALE plots are used here to assess the influence of the household and temporal characteristics.

ICE plots are the same as PDPs but instead of averaging, ICEs show one curve for each data instance (each day and household group). In other words, an ICE plot shows the response of the dependent variable (daily water consumption) for a change in the independent variable (weather) for each data instance. Because there are 56,020 different groups for all days in the data, the same amount of curves are represented in one plot, which makes it very difficult to distinguish between them. Therefore, these curves are aggregated for each plot into three groups, using k-means clustering (Steinley 2006). The ICE plots are used to capture the varying effect of the weather variables across different types of households and days in the data (Xenochristou et al. 2020b).

More details and explanations regarding these three methods (variable permutation, ICE curves, and ALE plots) have been given by Molnar (2019). All of the preceding analyses were performed using the R version 3.5.0 programming language, particularly the RandomForest (Liaw and Wiener 2002) and iml (Molnar 2019) packages.

Technical Implementation

Because the methods described previously (variable permutation and ICE curves) are affected by variable interactions, the correlations between the predictors need to be assessed. Many household variables are indicative of the socioeconomic status of the household's residents, thus the correlations between them are evaluated using a chi-square (χ^2) test of independence (Table 1). The χ^2 varies between 1 and −1, indicating a perfect positive or negative correlation, respectively. According to Table 1, the council tax band is the most highly interrelated variable. Properties that are under higher paying council tax bands have higher rateable values, larger gardens, and residents with higher socioeconomic status. Properties with larger gardens have a higher rateable value and are occupied by residents in higher acorn groups (Table 1). Although there are clear relationships among the household variables, these were not considered strong enough in order to remove one of them as input.

An investigation into the weather variable interactions (Xenochristou et al. 2020b) showed that sunshine hours and humidity, rainfall, and days without rain, as well as air and soil temperature, are correlated. Temporal variables such as the type of day (working day versus weekend/holiday) and the weekday, as well as the season and the month, are by definition also heavily correlated. Past consumption data are also autocorrelated from one day to the next one.

Table 1. Chi-square correlation statistic between each one of the six household variables

Household characteristics	Garden size	Rateable value	Metering status	Acorn groups	Occupants	Council tax band
Garden size	1	−0.41	0.16	0.33	−0.12	−0.48
Rateable value	−0.41	1	0.09	−0.30	−0.07	0.57
Metering status	0.16	−0.20	1	0.17	0.29	−0.15
Acorn groups	0.33	−0.30	0.17	1	−0.04	−0.58
Occupants	−0.12	0.10	0.29	−0.04	1	0.13
Council tax band	−0.48	0.57	−0.15	−0.58	0.13	1

Table 2. Input variables for Models 1–7

Variable group	Model input variables	1	2	3	4	5	6	7
Past consumption	Consumption 1–7 days ago	X	—	—	—	—	X	—
	Consumption 1 day ago	—	X	—	—	—	—	—
Temporal	Type of day	X	X	X	X		X	X
	Weekday	X	—	X	—	X	—	—
	Month	X	—	X	—	X	—	—
	Season	X	X	X	X	—	—	—
Household	Acorn	X	X	X	X	X	—	X
	Garden size	X	X	X	X	X	—	X
	Metering status	X	X	X	X	X	—	X
	Rateable value	X	X	X	X	X	—	X
	Council tax band	X	X	X	X	X	—	X
	Occupancy rate	X	X	X	X	X	—	X
Weather	Sunshine hours	X	X	X	X		—	—
	Soil temperature	X	—	X	—	X	—	—
	Air temperature	X	X	X	X	—	—	—
	Humidity	X	—	X	—	X	—	—
	Days without rain	X	—	X	—	X	—	—
	Rainfall	X	X	X	X	—	—	—
Total input variables		23	12	16	11	11	8	7

Based on the preceding investigations, two groups of RF models are developed for daily predictions 1 day into the future (Table 2). Models 1, 2, and 6 incorporate past consumption data, whereas Models 3, 4, 5, and 7 use a combination of temporal, household, and weather characteristics. Consumption data are of high interest for two reasons. Firstly, water utilities do not always have access to these data, and therefore it is important to account for this scenario and develop an alternative strategy. Secondly, past consumption incorporates many qualities that are characteristic of the household or the day the consumption corresponds to and therefore can mask the effect of other predictors.

The input variable configuration for Models 1–7 is chosen according to the following. Model 1 (with past consumption) and Model 3 (without past consumption) include all temporal, weather, and household variables. To reveal the influence of each variable without being concealed by overlapping information, Models 2, 4, and 5 exclude strongly correlated inputs (Table 2). Finally, results regarding the most important predictors from Models 1–5 are used to build Models 6 and 7, based on the simplest model configuration that would not compromise the modeling accuracy (Table 2).

In order to start the modeling process, the data set is shuffled and divided randomly into a training set (70% of the data) used to train the models and a test set (30% of the data) used to assess their performance on unseen data, i.e., data that are not used during the model-building phase.

Results and Discussion

Preliminary Analysis

The preliminary data analysis demonstrates how consumption varies across different household and temporal categories. Modeling results can be strongly influenced by interactions among variables as well as the model structure itself. Therefore, it is important to have an initial view of which are the variables with the highest effect on water consumption and test if these conclusions align with the modeling results.

Fig. 3 shows the distribution of consumption for each variable category and each day in the data set. The most distinct difference in consumption is observed when households are grouped based on their occupancy rate, with low-occupancy households (one resident) consuming significantly more per capita compared with high-occupancy ones (three or more residents) [Fig. 3(a)]. Differences also appear between households in different council tax bands [Fig. 3(b)], with houses in Bands A–C (lower council tax bands) consuming less water per capita than houses in Bands F–H (higher council tax bands).

Fig. 3 also shows that distributions of household categories that relate to higher consumption are generally more spread out, whereas the low-consumption curves tend to have a higher peak and a much smaller variance. This is likely because lower consumption constitutes base consumption, i.e., water used in order to perform essential day-to-day activities such as toilet flushing, showering, and cooking. On the other hand, higher demand values and variance, typically found in higher council tax band households, are due to additional, nonbase water consumption activities such as garden watering that occur on some days but not on others. The high variance in the case of the occupancy rate is due to the consumption in single-occupancy properties being more erratic, as it only depends on one person. In the case of two, three, or more residents, the per capita consumption (PCC) is calculated as the mean between the occupants of the property, thus averaging out any differences in consumption behavior from one day to the next one.

Fig. 4 shows the distribution of daily PCC for different categories of four temporal characteristics (month, day of the week, type of day, and season). Demand is time-dependent because it increases during certain times of the week or the year. Consumption is higher over weekends and holidays as opposed to weekdays, with Sundays claiming the highest weekly consumption [Figs. 4(a and d)]. A milder influence is observed throughout the year, with water demand over the summer months and December slightly higher than any other time of the year [Figs. 4(b and c)].

Prediction Accuracy

A summary of the modeling results for the training and test data sets is given in Table 3. Model 6 has the best performance (MAPE = 17.9% and $R^2 = 54.9\%$). Model 7, which does not include data on past consumption, can still explain 49% of the variance in the

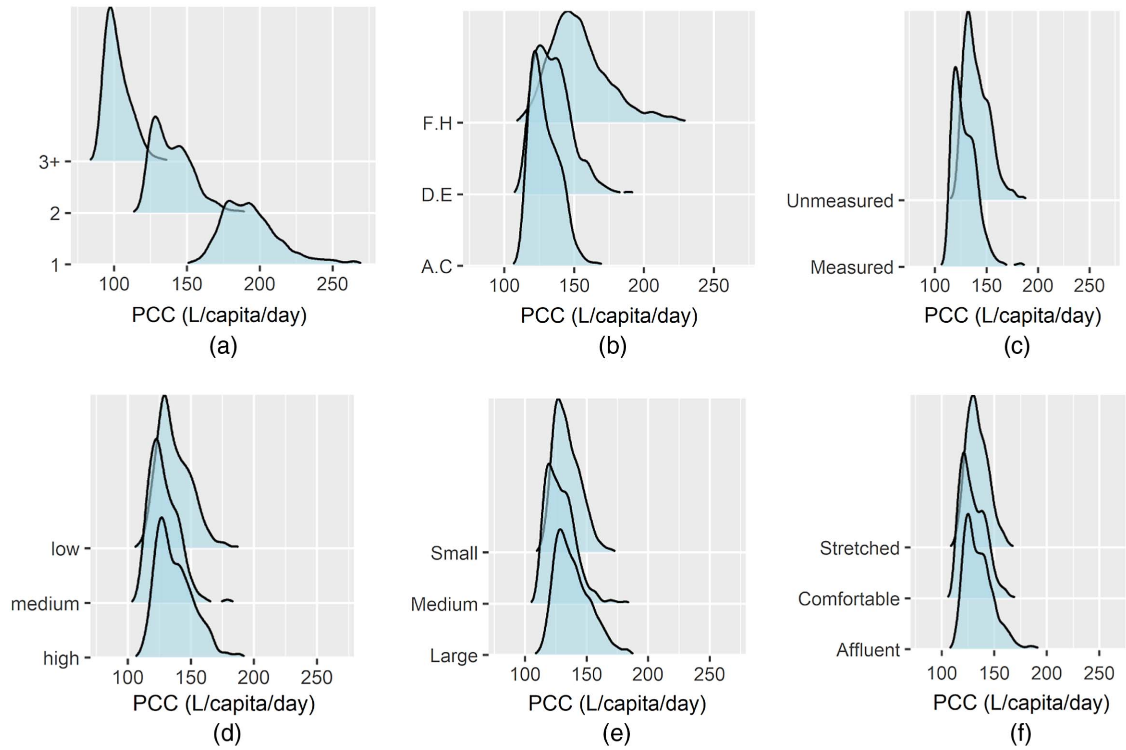

Fig. 3. Distribution of consumption for different categories of 6 household characteristics. Each distribution shows the mean daily PCC among all properties with the corresponding characteristic for each day in the data, for different (a) occupancy rates; (b) council tax bands; (c) metering statuses; (d) rateable values; (e) garden sizes; and (f) acorn groups.

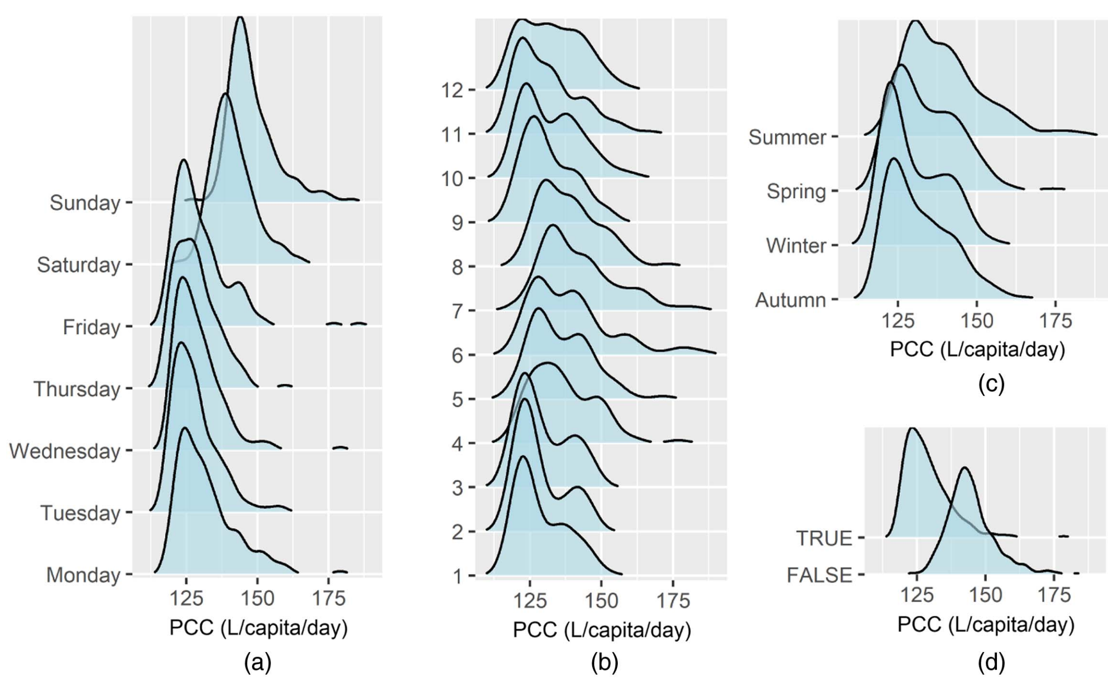

Fig. 4. Distribution of consumption for different categories of temporal characteristics. Each distribution shows the mean daily PCC among all properties for each day in the data for different (a) weekdays; (b) months; (c) seasons; and (d) day types (weekend/bank holiday).

model (MAPE = 19.7% and $R^2 = 49.0\%$). For comparison, the model that assumes water demand for each day in the data is equal to mean demand averaged across all days has a MAPE = 29.3% and $R^2 = 0$. The model that predicts consumption for each household group to be equal to the previous day has a MAPE = 23.7% and $R^2 = 34.9\%$. Unlike the two simple benchmark models, the RF model is able to predict a significant portion of the consumption variance ($R^2 = 54.9\%$), despite the relatively high amount of randomness associated with this level of aggregation. This is due to the RF's ability to learn the consumption patterns in the training data, even in the presence of noise.

Taking into account the small temporal (daily) and spatial (~3.8 household on average) scales for which predictions are made, the models can predict a significant portion of the variance in household consumption despite the amount of randomness associated with this level of aggregation. Previous studies specifically accounted for this effect of spatial scale on prediction accuracy and highlighted the reduction in predictive performance associated with small-scale household consumption (Xenochristou and Kapelan 2020). When predicting household consumption at the monthly scale and household level, previous studies achieved a maximum R^2 of 44% (Williamson et al. 2002), or ~30% (Jorgensen et al. 2009). Therefore, the RF model developed here ($R^2 = 54.9\%$) performs significantly better at a very high temporal (daily scale) and spatial (3.8 households) resolution.

According to Table 3, reducing the number of explanatory variables does not (in most cases) influence the results, whereas in some cases, it even improves the model's accuracy. Removing correlated weather and temporal variables has hardly any effect on the result (Table 3, Models 3–5), whereas excluding 6 days of past consumption from Model 1 leads to increased forecasting errors (Table 3, Model 2). Model 7, which includes only six household variables and the type of day as input, performs better than Model 3, which has additional temporal and weather variables. Removing all variables other than past consumption and the type of day from Model 1 also slightly increases the prediction accuracy (Table 3, Model 6). In both cases, this is likely due to overfitting problems, i.e., the model learning patterns from the variables that do not influence consumption.

Based on the preceding findings, for the purposes of demand prediction, water utilities do not necessarily need to rely heavily on extensive smart metering programs over the whole network, although there are potential benefits of smart metering data beyond demand forecasting. These benefits include reduced consumption, leakage detection, and deriving a greater understanding of household water consumption for individual water users. In terms of demand forecasting, smaller-scale metering programs may be sufficient to develop useful predictive models that could then be upscaled with data on customer and property characteristics. This finding is particularly valuable for water utilities in the UK, where almost half of the properties are unmetered and overall smart meter penetration is significantly lower.

Variable Permutation

One variable is permutated at a time for each model, and results appear in Fig. 5 (models with past consumption) and Fig. 6 (models without past consumption). The x-axis demonstrates the importance factor, i.e., the factor by which the MSE increases (denoting decline in model performance), when an input variable is permutated. The variables are ranked on the y-axis based on this importance factor. Because the shuffling is repeated multiple times in order to increase the robustness of the outcome, several importance factors are calculated for each variable. The error bar corresponds

Table 3. Model configuration and prediction accuracy for Models 1–7

		Model parameters			Training			Testing		
Models	Consumption data	mtry	nodesize	ntrees	MAPE (%)	MSE (L^2/day^2)	R^2 (%)	MAPE (%)	MSE (L^2/day^2)	R^2 (%)
1	Yes	5	200	300	16.1	742	64.3	17.9	952	54.7
2	Yes	4	200	300	18.1	936	54.7	19.0	1,055	50.0
3	No	8	200	300	18.7	983	53.1	19.7	1,115	47.6
4	No	6	200	300	19.3	1,027	51.3	20.0	1,132	47.3
5	No	5	200	300	19.1	1,014	52.0	19.8	1,126	47.5
6	**Yes**	**3**	**200**	**300**	**16.7**	**809**	**61.0**	**17.9**	**934**	**54.9**
7	**No**	**3**	**200**	**300**	**19.6**	**1,069**	**48.5**	**19.7**	**1,067**	**49.0**

Note: Two best models, i.e., with and without past consumption data, are highlighted in bold.

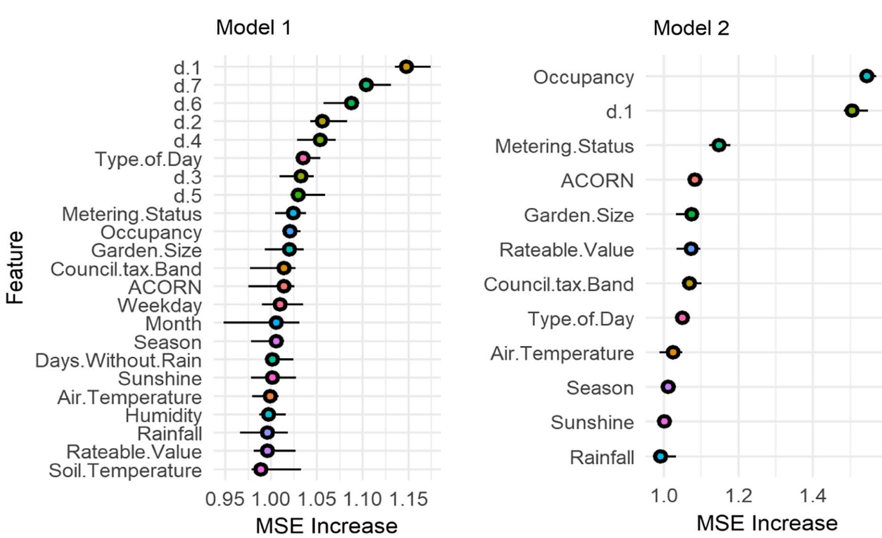

Fig. 5. Factor by which the mean square error (MSE) increases when each feature is permutated for Models 1 and 2.

J. Water Resour. Plann. Manage., 2021, 147(4): 04021004

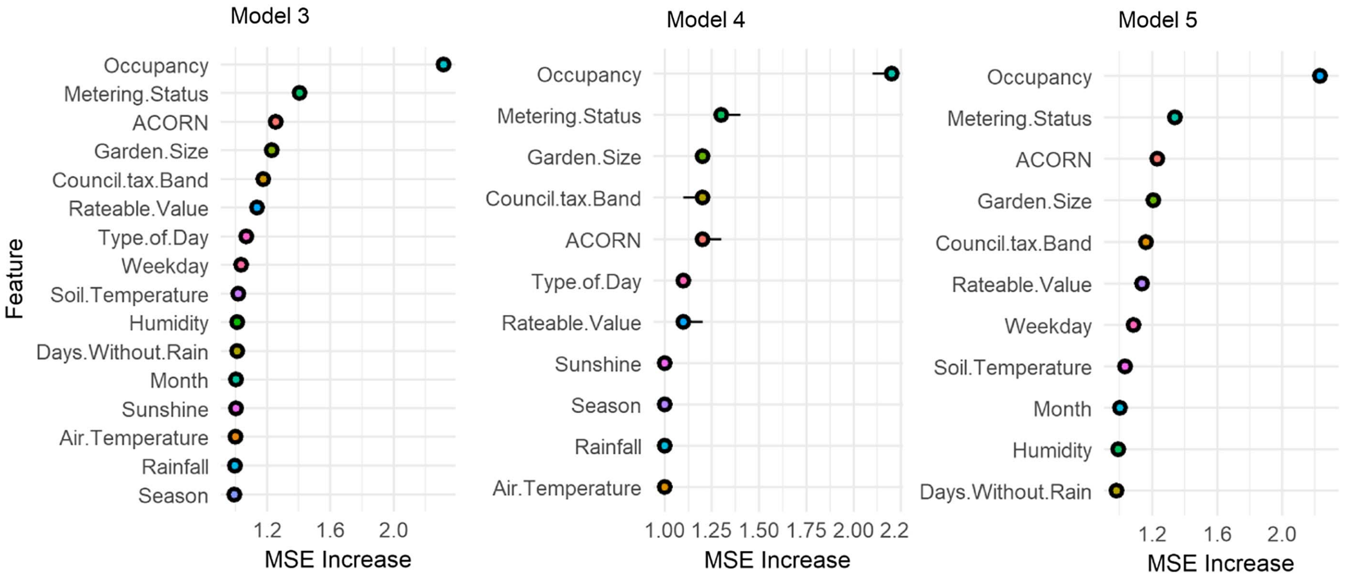

Fig. 6. Factor by which the mean square error (MSE) increases when each feature is permutated for Models 3–5.

to the importance at 5% and 95% of the repetitions, whereas the dot corresponds to the median. A factor of 1 means that excluding the variable from the model does not influence accuracy.

According to Fig. 5, when 7 days of past consumption are included as model input, they are by far the most important predictors (Fig. 5, Model 1). Demand 1 day in the past (d.1) has the highest explanatory value, followed by demand on the same day of the week but 7 days prior (d.7). The day of the week is the only other important variable, whereas the other predictors have a mild to zero influence. However, even when the variable with the highest importance (d.1) loses its predictive capacity, the MSE increases only by a factor of 1.15. Because Model 1 already includes 7 days of past consumption that carry overlapping information, excluding any one of them individually does not have a major effect on the output.

However, things are different for Model 2 (Fig. 5), which excludes highly correlated predictors. In this case, both consumption 1 day ago (d.1) and the occupancy rate are highly important, and excluding either from the model increases the MSE by a factor of 1.50–1.53 (i.e., by 50%–53%), a much higher rise compared with Model 1. In addition, the significance of the rest of the household characteristics as well as the type of day also increases (Fig. 5, Model 2).

Fig. 6 demonstrates the same results when past consumption data are not used as input (Models 3–5). In this case, household characteristics, particularly the occupancy rate, are the most important predictors, followed by temporal information (type of day or weekday) (Fig. 6). Similarly to Fig. 5, all other variables, including the weather and the rest of the temporal characteristics, are very close to a factor of 1. This means that even when past consumption is not included as model input, excluding these variables from the model does not influence accuracy.

Although there are slight differences among Models 3–5 (Fig. 6), the importance factors relating to each predictor are very similar. Removing correlated predictors (e.g., the season, month, and various weather variables) from Models 4 and 5 in this case did not increase their importance.

Notably, there is a large difference in the scale of feature importance between Fig. 5 (with past consumption) and Fig. 6 (without past consumption). When the explanatory factors contain overlapping information, excluding one of them only marginally reduces accuracy, resulting in low feature importance factors (Fig. 5). When information about past consumption is not available, the occupancy rate is the only variable carrying this information, resulting in an importance factor of up to 2.3 (Fig. 6, Model 3). This means that excluding information about the occupancy rate of a household when past consumption is not available will increase the MSE ~2.3 times or 130%.

These findings provide a good overview of variable importance and interactions and can be used as a guide on what variables to include in the model under different conditions, i.e., based on what other relevant information is available in each case.

Influence of Household Variables

Next, the effect that different household characteristics have on the predictions is explored using the ALE plots (Fig. 7). The y-axis shows different categories of each explanatory variable, and the x-axis demonstrates the deviation from the mean predicted consumption for each household category (Fig. 7). When the ALE value of the x-axis is positive, the corresponding category is predicted to have a consumption higher than average, whereas the opposite is true when the ALE value is negative.

As can be seen from Fig. 7, the results are in agreement with previous analysis that explored the distribution of consumption for each household category (Fig. 3). Occupancy has by far the highest influence on predicted consumption, as properties with low occupancy rate (one resident) are predicted to consume ~75 L/capita/day of water more than properties with high occupancy (three or more residents) [Fig. 7(a)]. The next most influential variable is the council tax band [Fig. 7(b)]. Higher paying bands (F–H) have a predicted consumption of ~26.5 L/capita/day more than lower bands (A–C), and unmetered customers are also on the higher end, with ~19.5 L/capita/day more than metered customers [Fig. 7(c)]. A smaller influence is identified for the acorn group, garden size, and rateable value. Financially stretched customers have the highest predicted consumption, which is ~9 L/capita/day more than customers in the comfortable acorn group [Fig. 7(f)]. Properties with large gardens are predicted to consume ~5 L/capita/day more than the ones with small gardens [Fig. 7(e)], whereas properties with high rateable values are

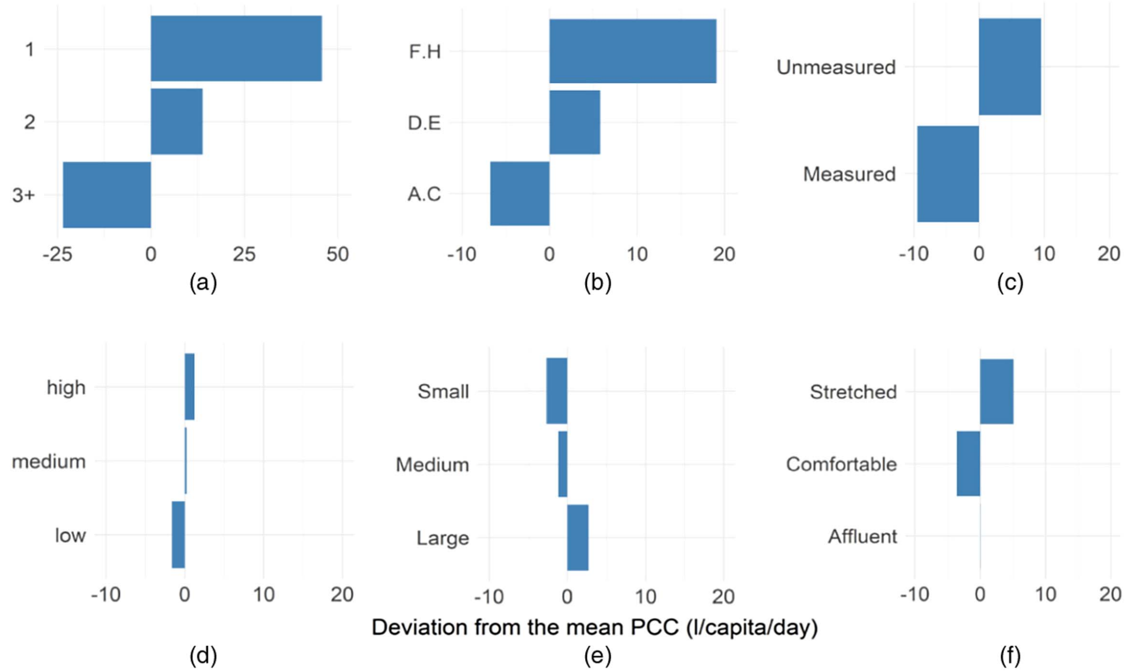

Fig. 7. Influence of 6 household characteristics on predicted water consumption, accumulated local effects (ALE) plots: (a) occupancy rate; (b) council tax band; (c) metering status; (d) rateable value; (e) garden size; and (f) acorn group.

predicted to consume ~3.5 L/capita/day more than the low ones [Fig. 7(d)].

These results are in general agreement with studies in the literature (Russac et al. 1991; Edwards and Martin 1995; Bellfield 2001; Butler and Memon 2006). Edwards and Martin (1995) concluded that lower acorn classes are associated with lower PCC, whereas other studies (Russac et al. 1991; Bellfield 2001) found no strong relationship between the acorn group and water use. Although some studies (Russac et al. 1991) observed that as the rateable value increases, so does water consumption, others (Bellfield 2001) did not find any relationship between the two. Finally, the relationship between the garden size and water consumption has been so far difficult to establish (Bellfield 2001; Gato et al. 2007).

Influence of Temporal Variables

The effect of four temporal characteristics on the model's result is also investigated using the ALE plots (Fig. 8). According to Fig. 8, the type of day (if it is a weekend/bank holiday or not) and day of the week have the highest impact on predicted water demand, whereas the month and season have almost no influence. Overall, water consumption on weekends and holidays is predicted to be ~11 L/capita/day higher than on working days [Fig. 8(c)]. Water demand gradually declines from Monday to Friday, to then increase again on Saturday and Sunday. Sundays claim almost 8 L/capita/day more on average compared with Fridays, the day with the lowest predicted consumption [Fig. 8(a)]. Although the month and season have almost no influence on the model's result, summers cause a slight increase in consumption (<1 L/capita/day). An even smaller influence is observed for December (<0.5 L/capita/day), the month associated with the highest increase in predicted consumption. This is likely due to the holiday season because people tend to spend more time at home.

Time variations in water-use patterns are widely recorded in the literature (Edwards and Martin 1995; Hartley 1995; Kowalski and Marshallsay 2005; Gato et al. 2007; Billings and Jones 2008; Parker and Wilby 2013). Water use is higher in the weekends because this is when people tend to be more regularly at home (Edwards and Martin 1995; Hartley 1995; Bellfield 2001; Gato et al. 2007; Parker and Wilby 2013). Typically, water use peaks over the summer months, although lower peaks have also been observed over the winter (Billings and Jones 2008; Parker and Wilby 2013). However, in a temperate climate like the UK with lack of strong seasonality and rainfall well distributed over the year, it is expected that the seasonal pattern is going to be weaker than in other countries.

Influence of Weather Variables

The influence of four weather variables on the model's response, i.e., the daily water consumption, is assessed using the ICE plots (Fig. 9). Previous work (Xenochristou et al. 2020b) concluded that the rainfall amount and soil temperature have a limited effect on water demand; thus, only the ICE curves corresponding to air temperature, humidity, sunshine duration, and days without rain are presented here. To avoid significant interactions from correlating weather predictors, only one weather variable at a time is considered as model input when creating the ICE plots, along with past consumption data and the type of day. For each plot in Fig. 9, the *y*-axis represents the change in PCC compared with the mean, when the variable of interest (in this case one of the four weather variables), varies within its whole range of values (*x*-axis). The percentage associated with each curve represents the percentage of data points that belong to each cluster.

According to Fig. 9, the weather variable that causes the biggest spike in water consumption is air temperature [Fig. 9(a)]. This effect is nonlinear and becomes significant when temperature exceeds approximately 18°C and to a lesser extent for near-freezing temperatures, which is likely due to water used to prevent pipes from freezing (Billings and Jones 2008) or leakages between the meter and the property. Although water consumption starts increasing for temperatures over the 18°C threshold, the rate of increase varies

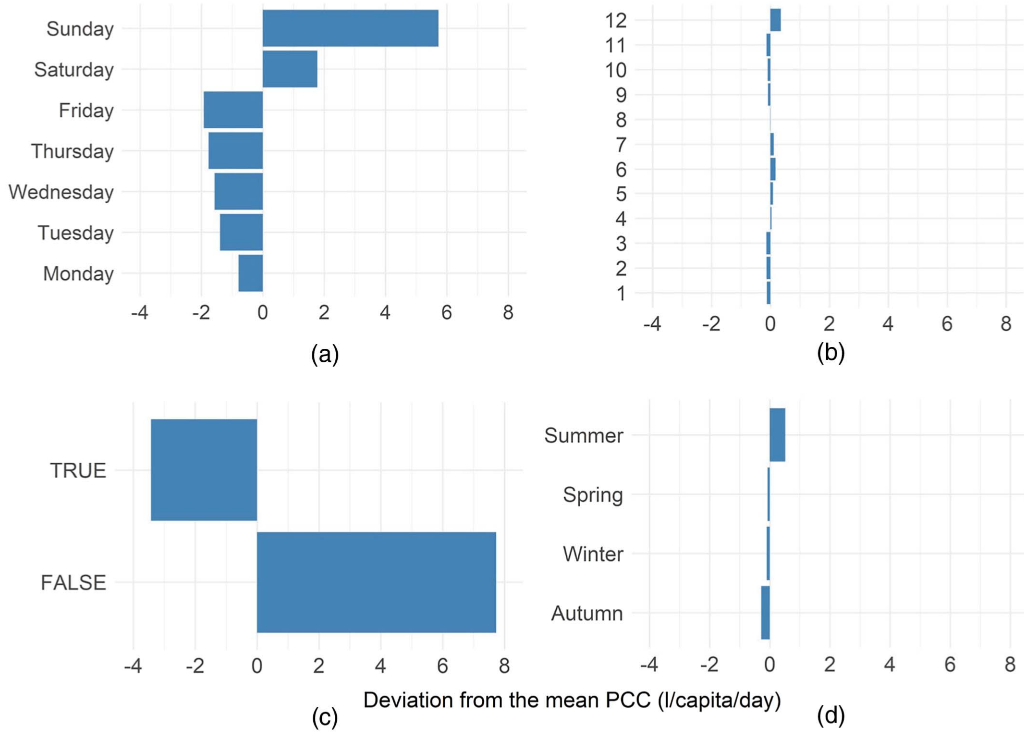

Fig. 8. Influence of 4 temporal characteristics on predicted water consumption, accumulated local effects (ALE) plots: (a) day of the week; (b) month; (c) type of day; and (d) season.

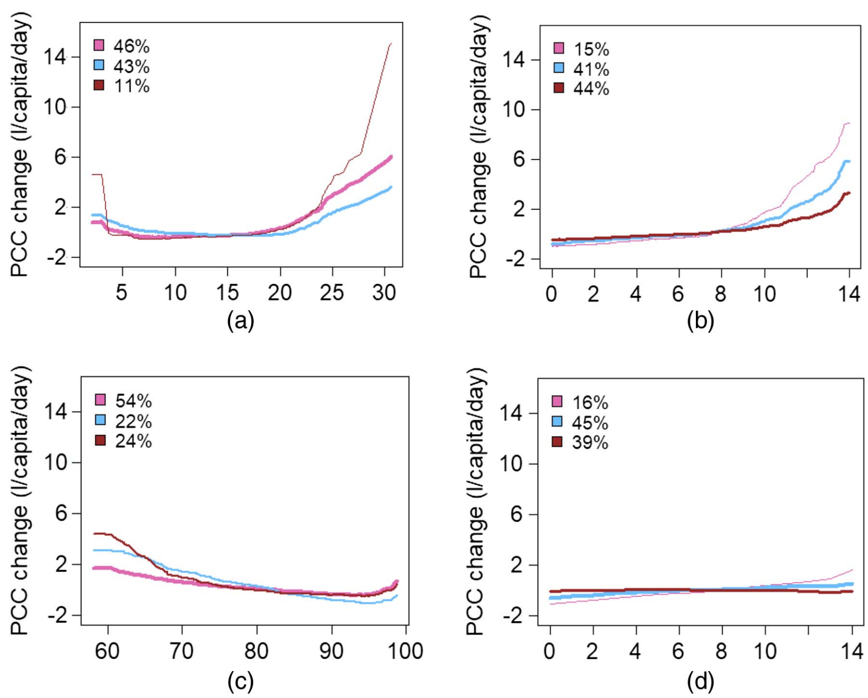

Fig. 9. Influence of 4 weather variables on predicted water consumption, individual conditional expectation (ICE) plots: (a) air temperature; (b) sunshine duration; (c) relative humidity; and (d) days without rain. The percentage associated with each curve represents the percentage of data points that belong to each cluster.

significantly [Fig. 9(a)]. Different days and households have different sensitivity to weather changes. Here, only for 11% of data instances (one data instance is 1 day and one household group), the model predicts an increase in water use of up to 15 L/capita/day for an increase in air temperature from 18°C to 30°C. For the other 89% of the days and household types, the predicted increase in consumption is between 2.5 and 6.0 L/capita/day [Fig. 9(a)].

For the rest of the weather variables, the predicted increase in consumption is lower than for air temperature. The maximum increase in water consumption caused by sunshine duration is

9 L/capita/day, which is 6 L/capita/day lower than for air temperature, but this increase relates to 15% of data instances. The relative humidity has an even smaller effect, with a maximum change of 4 L/capita/day. However, this change applies to ~54% of all days and household types, whereas for 22% of them, there is a near-steady decline over the whole range of humidity values [Fig. 9(c)]. For the rest 24% of data points, water consumption drops by 4 L/capita/day for an increase in humidity from 60% to 70%, whereas it only marginally decreases after this point.

The number of consecutive days without rain has the smallest effect on the predictions. Consumption starts increasing after 12 days without rain, reaching a maximum increase of 3 L/capita/day, for 16% of data points. This could potentially cause problems in the future if the length of droughts increase. There is also a marginal but steady increase in consumption over the whole range of days without rain (x axis) for 61% of the data points, whereas there is no visible change for 39% of the days and households.

In previous studies, the temperature (Bellfield 2001; Parker and Wilby 2013; Dos Santos and Pereira 2014), sunshine hours (Bellfield 2001), and humidity (Dos Santos and Pereira 2014), have all been found to influence water demand, whereas the rainfall amount had a lesser effect on water consumption (Bellfield 2001; Schleich and Hillenbrand 2008). One reason that could explain this low impact of weather on prediction accuracy could relate to the mild UK climate, which lacks seasonal extremes. In this region, household demand uplifts associated with the weather are typically in the order of 5% during hot summer periods; thus, weather induced demand is overall limited. Even more so, the years included in this study did not capture a particularly hot dry summer. During the record summer temperatures of 2018, the nonlinear influence between weather and demand was seen at a broader aggregation—e.g., from DMA to company level. Therefore, stronger weather effects could have been observed if the analysis included 2018 data.

Another reason for the limited weather effect could be the small size of household groups (a mean of 3.8 properties/group). At this level, the noise in the consumption signal might be too strong to allow for the subtle changes due to weather to show. Previous work showed that the effect of weather becomes noticeable only for certain households, days, and times (Xenochristou et al. 2020b), as well as for certain aggregation levels (Xenochristou et al. 2020a).

Summary and Conclusions

This study has demonstrated a novel approach that combines the high accuracy of machine learning models with the interpretability of statistical methods. As part of this work, a RF model is developed that predicts daily water consumption 1 day ahead for homogenous groups of properties (~3.8 households/group). A variety of interpretable machine learning techniques (variable permutation, ALE, and ICE curves) is used in order to assess the contribution of the predictors on the forecasting accuracy and predicted water consumption.

Based on the results obtained, the following conclusions can be drawn:

- The RF-based short-term demand forecasting model is able to accurately capture the complex and nonlinear dependencies between water consumption and different explanatory variables such as temporal, household, and weather characteristics.
- When past consumption is not available, credible forecasting models can be developed using household and temporal characteristics, but weather input does not further improve results. The best-performing forecasting model in this case is the one including six household variables (occupancy rate, council tax band, metering status, rateable value, acorn, and garden size) as well as the type of day as inputs.
- When past consumption is not available, the property's occupancy rate is the most influential input variable, followed by the council tax band and metering status. The acorn group, garden size, and rateable value have the smallest effect. The weekly pattern of consumption also becomes evident. Weekends and holidays have a higher predicted consumption compared with working days, although the monthly and seasonal patterns are very weak.
- When past consumption data are included in the demand forecasting model, no other variable can significantly improve the prediction results. The best-performing model in this case is the one using 7 days of past consumption and the type of day as inputs.
- Although weather input does not improve the forecasting accuracy, relationships are identified between water consumption and air temperature, sunshine duration, humidity, and to a lesser extent for days without rain. This influence, however, is limited to only certain household groups and days in the data, and in most cases, it is triggered when the weather variable exceeds a certain threshold. This nonlinearity is important to identify and is relevant to help understand and predict changes in household consumption under potential changes in the UK climate.

These results help identify the factors that can explain consumption variability among households. Thus, they may assist with effectively targeting water conservation strategies, testing new tariffs, and assessing the impact of population and lifestyle changes, as well as evaluating the effect of potential changes in the climate at the household level. In addition, this methodology can lead to the development of improved water demand forecasting models and enhance the usefulness of machine learning models even when past consumption is not available.

The same methodology can be adopted and applied in different studies in order to determine the predictors of water demand with respect to the characteristics of each individual case. However, the results of each study are specific to and dependent on its individual characteristics that can relate to environmental factors such as climatic variables, as well as household characteristics, customs, and habits, and the interactions among them. Therefore results should always be interpreted within the context of the specific case study.

In addition, this work uses a certain level of temporal (daily) and spatial (~3.8 households/group) aggregation. The small temporal and spatial scales implemented here allow to maintain the heterogeneity of the data set and account for the influence of the different household, temporal, and weather variables, as well as their interactions, on the model's output. However, this choice might have influenced the results. Increasing the level of spatial aggregation decreases the range of demand values and thus it reduces forecasting errors. In addition, variable importance also changes at different aggregation levels (Xenochristou et al. 2020a).

Finally, the RF model was selected for this analysis due to its accuracy and ease of implementation. However, forecasting accuracy may further improve if a different model is used instead. The performance of RFs with respect to the characteristics of the problem, such as the temporal and spatial scale, forecast horizon, and data availability, compared with other machine learning models, has been the topic of further work (Xenochristou and Kapelan 2020).

Data Availability Statement

Some or all data, models, or code used during the study were provided by a third party. Direct requests for these materials may be made to the provider as indicated in the Acknowledgments.

Acknowledgments

This study was funded as part of the Water Informatics Science and Engineering Centre for Doctoral Training (WISE CDT) under a grant from the Engineering and Physical Sciences Research Council (EPSRC) (Grant No. EP/L016214/1). The data for this study were made available by Wessex Water.

References

Adadi, A., and M. Berrada. 2018. "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)." *IEEE Access* 6: 52138–52160. https://doi.org/10.1109/ACCESS.2018.2870052.

Anele, A. O., Y. Hamam, A. M. Abu-Mahfouz, and E. Todini. 2017. "Overview, comparative assessment and recommendations of forecasting models for short-term water demand prediction." *Water* 9 (11): 887. https://doi.org/10.3390/w9110887.

Apley, D. W., and J. Zhu. 2020. "Visualizing the effects of predictor variables in black box supervised learning models." *J. R. Stat. Soc., Ser. B* 82 (4): 1059–1086. https://doi.org/10.1111/rssb.12377.

Bellfield, S. 2001. "Short-term domestic water demand: Estimation, forecasting and management." Ph.D. thesis, Faculty of Environment, School of Geography, Univ. of Leeds.

Berk, R. A., S. B. Sorenson, and G. Barnes. 2016. "Forecasting domestic violence: A machine learning approach to help inform arraignment decisions." *J. Empirical Legal Stud.* 13 (1): 94–115. https://doi.org/10.1111/jels.12098.

Billings, R. B., and C. V. Jones. 2008. *Forecasting urban water demand*. 2nd ed. Denver: American Water Works Association.

Breiman, L. 2001. "Random forests." *Mach. Learn.* 45 (1): 5–32.

Butler, D., and F. A. Memon. 2006. *Water demand management*. London: International Water Association Publishing.

CACI (Consolidated Analysis Center, Incorporated). 2014. "The acorn user guide." Accessed November 29, 2020. https://acorn.caci.co.uk.

Carmichael, I., et al. 2019. "Joint and individual analysis of breast cancer histologic images and genomic covariates." Accessed January 15, 2019. https://arxiv.org/abs/1912.00434.

Carvalho, D. V., E. M. Pereira, and J. S. Cardoso. 2019. "Machine learning interpretability: A survey on methods and metrics." *Electronics* 8 (8): 832. https://doi.org/10.3390/electronics8080832.

Chen, G., T. Long, J. Xiong, and Y. Bai. 2017. "Multiple random forests modelling for urban water consumption forecasting." *Water Resour. Manage.* 31 (15): 4715–4729. https://doi.org/10.1007/s11269-017-1774-7.

Choi, E. 2018. "Doctor AI: Interpretable deep learning for modeling electronic health records." Ph.D. thesis, School of Computational Science and Engineering, Georgia Institute of Technology.

Clarke, G. P., A. Kashti, A. McDonald, and P. Williamson. 1997. "Estimating small area demand for water: A new methodology." *Water Environ. J.* 11 (3): 186–192. https://doi.org/10.1111/j.1747-6593.1997.tb00114.x.

Cremona, F., L. Tuvikene, J. Haberman, P. Nõges, and T. Nõges. 2018. "Factors controlling the three-decade long rise in cyanobacteria biomass in a eutrophic shallow lake." *Sci. Total Environ.* 621 (Apr): 352–359. https://doi.org/10.1016/j.scitotenv.2017.11.250.

Culos, A., et al. 2020. "Integration of mechanistic immunological knowledge into a machine learning pipeline improves predictions." *Nat. Mach. Intell.* 2: 619–628. https://doi.org/10.1038/s42256-020-00232-8.

Doshi-Velez, F., and B. Kim. 2017. "Towards a rigorous science of interpretable machine learning." Preprint, submitted February 28, 2017. https://arxiv.org/abs/1702.08608.

Dos Santos, C. C., and F. A. J. Pereira. 2014. "Water demand forecasting model for the metropolitan area of São Paulo, Brazil." *Water Resour. Manage.* 28 (13): 4401–4414. https://doi.org/10.1007/s11269-014-0743-7.

Dresner, S., and P. Ekins. 2006. "Design of environmentally and socially conscious water metering tariffs for the UK." *J. Environ. Plann. Manage.* 49 (6): 909–928. https://doi.org/10.1080/09640560600947048.

Duerr, I., H. R. Merrill, C. Wang, R. Bai, M. Boyer, M. D. Dukes, and N. Bliznyuk. 2018. "Forecasting urban household water demand with statistical and machine learning methods using large space-time data: A Comparative study." *Environ. Modell. Software* 102 (Apr): 29–38. https://doi.org/10.1016/j.envsoft.2018.01.002.

Edwards, K., and L. Martin. 1995. "A methodology for surveying domestic water consumption." *Water Environ. J.* 9 (5): 477–488. https://doi.org/10.1111/j.1747-6593.1995.tb01486.x.

Gato, S., N. Jayasuriya, and P. Roberts. 2007. "Temperature and rainfall thresholds for base use urban water demand modeling." *J. Hydrol.* 337 (3–4): 364–376. https://doi.org/10.1016/j.jhydrol.2007.02.014.

Goldstein, A., A. Kapelner, J. Bleich, and E. Pitkin. 2015. "Peeking inside the black box: Visualizing statistical learning with plots of individual conditional 166 expectation." *J. Comput. Graphical Stat.* 24 (1): 44–65. https://doi.org/10.1080/10618600.2014.907095.

Goodchild, C. 2003. "Modelling the impact of climate change on domestic water demand." *Water Environ. J.* 17 (1): 8–12. https://doi.org/10.1111/j.1747-6593.2003.tb00423.x.

Hartley, J. A. 1995. "A neural network and rule based system application in water demand forecasting." Ph.D. thesis, Dept. of Electronic and Computer Engineering, Brunel Univ.

Herrera, M., L. Torgo, J. Izquierdo, and R. Pérez-García. 2010. "Predictive models for forecasting hourly urban water demand." *J. Hydrol.* 387 (1–2): 141–150. https://doi.org/10.1016/j.jhydrol.2010.04.005.

Jorgensen, B., M. Graymore, and K. O'Toole. 2009. "Household water use behavior: An integrated model." *J. Environ. Manage.* 91 (1): 227–236. https://doi.org/10.1016/j.jenvman.2009.08.009.

Jorgensen, B. S., J. F. Martin, M. W. Pearce, and E. Willis. 2014. "Predicting household water consumption with individual-level variables." *Environ. Behav.* 46 (7): 872–897. https://doi.org/10.1177/0013916513482462.

Kowalski, M., and D. Marshallsay. 2005. "Using measured microcomponent data to model the impact of water conservation strategies on the diurnal consumption profile." *Water Sci. Technol. Water Supply* 5 (3–4): 145–150. https://doi.org/10.2166/ws.2005.0094.

Liaw, A., and M. Wiener. 2002. "Classification and regression by randomForest." *R News* 2 (3): 18–22.

Met Office. 2006a. *MIDAS: UK daily rainfall data*. Oxford, UK: NCAS British Atmospheric Data Centre.

Met Office. 2006b. *MIDAS: UK daily temperature data*. Oxford, UK: NCAS British Atmospheric Data Centre.

Met Office. 2006c. *MIDAS: UK daily weather observation data*. Oxford, UK: NCAS British Atmospheric Data Centre.

Met Office. 2006d. *MIDAS: UK hourly weather observation data*. Oxford, UK: NCAS British Atmospheric Data Centre.

Met Office. 2006e. *MIDAS: UK soil temperature data*. Oxford, UK: NCAS British Atmospheric Data Centre.

Molnar, C. 2019. "Interpretable machine learning." Accessed November 29, 2020. https://christophm.github.io/interpretable-ml-book/.

Parker, J. M. 2014. "Assessing the sensitivity of historic micro-component household water-use to climatic drivers." Ph.D. thesis, Geography and Environment, Loughborough Univ.

Parker, J. M., and R. L. Wilby. 2013. "Quantifying household water demand: A review of theory and practice in the UK." *Water Resour. Manage.* 27 (4): 981–1011. https://doi.org/10.1007/s11269-012-0190-2.

Russac, D. A. V., K. R. Rushton, and R. J. Simpson. 1991. "Insights into domestic demand from a metering trial." *Water Environ. J.* 5 (3): 342–351. https://doi.org/10.1111/j.1747-6593.1991.tb00628.x.

Schleich, J., and T. Hillenbrand. 2008. "Determinants of residential water demand in Germany." *Ecol. Econ.* 68 (6): 1756–1769. https://doi.org/10.1016/j.ecolecon.2008.11.012.

Scornet, E. 2017. "Tuning parameters in random forests." In Vol. 60 of *ESAIM: Proc., and Surveys*, 144–162. Les Ulis, France: EDP Sciences. https://doi.org/10.1051/proc/201760144.

Stanley, N., et al. 2020. "VoPo leverages cellular heterogeneity for predictive modeling of single-cell data." *Nat. Commun.* 11: 3738. https://doi.org/10.1038/s41467-020-17569-8.

Steinley, D. 2006. "K-means clustering: A half-century synthesis." *Br. J. Math. Stat. Psychol.* 59 (1): 1–34. https://doi.org/10.1348/000711005X48266.

Williamson, P., G. Mitchell, A. McDonald, and P. Wattage. 2002. "Domestic water demand forecasting: A static microsimulation approach." *Water Environ. Manage.* 16 (4): 243–248. https://doi.org/10.1111/j.1747-6593.2002.tb00410.x.

Wong, J., Q. Zhang, and Y. D. Chen. 2010. "Statistical modeling of daily urban water consumption in Hong Kong: Trend, changing patterns, and forecast." *Water Resour. Res.* 46 (3). https://doi.org/10.1029/2009WR008147.

Xenochristou, M. 2019. "Water demand forecasting using machine learning on weather and smart metering data." Ph.D. thesis, College of Engineering, Mathematics, and Physical Sciences, Centre for Water Systems, Univ. of Exeter.

Xenochristou, M., M. Blokker, I. Vertommen, J. Urbanus, and Z. Kapelan. 2018. "Investigating the influence of weather on water consumption: A Dutch case study." In *Proc., WDSA/CCWI 2018 Joint Conf.*, 26–35. Sheffield, UK: Univ. of Sheffield.

Xenochristou, M., C. Hutton, J. Hofman, and Z. Kapelan. 2020a. "Water demand forecasting accuracy and influencing factors at different spatial scales using a Gradient Boosting Machine." *Water Resour. Res.* 56 (8): e2019WR026304. https://doi.org/10.1029/ 2019WR026304.

Xenochristou, M., and Z. Kapelan. 2020. "An ensemble stacked model with bias correction for improved water demand forecasting." *Urban Water J.* 17 (3): 212–223. https://doi.org/10.1080/1573062X.2020.1758164.

Xenochristou, M., Z. Kapelan, and C. Hutton. 2020b. "Using smart demand metering data and customer characteristics to investigate the influence of weather on water consumption in the UK." *J. Water Resour. Plann. Manage.* 146 (2): 04019073. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001148.

Zubaidi, S. L., S. K. Gharghan, J. Dooley, R. M. Alkhaddar, and M. Abdellatif. 2018. "Short-term urban water demand prediction considering weather factors." *Water Resour. Manage.* 32 (14): 4527–4542. https://doi.org/10.1007/s11269-018-2061-y.