Contents lists available at ScienceDirect

Environmental Modelling and Software

journal homepage: www.elsevier.com/locate/envsoft

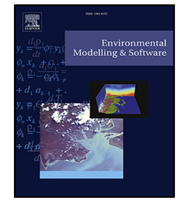

Smart meters data for modeling and forecasting water demand at the user-level

Jorge E. Pesantez [a,*], Emily Zechman Berglund [a], Nikhil Kaza [b]

[a] Department of Civil, Construction and Environmental Engineering, North Carolina State University, Raleigh, NC, USA

[b] Department of City and Regional Planning, Environment, Ecology & Energy Program, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

A R T I C L E I N F O

Keywords:
Smart water meters
AMI
Forecasting model
Hourly water demand
User-level data
Water demand management
Machine learning
Urban water systems

A B S T R A C T

Smart meters installed at the user-level provide a new data source for managing water infrastructure. This research explores the use of machine learning methods, including Random Forests (RFs), Artificial Neural Networks (ANNs), and Support Vector Regression (SVR) to forecast hourly water demand at 90 accounts using smart-metered data. Demands are predicted using lagged demand, seasonality, weather, and household characteristics. Time-series clustering is applied to delineate data based on the time of day and day of the week, which improves model performance. Two modeling approaches are compared. Individual models are developed separately for each meter, and a Group model is trained using a data set of multiple meters. Individual models predict demands at meters in the original data set with lower error than Group models, while the Group model predicts demands at new meters with lower error than Individual models. Results demonstrate that RF and ANN perform better than SVR across all scenarios.

1. Introduction

Designing and operating water distribution systems rely on models and forecasts of water consumption. Water managers use operational or short-term water demand forecasts, ranging from one day to a few weeks, to efficiently manage devices, such as pumps and valves. Short-term forecasting models are based on data collected at the account-level using data that have historically been collected at monthly or quarterly intervals, corresponding to billing cycles. Recently, municipalities and utilities have deployed smart water meters in the context of the smart cities paradigm (Berglund et al., 2020), providing new data about account-level demands at hourly or sub-hourly frequencies. Some utilities use a time resolution for reporting demands in the range of 15 s to 30 min. Battery life and transmitting issues, however, limit the frequency of data collection, and other utilities collect water demand at an hourly frequency (Beal and Flynn, 2015). The use of smart meter data for developing models for forecasting demands at sub-hourly or hourly frequencies is limited to date. Unlike the energy sector, in which smart meters have been extensively deployed to forecast consumption as part of Advanced Metering Infrastructure (AMI) projects (Kavousian et al., 2013), the water sector has not benefited from the development and use of models that forecast water demand at the account or user-level with high temporal resolution.

AMI can be deployed by water utilities to gain insight into water consumption at high spatial and temporal resolutions and to implement advanced capabilities for water management (March et al., 2017; Stewart et al., 2018). AMI provides the technology to collect big data about water consumption and to communicate unusual water consumption to consumers for identifying post meter water leaks (Giurco et al., 2010; Luciani et al., 2019). Smart meter data have been used to support the development of water demand management policies (Cominola et al., 2015), near real-time water distribution system models (Arandia-Perez et al., 2014; Gurung et al., 2017), and enhanced hydraulic and water quality models (Gurung et al., 2014; Creaco et al., 2017b). AMI data was also used to develop descriptive water demand models that were applied to identify appliance-level end uses (Cardell-Oliver, 2013; Nguyen et al., 2014; Gurung et al., 2015), to determine demand patterns for daily consumption profiles and hourly peak values, (Beal and Stewart, 2014; Cominola et al., 2018b), and to group households with similar consumption behaviors (Cardell-Oliver et al., 2014). Forecasting models were developed to predict water consumption at the next time step using data from smart meters that were placed at District Metered Areas (DMAs) or at main pipes, with sub-minute reporting frequencies (Brentan et al., 2018b; Donkor et al., 2014). These models were developed using lagged demands or past consumption as predictor variables (Romano and Kapelan, 2014; Chen and Boccelli, 2018), in addition to exogenous variables, such as weather variables and social characteristics (Sebri, 2016; Hussien et al., 2016). Data reported at

* Corresponding author.
E-mail address: jpesant@ncsu.edu (J.E. Pesantez).

https://doi.org/10.1016/j.envsoft.2020.104633
Received 6 April 2019; Received in revised form 10 January 2020; Accepted 13 January 2020

monthly and annual frequencies at the account-level have been analyzed to explore the effects of pricing on water consumption (Arbues et al., 2003) and the effects of rebate programs on the adoption of low-flow appliances (Price et al., 2014). Account-level data collected at high (sub-hourly) resolution allows researchers to parameterize residential water demand models (Alvisi et al., 2014; Gargano et al., 2016; Kofinas et al., 2018; Creaco et al., 2016), model water quality (Blokker et al., 2008), analyze water end-uses (Blokker et al., 2009; Buchberger and Wells, 1996; Creaco et al., 2017a; Mostafavi et al., 2018), develop models to describe demand (Gurung et al., 2017), develop conservation policies (Maas et al., 2017), and evaluate feedback strategies to customers about their water consumption (Sonderlund et al., 2016).

Account-level data collected at medium resolution can be used to develop forecasting models. Forecasting models can be used within a portfolio of management tools to identify leaks, explore water restriction policies during water supply interruptions, and design demand-management strategies to reduce peak demands (Monks et al., 2019). The high spatial resolution (at the user level) and temporal resolution of smart meter data increase variability and the presence of zero-valued data points in the data set. These characteristics lead to difficulties in forecasting water consumption (Cominola et al., 2018a). Variability emerges in the data due to factors including diverse end-uses, seasonality, and socio-economic conditions (Boyle et al., 2013). As a result, a limited number of studies have used account-level consumption data collected at a medium temporal resolution to develop models to forecast or classify water demand (Aksela and Aksela, 2011; Walker et al., 2015; McKenna et al., 2014; Candelieri, 2017). In this research, new forecasting models are developed using machine learning (ML) methods for hourly smart water data. ML has shown promising results for building predictive models for high resolution demands (Savic et al., 2014). Unlike mechanistic regression models, ML techniques do not require the definition of an explicit relationship between water consumption and independent variables. ML methods have been applied to model and forecast water demand using traditionally available water demand data, such as billing records of water consumption at monthly or quarterly time steps (Jain and Ormsbee, 2002; Adamowski, 2008; Odan et al., 2012; Duerr et al., 2018). Support Vector Regression (SVR), Artificial Neural Networks (ANNs), and Random Forests (RFs) have been applied to model water demands at aggregate levels, such as the system and the DMA levels (Mouatadid and Adamowski, 2017; Gagliardi et al., 2017; Antunes et al., 2018). ANNs have performed better than traditional regression models in forecasting short-term water demands (Bougadis et al., 2005). Herrera et al. (2010) compared the performance of different models to forecast hourly water demand collected at a DMA in Spain and found that SVR, ANNs, and RFs performed similarly well to predict water demand.

The research presented here tests the application of three ML models, including RFs, ANNs, and SVR, to forecast hourly water demand based on the information provided by smart meters. Data were retrieved from a set of 90 smart meters located in Cary, North Carolina, that reported hourly consumption in increments of 10 gallons per hour (gph) [1 gallon = 3.78 l] for a 12-month period in 2017. Time series clustering is explored to improve the accuracy of forecasts by creating separate models for distinct hours of the day. Models are explored with observed demands, the inclusion of weather, and social variables as predictor variables. Results demonstrate that RF and ANN models perform better than SVR in accurately predicting water demands. Time series clustering improves model predictability. Two modeling approaches are compared. In the first approach, ML methods are developed for each smart meter separately to explore the level of predictability. In the second approach, the entire data set of smart meters is used to train a model to forecast water demand at any meter. Results demonstrate that the performance of the two modeling approaches is relatively similar; however, the individual models show a slightly lower error for existing meters, and the Group model predicts demands at new meters with lower error than individual models. The models developed through this research provide new tools for water management by providing demand forecasts at both existing and new accounts.

2. Background

Four research studies explored the development of forecasting and classification models using smart water meter data. These studies analyze hourly data collected at the account-level using Gaussian Mixture models (Aksela and Aksela, 2011; McKenna et al., 2014), a coupled evolutionary algorithm and ANN approach (Walker et al., 2015), and a Support Vector Machine model (Candelieri, 2017). These studies applied clustering to reduce the variability in data sets by grouping vectors of data based on consumption (Aksela and Aksela, 2011; McKenna et al., 2014) or time of day (Candelieri, 2017). To further reduce variability, McKenna et al. (2014) excluded weekends from the data set. The four data sets described by these studies varied in size: data were collected at 81 meters over a three-month period (Aksela and Aksela, 2011), 85 meters over a six-month period (McKenna et al., 2014), nine meters over a two-month period (Walker et al., 2015), and 26 meters over a four-month period (Candelieri, 2017). These studies tested the time of day, lagged demand, and average consumption in a range of forms as demand predictors. Aksela and Aksela (2011) used average weekly consumption to forecast water demand a week ahead, whereas McKenna et al. (2014) classified daily demand patterns. Walker et al. (2015) explored model inputs including the reported water demand at the previous hour, the average consumption of the prior seven days, and the time of the day to improve model performance. Candelieri (2017) developed models to use the first six hours of consumption as predictors for the remaining 18 h of a day. All the models showed the importance of clustering water demand and showed the use of some form of lagged demands (previous demands) in developing predictive models.

Related research also explored how alternative predictors affect hourly or sub-hourly demands. Research demonstrated that aggregate (e.g., system-level) water consumption is significantly correlated with weather data (Praskievicz and Chang, 2009; House-Peters et al., 2010) and property characteristics (Aitken et al., 1994). Whereas the models described in the paragraph above used only lagged water demands to forecast future demands, recent work tested the correlation of weather and property characteristics parameters with account-level hourly water demand reported by smart meters (Xenochristou et al., 2019). More than 1500 smart meters reporting consumption for a 20-month period were analyzed to test predictors, including the building area, number of occupants, household income, and maximum daily temperature for correlation with hourly water demand. Results demonstrated that water consumption is correlated to air temperature, especially during working days in the spring and summer seasons. Further research explored additional weather variables for correlation and demonstrated that precipitation did not influence demands (Xenochristou et al., 2018).

3. Methods and materials

This section describes the procedure applied to forecast water demand one hour ahead at the user-level using hourly data reported by smart meters. The first subsection describes smart meter data and exogenous predictor variables, including weather data and characteristics of the households. The second subsection describes the time-series clustering approach that is applied to group data based on time of day and day of the week, and the third subsection describes the ML methods that are applied to forecast water demand.

3.1. Data

3.1.1. Water consumption data

The data used in this study are collected from a set of 100 smart meters that report hourly consumption at the account-level. The data set represents a small sample of an ongoing project which has already installed more than 60,000 smart meters in Cary, North Carolina, USA (Town of Cary, NC). The accounts represented in this data set are

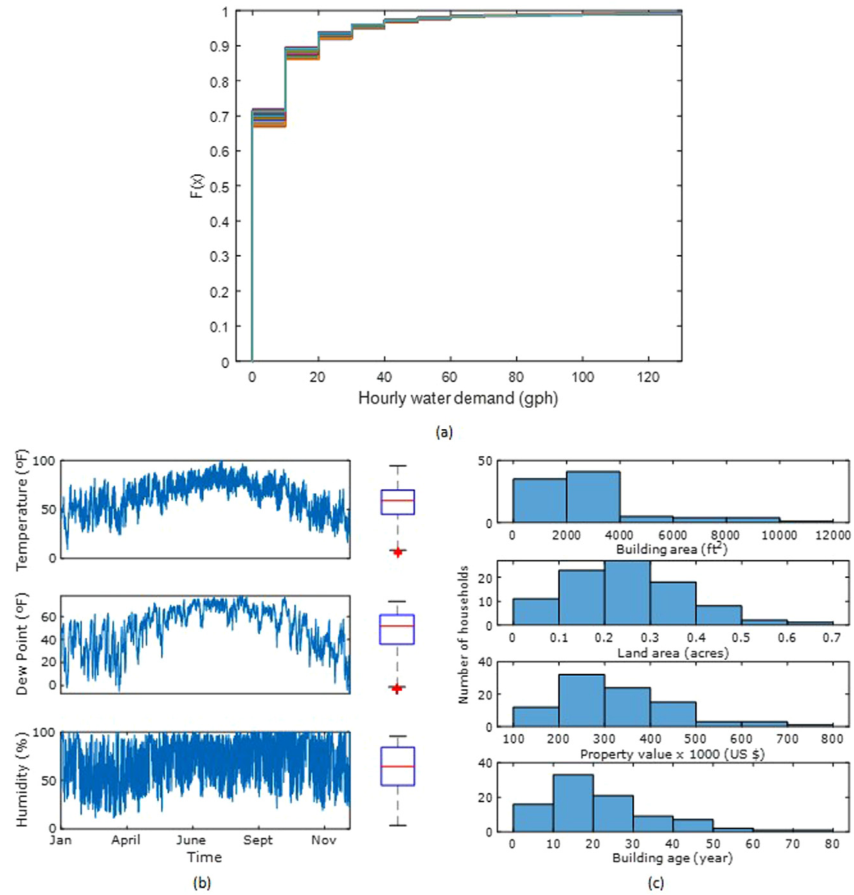

Fig. 1. Data sets used in developing ML models. (a) An empirical cumulative distribution function demonstrates hourly water consumption for each of 85 smart meters. Colors that can be seen in the online version are associated with different meters. The x-axis is restricted to 0 to 130 gph for visualization, (b) weather variables, and (c) histogram of households characteristics.

located throughout the city and include residential and non-residential users. Water consumption was captured hourly during a 12-month period starting on January 1st of 2017. An iPerl meter, produced by Sensus, was installed at each house with a Smart Point transmitter and FlexNet data transmission technology. Meters transmit the hourly data using a data logger to the water utility after a four-hour period. To preserve battery life, the meters report data only when the hourly water consumption of the building is at least 10 gallons. The hourly water demand of the set of smart meters is characterized by high variability. The median of the data set is zero, and the meters periodically report hourly consumption values greater than 10 gallons. An empirical cumulative distribution function is used to represent the water consumption across all meters (Fig. 1a). As shown in Fig. 1a, the majority of data points are zero for all meters, which creates difficulties in predicting hourly demands. The variability of the data set also contributes to challenging issues in modeling water demand.

Preliminary analysis was performed to preserve data continuity. The maximum number of data points reported by each meter is 8760 (number of hours in 2017), and a meter was excluded from the analysis if it reported zero consumption for more than 720 consecutive hours (approximately a month), based on the assumption that those meters correspond to an empty building or a malfunctioning data logger. Ten smart meters were discarded as a result of this screening process. Of the 90 remaining meters, ten meters are associated with accounts that are not single-family residential accounts. Data from 85 meters are used in training, validating, and testing the models. The data from the remaining five meters, which are all associated with residential accounts, are used in testing the models for performance at new meters.

Outliers were identified using two criteria. During one week of the observed period, most of the data loggers stopped transmitting information due to a wide-scale power outage. When meters resumed working, the data loggers reported cumulative values of water consumption, generating misleading peaks. These outliers were removed from the data set. The second criterion applied a threshold consumption of 500 gph, based on the experience of the utility. The threshold value was exceeded by a minor number of data points (less than 5%), and no meters were discarded due to this criterion.

3.1.2. Weather data

Weather data were retrieved from the nearest meteorological station located at Raleigh–Durham international airport (RDU), which is approximately seven miles (12.5 km) away from the location of the households. Weather data that are included in this research are the last 24 h of air temperature (°F) [32°F = 0 °C], dew point (°F), relative humidity (%), the maximum hourly temperature of the last 24 h (°F), hourly precipitation of the last 24 h (in) [1 in = 25.4 mm], and the occurrence of precipitation during the last 24 h as a binary variable. Hourly precipitation during 2017 in Cary reports mostly zero values with some hours reporting up to one inch of rain. The time sampling resolution of the weather variables is hourly (Fig. 1b).

3.1.3. Household characteristics

A set of data was collected describing household characteristics, including the building area (ft^2) [1 m^2 = 10.76 ft^2], the lot area (acres) [1 acre = 4047 m^2], the building age (years), and the property value (U.S. $). Data were retrieved from a GIS database made available through the Town of Cary (Town of Cary, 2013). Variation in household characteristics represented by the data set is shown in Fig. 1c.

3.2. Time series clustering

Similar to the work presented by Candelieri (2017), this study uses a clustering algorithm to group the average hourly water demand of the households into clusters that are based on the time of day. The K-means clustering algorithm (Lloyd, 1982) was applied to the average consumption at each hour of the day, where the average is calculated across all meters. The number of clusters was explored for its effects on the model performance, and Silhouette analysis was used to quantitatively assess the most efficient number of clusters (Arbelaitz et al., 2013). The k-means++ algorithm (David and Vassilvitskii, 2007) was used to initialize cluster centers, which improves the running time of the K-means algorithm, and a random seed was used to initialize random clusters for a set of simulations. Clusters were generated using the sum of absolute differences as the distance metric. To account for variability due to the random generation of cluster centers, the clustering subroutine was run multiple times, and the most commonly repeated clusters were selected.

3.3. Machine learning methods

Three ML methods were applied to forecast short-term water demand. RFs, ANNs, and SVR were implemented using methods available through MATLAB 2019a (mathworks.com). Due to the range of the magnitude of the variables, the input data were normalized for the ANN models. Min–Max scaling was applied to normalize the data and convert the data into a range of −1 and +1 values. For the SVR models, data were standardized by centering and scaling each column of the predictor set using the mean and the standard deviation of each predictor, respectively. For the RF model, input data were used directly.

3.3.1. Random forests

RFs are represented by ensembles of decision trees, which are expanding structures of nodes with the application of binary splits (Breiman et al., 1984). Each node represents a predictor variable. The initial value of each node is the average of the response variable over all the observations of that variable. Splits are formed by using the inequality condition, and the performance of a split is evaluated through the Gini index, which measures how diverse the data are until a terminal node is reached. The tree size is determined based on the number of nodes, which is used as input to minimize the variance of each split. RFs are created through a bootstrap aggregation (bagging) process (Breiman, 1996), and data are re-sampled randomly with replacement. The use of ensemble modeling can improve the overall performance of the model, though overfitting and complexity issues may emerge (Breiman, 2001).

Algorithmic settings for RF models include the number of trees in an ensemble and the leaf size, which is the minimum number of observations per terminal node (Herrera et al., 2010; Villarin and Rodriguez-Galiano, 2019). RFs are modeled by joining several individual decision trees, which has proved to provide better results in terms of accuracy. An ensemble of decision trees creates a model that can be considered a "gray-box", where understanding which parameters lead to a good performance is more difficult, compared to models that are built using a single decision tree. The leaf size refers to the number of observations evaluated at each node, where a low number of observations will generate deep trees that may overfit data. On the other hand, a high number of observations may lead to poor model performance. These settings affect model accuracy and the computing time.

3.3.2. Artificial neural networks

ANNs are widely used in water systems applications (Adamowski, 2008; Herrera et al., 2010; Romano and Kapelan, 2014). ANNs are modeled after the human brain to simulate the mechanisms of human neurons to collect, analyze, and transmit information through different layers (Haykin, 2008). This study uses a feed forward neural network with input, hidden, and output layers. The neurons of the input layer correspond to each of the predictor variables. The output layer represents the response variable, which is the forecasted water demand, and the hidden layer nodes apply the activation function, bias component, and weights to transform the input data. Mathematically, the process is described as:

$$y_i = f(\sum_{j=1}^{m}(w_{ij}x_j) + b_i) \tag{1}$$

where y_i is the response variable; i is the corresponding data point; f represents the activation function, which typically corresponds to an S-shaped function; m is the number of inputs; w_{ij} is the weight applied to the jth input signal; x_j is the jth input signal (predictor value); and b_i represents the bias applied to the data at the i data point.

A back-propagation algorithm is used to iteratively adjust the connection among neurons, bias, and weights and improve the value of the Mean Squared Error (MSE) that is calculated based on the modeled and the observed values in the training process. ANN settings that should be identified are the number of hidden layers and the number of neurons within each layer. Hidden layers separate nonlinear data to improve the predictive capabilities of the model. Each hidden layer uses Eq. (1) in the learning process, and the weights and bias terms are stored in the neurons of each hidden layer. The number of neurons in the hidden layer generally corresponds to the double of the number of predictors (Cutore et al., 2008). More predictors increase the number of neurons and the complexity of an ANN model.

3.3.3. Support vector regression

SVR applies a transformation or kernel function to map a nonlinear data set into a linear function in a high dimensional feature space (Haykin, 2008). Gaussian, linear, and polynomial functions can be used within SVR to transform data. The mathematical formulation of SVR is represented as (Smola and Scholkopf, 2004):

$$f(x) = \langle w, \phi(x) \rangle + b \tag{2}$$

where $f(x)$ is the response value that should fall within the bandwidth defined by an allowable margin (ϵ) for all modeled data. The support vectors define the feature space from $-\epsilon$ to $+\epsilon$. ϕ is the mapping function to transform non-linear relations into linear functions. In Eq. (2), $\langle w, \phi(x) \rangle$ represents the dot product of the weight vector (w) and the transformed input data set ($\phi(x)$), and b is the bias applied to the function approximation. A convex optimization problem is solved to identify the smallest value of the error between modeled and observed data. A cost or box constraint controls with a positive numeric value the penalty for outputs that lie outside of the allowable margin (ϵ) and helps to prevent overfitting. To apply SVR, the user should select the type of transformation or kernel function (e.g., linear, polynomial or Gaussian function), the value of the cost or box constraint, the bandwidth margin (ϵ), and the kernel scale. Scaling the kernel function plays an important role on the performance of SVR, as explained by Williams et al. (2005). In this research, the kernel scale is included as a setting to evaluate model performance.

4. Time clustering results

The K-means clustering algorithm was applied to cluster the average hourly water demand based on the time of day. The number of clusters was selected based on two criteria. First, a quantitative approach with different numbers of clusters (2, 3, and 4) was applied and evaluated using the Silhouette analysis as described in Section 3.2. Two clusters reported the highest Silhouette value (0.72), similar to results reported by Candelieri (2017), who also used two clusters to group water demand data. The use of two clusters also agrees with the emergent characteristics of daily water demands. That is, water consumption data typically follows a diurnal pattern in which two peak values occur: one peak in the morning around 7 AM and another peak in the early evening around 6 PM (Adamowski, 2008; Buchberger et al., 2017). The smart meter data set also demonstrates these peaks in the average daily demand in the morning and evening. During the weekdays, the first cluster starts at 3 AM and continues through the 9 AM hour (7 h). The peak consumption as shown in Fig. 2a occurs at 7 AM, with an average consumption of up to 15 gph between 7 AM and 8 AM on Wednesdays. For weekdays, the evening peak occurs at 8 PM, except for Thursday and Friday, where the peak consumption is at 5 pm. The second cluster begins at 10 AM and continues to 2 AM (17 h). On weekend days, the clusters and the peaks change. The morning peak consumption occurs at 10 AM. Cluster 1 begins at 2 AM and continues to 10 AM (9 h), and the duration of Cluster 2 is 15 h. Similarly, the night peak consumption during the weekends (as shown in Fig. 2b) occurs at 9 PM with a pronounced difference between Saturday and Sunday.

5. Modeling results

5.1. Experimental design

The predictive model relies on several predictors to forecast water demand one hour ahead at the user level. A set of experiments was conducted to test the importance of these predictor variables, data clusters, and size of data sets on model performance. Three sets of predictor variables are created to forecast water demand, grouped as demand and seasonality (DS) variables, weather (W) variables, and property characteristics variables (CH) (Table 1). Four input data sets are created with data grouped in alternative sets of clusters (Table 2). The first input (all_data) includes all data points from a data set, without the use of any clustering. The second data set (wd_we) clusters data into two clusters, based on weekdays and weekend days. For the third and fourth data sets (hour and wd_we_hour), clusters were used that are based on the time of day and identified through the use of the K-means clustering algorithm (as shown in Section 4). Finally, two different data set sizes are explored. For individual data sets (Individual), one model is trained for each smart meter, where each Individual model has 8760 data points corresponding to the number of hours in 2017. For the group data set (Group), one model is trained using the entire data set of 85 smart meters, where the Group model has 744,600 data points, corresponding to the product of the number of hours in a year and the number of meters.

This study evaluates the effects of smart metered data as the main predictor to forecast water demand (DS predictor). The inclusion of additional sets of predictors (DS $\cup$ W, DS $\cup$ CH, and DS $\cup$ W $\cup$ CH), different clusters based on time of day and day of the week (all_data, wd_we, hour, wd_we_hour), and different size of the data sets (Individual, Group) are explored as they affect the capabilities of the forecasting model. Experiments for each combination of predictor set, cluster, and data set size are conducted (Table 3). For example, one model is developed to forecast water demand using only previously recorded consumption and seasonality data clustered by hour of the day during all days of the week for a single household. This combination corresponds to the settings DS, hour, Individual (Table 3). Based on the combination of predictors, clusters, and data set size, a total of 24 experiments are created, and for each experiment, the three ML methods are applied and evaluated. Similar to the work presented by Walker et al. (2015), to initialize each model, the first week of water consumption is stored to use as predictors (*pwsh* in Table 1) of the forecasting model.

Table 1
Predictor variable set, names, and definitions.

Predictor set	Predictor variable name	Predictor variable definition	Units
Demand and Seasonality Variables (DS)	$pwsh$	Water demand of the previous week same hour	gph
	$pdsh$	Water demand of the previous day at the same hour	gph
	$av24hr$	Average water demand of the last 24-p hours (p = 4)	gph
	$time$	Hour of the day (0 – 23)	NA
	day	Day of the week (0 – 6)	NA
Weather Variables (W)	T	Temperature	°F
	DP	Dew point	°F
	H	Humidity	%
	P	Hourly precipitation	in.
	$DailyMaxT$	Max temperature of the previous day	°F
	$DailyMaxH$	Max humidity of the previous day	°F
	$PrecipOccu$	Occurrence of precipitation in Previous 24-h period (0 = no, 1 = yes)	NA
Property Characteristic Variables (CH)	$LArea$	Property land area	acre
	$PropVal$	Property value	US $
	$BAge$	Building age	year
	$BArea$	Building area	ft^2

The Root Mean Squared Error (RMSE) is used as a metric of performance to evaluate the models. For Individual models, RMSE is calculated for each of the 85 meters and for each experiment as shown in Eq. (3):

$$RMSE_j = \sqrt{\frac{\sum_{i=1}^{N_j} (D_{pred_i} - D_{obs_i})^2}{N_j}} \tag{3}$$

where RMSE$_j$ is the RMSE for the jth meter for the training, validation, or test data set, and there are N_j data points associated with the jth meter. D_{pred_i} is the ith demand predicted using a forecasting model, and D_{obs_i} is the ith demand observed. Note that different data points (D_{pred_i}) may be predicted using different models at one meter if data were clustered. Each model was trained 30 times for one meter to account for stochasticity of the ML methods, and the average RMSE associated with each meter is reported as the average across the 30 trials in gph. To aggregate the RMSE value for comparison among the combination of model settings, we report the median across the meters of the average RMSE values.

For Group models, RMSE is evaluated across the entire data set as shown in Eq. (4):

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N} (D_{pred_i} - D_{obs_i})^2}{N}} \tag{4}$$

where RMSE is calculated for the training, validation, or test data set of N data points across multiple meters. Again, different data points may be predicted using different models if data were clustered. For Group models, RMSE is the average performance across 30 trials.

The Spearman's Rank-order Correlation (r_s) is used as a second metric to define the strength of a monotonic relationship between observed and predicted water demand. The Spearman's rank correlation is used instead of the Pearson's correlation (R), because a linear relationship is not apparent in this data, which is a zero-inflated date set (Myers and Sirois, 2004). The Spearman's rank correlation was not aggregated and is reported in the Results section for each meter.

For each of the 24 experiments (Table 3), data are divided into training, validation, and test sets for model developing. Training data are used for the model to learn from the data. The validation data set is used to identify the model parameters to best fit the modeled outputs with the observed data while reducing overfitting. Finally, the model is applied to a test data set to evaluate its performance (James et al., 2013). A random sampling without replacement algorithm was used to divide the data set, ensuring that no overlapping occurs. Based on

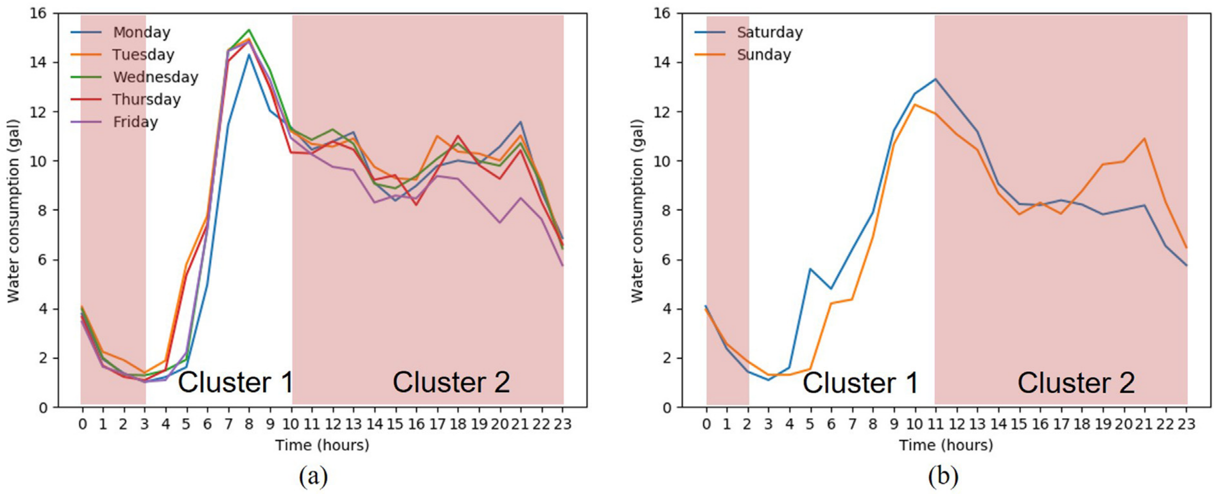

Fig. 2. Clusters identified for (a) weekdays and (b) weekends using the hourly average consumption of the set of meters.

Table 2
Clustered input.

Clusters	Description	Number of clusters
all_data	All data points are included in the model (no clustering)	1
wd_we	Data are clustered by days of the week	2
hour	All data points are clustered by time of day	2
wd_we_hour	Data are clustered by day of the week and time of day	4

Table 3
Combinations of settings for model input used for experiments.

Predictor variable set	Clusters	Size
DS	all_data wd_we hour wd_we_hour	Individual Group
DS ∪ W	all_data wd_we hour wd_we_hour	Individual Group
DS ∪ CH	all_data wd_we hour wd_we_hour	Group
DS ∪ W ∪ CH	all_data wd_we hour wd_we_hour	Group

previously conducted studies (Mouatadid and Adamowski, 2017; Guo et al., 2018), 80% of data were used for training, 10% for validation, and 10% for testing. Because of the noisiness in the water demand data, large training sets are needed to guide the learning procedure that is used to develop the model. Only a few settings are determined based on the performance for validation data, and a small data set is sufficient to make those selections. To evaluate the effects of random sampling, the data partition was re-initialized for each of the 30 trials. The approach leaves 10% of data out for each run, thus applying a hold-out cross validation technique to generalize the model.

5.2. Machine learning settings analysis

Preliminary analysis was explored to set algorithmic parameters for each ML approach. A set of 10 smart meters was randomly selected from the set of 85 meters. For this analysis, the predictor variable was set as DS; no clustering was applied (cluster: all_data); and Individual data sets were used to train one Individual model separately for each meter. A selection of settings were evaluated for each ML method based on previously conducted work (Herrera et al., 2010; Antunes et al., 2018). Each ML method was run 30 times, and the average and standard deviation of the RMSE values for each meter were calculated. The settings that produced the lowest average RMSE across the 30 trials were used for the remainder of the study described in this manuscript. Because the differences in the mean of RMSE values reported in Tables 4 and 5 are small for each combination of settings, the computational complexity (e.g., computational time) of executing each method was also considered when selecting settings.

5.2.1. Random forest settings

The RF model was developed using the bootstrap aggregation technique (Breiman, 1996), and algorithmic settings include the number of trees and the minimum number of observations per tree leaf. Three values were evaluated for both parameters: 50, 100, and 200. The range for evaluating the number of trees was based on research showing that 50 or fewer trees lead to accurate predictions (Antunes et al., 2018). Another study demonstrates that using more than 200 trees increases the execution time beyond practical limits (Villarin and Rodriguez-Galiano, 2019). Within this range, the best performance was found for 100 trees and a leaf size of 50 (Table 4).

5.2.2. Artificial neural network settings

An ANN was applied to explore settings for the number of hidden layers and the number of neurons per layer. Based on previous applications of ML methods for water demand (Herrera et al., 2010; Mouatadid and Adamowski, 2017; Antunes et al., 2018), potential settings for both parameters were identified as 1, 20, and 30, and a total of nine combinations were analyzed. The best settings are selected based on both model performance and computing time, as one hidden layer and 20 neurons (Table 4). Previous predictive models have also used relatively low numbers of neurons, as a characteristic of shallow neural networks (Cutore et al., 2008; Herrera et al., 2010; Antunes et al., 2018).

Table 4
Average RMSE ($\overline{RMSE}$) and standard deviation (σ) reported for each combination of model settings for RF and ANN.

RF settings									
Num of trees	50	50	50	**100**	100	100	200	200	200
Leaf size	50	100	200	**50**	100	200	50	100	200
$\overline{RMSE}$	13.06	11.97	13.86	**11.73**	13.5	11.93	13.08	13.53	13.75
σ	2.45	1.63	1.85	1.38	1.42	1.62	0.86	1.99	2.67
ANN settings									
Num of hidden layers	1	**1**	1	20	20	20	30	30	30
Num of neurons	1	**20**	30	1	20	30	1	20	30
$\overline{RMSE}$	13.01	**12.10**	12.77	13.35	12.49	12.32	12.13	13.32	13.11
σ	1.38	2.24	1.63	1.11	2.18	1.66	0.80	0.81	2.69

Table 5
Average RMSE ($\overline{RMSE}$) and standard deviation (σ) reported for each combination of model settings for SVR.

SVR settings									
Box constraint	**50**	50	50	50	50	50	50	50	50
Kernel scale	**10**	10	10	50	50	50	100	100	100
ϵ	**10**	50	100	10	50	100	10	50	100
$\overline{RMSE}$	**14.70**	26.65	23.46	16.28	27.25	23.25	16.75	27.68	23.51
σ	1.12	0.95	1.84	1.52	0.97	1.29	1.42	0.60	1.54
Box constraint	500	500	500	500	500	500	500	500	500
Kernel scale	10	10	10	50	50	50	100	100	100
ϵ	10	50	100	10	50	100	10	50	100
$\overline{RMSE}$	15.51	25.97	23.77	16.77	26.68	23.15	15.95	25.93	23.61
σ	2.19	1.21	1.18	1.98	0.80	0.74	1.89	3.49	0.79
Box constraint	1000	1000	1000	1000	1000	1000	1000	1000	1000
Kernel scale	10	10	10	50	50	50	100	100	100
ϵ	10	50	100	10	50	100	10	50	100
$\overline{RMSE}$	15.92	25.84	23.07	14.69	26.65	23.74	16.86	26.94	23.65
σ	1.34	0.79	0.93	1.05	1.18	1.86	1.50	0.70	1.29

5.2.3. Support vector regression settings

The box constraint (bc) and bandwidth (ϵ) were evaluated based on procedures shown by Mouatadid and Adamowski (2017) and Fan et al. (2005). The effects of the kernel scale (ks) were also evaluated, based on previous work (Williams et al., 2005). The setting for the box constraint is based on the value of the hourly water demand, which ranges from 0 to 500 gph. The box constraint was tested at settings of 50, 500, and 1000. Similarly, the allowable margin, ϵ, that defines the feature space was evaluated using values in the same order of magnitude as the reported hourly water demand, at 10, 50, and 100. The kernel scale was evaluated at values of 10, 50, and 100. This analysis includes a total of 27 combinations. The settings of $bc = 50$, $ks = 10$, and $\epsilon = 10$ generate the lowest RMSE value (Table 5). Many of the RMSE values reported by SVR were higher than those reported by RF and ANN. A Bayesian optimization algorithm (Gelbart et al., 2014) is available in the MATLAB toolbox and was applied to improve the performance of the SVR models. The Bayesian optimization subroutine is constrained by the size of the data set and could only be applied to train Individual models. The performance of the SVR that was found using the optimization procedure is approximately the same as the best values (shown in boldface) reported in Table 5.

5.3. Models for individual meters

The performance of the models used for the individual data sets is reported in this section. Each of the 85 smart meters is included in the analysis, and predictor variable sets DS and DS ∪ W and alternative settings for clustering are evaluated. The settings defined in Section 5.2 are applied to train the Individual models: ANN with one hidden layer and 20 neurons; RF with 100 trees and at least 50 observations per leaf; and SVR with box constraint equal to 50, kernel scale equal to 10, and ϵ equal to 10.

The median of the average RMSE for each meter is reported in Table 6. All results are reported for test data. For each of the ML methods, the median of the RMSE values decreases when data are clustered. The lowest error is reported as 9.5 gph by the RF model using the DS predictor set and clustering by days and hours (wd_we_hour). For these settings, the average RMSE reported for 85 meters ranges from 4.3 gph to 80.1 gph. Out of the 85 meters, 28 meters report an average RMSE value less than 10 gph, which is the resolution of the data. Five meters report an average RMSE greater that 40 gph. For these five meters, the demand pattern is erratic, with multiple changes between low (e.g., 10 gph) and high (e.g., 100–500 gph) demands during a 24-h period. Three of these five meters correspond to "other" meters that are not associated with single family residential accounts, which may explain a lack of pattern in the demand data. The Individual models, however, do not perform poorly overall in simulating demands at "other" meters based on RMSE values; the distributions of RMSE values for single-family residential meters and other meters are similar.

Table 6
Median of the averages of RMSE (gph) reported for Individual models developed using different ML methods, predictor sets, and clusters. Results are shown for test data.

Cluster	DS			DS ∪ W		
	RF	ANN	SVR	RF	ANN	SVR
all_data	10.9	10.8	11.5	10.4	10.4	11.0
wd_we	10.7	10.6	11.5	10.3	10.3	11.0
hour	9.8	9.6	10.5	9.7	10.0	10.5
wd_we_hour	**9.5**	**9.5**	10.3	9.7	9.9	10.4

The ANN models with the same settings (DS predictor set and wd_we_hour cluster) reported similar RMSE values, with a median of the averages of 9.5 gph. The highest error is reported by SVR. The lowest RMSE value reported by the SVR models is 8% higher than the RMSE found using the RF and ANN models and corresponds to the DS predictor set with wd_we_hour cluster.

The longest computational time for training a model for one individual meter was approximately 30 s using the all_data cluster and the DS ∪ W predictor set, and the time required was similar for each of the ML methods. A PC with an i7 processor and 16.0 GB of RAM was used for the experiments.

The Spearman's rank correlation (r_s) is also shown for Individual models using the DS predictor set alone (Fig. 3), because this set performed best based on RMSE values, as shown above (Table 6). When the entire data set is used without clustering (all_data), the interquartile range of r_s for RF spans from 0.36 to 0.44, and meters reporting an r_s greater than 0.58 are considered as outliers. Similar to the results reported for RMSE values, RF reports the highest median of r_s (0.42) for the DS predictor set and the all_data cluster. Based on r_s values, the Individual models perform similarly for single-family residential meters and other meters.

Models based on data that are clustered using the hours cluster shows less variability than other clusters (shown in the bottom row of Fig. 3). While clustering improved the RMSE value associated with models, clustering does not similarly increase the value of r_s. A similar set of subplots was generated for the DS ∪ W predictor variables and the results are similar to those reported in this section.

5.4. Models for group data set

The results of the ML models applied to the Group data set are presented in this section. These models use the entire data set of water demand reported by the 85 smart meters during 2017, and the same set of clusters were tested for developing models (Table 7).

The lowest error is reported by RF, corresponding to the predictor variable set DS ∪ W ∪ CH and the wd_we_hour cluster. Using the all_data cluster, the RF models outperform ANN and SVR models across the different set of predictors. ANN models perform similar to RF models for the Days of week cluster, reporting an RMSE of 17.8 gph with the predictors DS ∪ CH and DS ∪ W, respectively. Using the hour cluster, the lowest error is found using RF and the DS ∪ W ∪ CH predictor set, with an RMSE of 16.7 gph (4% above the lowest value). Unlike the models trained using the individual data set, the models trained using the group data set do not show performance that monotonically improves with clustering.

The longest computational time for training the Group model varied among the ML methods. ANN training took thirty minutes, RF required four hours, and SVR required six hours, using the all_data cluster and the DS ∪ W ∪ CH predictor set for each method. The running time did not vary with the type of predictor set but it did vary with the type of cluster, where clusters reporting large data sets (e.g., all_data) took longer than small data sets (e.g., wd_we_hour). Similar to the results reported for the Individual models, the results presented in Tables 7 and 8 correspond to the test data set.

For the Group models, the highest average value of the Spearman's rank correlation is found when using the hour cluster (Table 8). The comparison between ML methods shows that RF report the highest r_s values across the predictor sets and the clusters. The predictor set that produces the highest r_s corresponded to DS. Similar to the Individual models, RF and SVR report the highest and lowest values of r_s, respectively. When the entire data set is used (all_data), r_s ranges from 0.37 to 0.52. In this case, clustering decreases the value of the Spearman's rank correlation. When the size of the data set is smallest (wd_we_hour), r_s is reported at the lowest value (0.15). Including other variables in the predictor set does not improve r_s.

5.5. Feature importance analysis

Analysis of feature importance was conducted for the RF models. The predictor sets were DS ∪ W and DS ∪ W ∪ CH for the Individual and Group models, respectively, and the all_data cluster was used. The importance of each predictor was found using the tree-based iterative input selection algorithm (Galelli and Castelletti, 2013). Fig. 4a shows the most repeated results of feature importance among the 85 Individual models. Hour of the day (*time*), the average consumption of the previous 24 h (*av*24*hr*), and day of the week (*day*) are the three most important predictors. Fig. 4b shows the feature importance results of the Group model. The most important features in the Group model are the consumption of the last week at the same hour (*pwsh*), followed by the average consumption of the previous 24 h (*av*24*hr*), and the hour of the day (*hour*). New Individual and Group models were generated using only the three most important predictors for each model type, and the performance of the models in terms of RMSE and r_s showed negligible improvement.

5.6. Comparing the performance of individual and group models

The analysis conducted above explores the best settings to obtain Individual and Group models. In this section, the Individual and Group models are compared to provide recommendations about developing predictive models for application in the field. Individual models are trained to match data from a specific meter, and it is expected that an Individual model could precisely model the behavior at that meter. Group models, on the other hand, are developed using a larger data set, which may improve the performance over all meters. In addition, Group models should perform better for new meters that have not been used to train the model. To compare the two approaches, we select the best Individual and Group models, based on the experiments which reported the highest r_s values in Sections 5.3 and 5.4, respectively. The experiments reporting the best performance for the Individual and Group models correspond to RF, the DS predictor set, and the all_data cluster.

First, the performance of the best Individual and best Group model for one smart meter are shown as the observed and modeled water demand of hourly consumption (Fig. 5). To provide a more detailed visualization of the time series, Fig. 6 shows the same data during May 2017. The meter was selected randomly from the set of 85 meters. Both models capture the trends in water demands over the 8760 h, but the peak values are not accurately identified. The RMSE value reported for the Individual model calculated over the 8760 h is 30 gph, and the RMSE for the Group model is 38 gph. The Spearman's rank correlation (r_s) is 0.87 and 0.75 for the Individual and Group models, respectively.

The best Individual and best Group models were then compared based on their capability to predict test data for each smart meter. The test data set of each meter includes 876 values of water demand values (the test data set is defined as 10% of the data). The cumulative distribution plot (Fig. 7) shows that for the Individual and Group models, approximately 30% of the meters report RMSE values below 12 gph. At the upper limit of the distribution of errors, 10% of the meters report RMSE values above 25 gph. A two-sample Kolmogrov–Smirnov hypothesis test (Marsaglia et al., 2003) is applied to the RMSE of the models, and the test does not reject the null hypothesis that the results from the Individual and Group models are from populations with the same distribution at the 5% significance level, as shown in Fig. 7. Therefore, the results obtained from the best Individual and Group models in terms of RMSE are not significantly different. The r_s coefficient is also calculated to compare the values of the Individual and Group models for the test data. Fig. 8 shows that the Group model reports a stronger monotonic relationship between observed and modeled data than the Individual model. The two-sample Kolmogrov-Smirnov hypothesis test confirms that the Group model generates higher r_s values than the Individual model at the 5% significance level.

Five meters from the original set of 90 meters were used to test the ability of the best Individual and Group models to predict demands for new data sets. Similar to the 85 original meters, each of the five meters reports 8760 hourly water demand values during 2017. The RMSE value generated by both models ranges from 4.73 to 28 gph. The Group model shows slightly lower errors than the Individual model (Fig. 9). This is because the best Group model was trained using a larger data set than the best Individual model, and the predictive capability of the best Group model for new data sets is higher. The two-sample Kolmogrov-Smirnov was applied to test the null hypothesis that RMSE from the best Individual and Group models comes from population

Table 7
Average RMSE (gph) of Group models using different ML methods, predictor sets, and clusters. Results are shown for test data.

Cluster	DS			DS ∪ W			DS ∪ CH			DS ∪ W ∪ CH		
	RF	ANN	SVR	RF	ANN	SVR	RF	ANN	SVR	RF	ANN	SVR
all_data	18.9	19.5	22.3	18.6	19.2	24.8	18.1	19.1	22.8	17.8	19.1	25.2
wd_we	18.6	19.1	22.4	17.8	18.3	24.5	16.8	17.8	21.7	18.2	18.9	25.6
hour	18.0	18.3	20.9	18.2	18.5	23.0	17.0	17.4	20.8	16.7	17.0	21.8
wd_we_hour	17.1	17.5	22.5	17.2	17.8	24.8	16.9	17.7	23.3	**16.1**	17.1	24.7

Table 8
Average Spearman's rank correlation (r_s) values for Group models using different ML methods, predictor sets, and clusters. Results are shown for test data.

Cluster	DS			DS ∪ W			DS ∪ CH			DS ∪ W ∪ CH		
	RF	ANN	SVR	RF	ANN	SVR	RF	ANN	SVR	RF	ANN	SVR
all_data	0.52	0.47	0.37	0.48	0.45	0.42	0.50	0.45	0.43	0.49	0.45	0.39
wd_we	0.41	0.33	0.22	0.30	0.26	0.18	0.37	0.35	0.31	0.32	0.30	0.29
hour	**0.53**	0.49	0.26	0.48	0.44	0.41	0.51	0.48	0.45	0.51	0.48	0.46
wd_we_hour	0.32	0.30	0.15	0.32	0.30	0.16	0.29	0.25	0.24	0.29	0.27	0.25

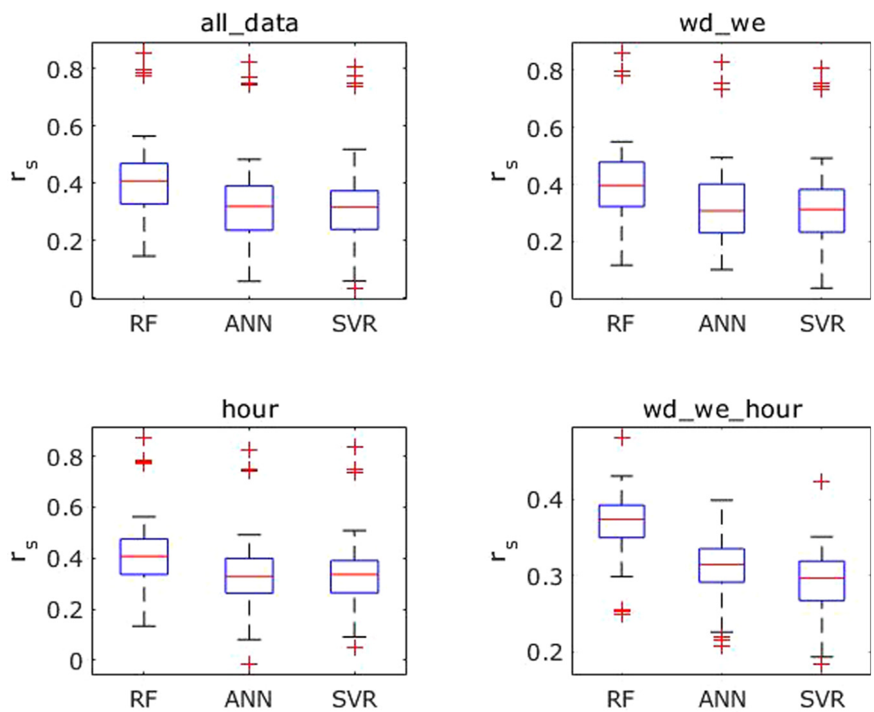

Fig. 3. Spearman's rank correlation (r_s) values of Individual models using DS as a predictor set. Results are shown for test data.

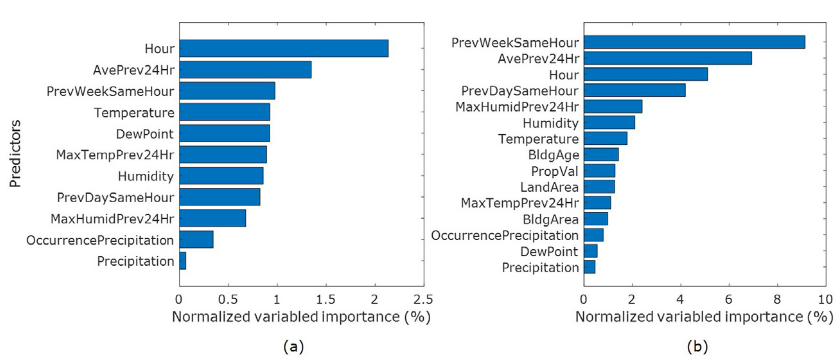

Fig. 4. Feature importance of the predictors: (a) Individual model, (b) Group model.

with the same distribution, and the result indicates that the test does not reject the null hypothesis at the 5% significance level. For the r_s coefficient, the median value reported by the Individual and Group models was 0.23 and 0.45, respectively. The two-sample test does not

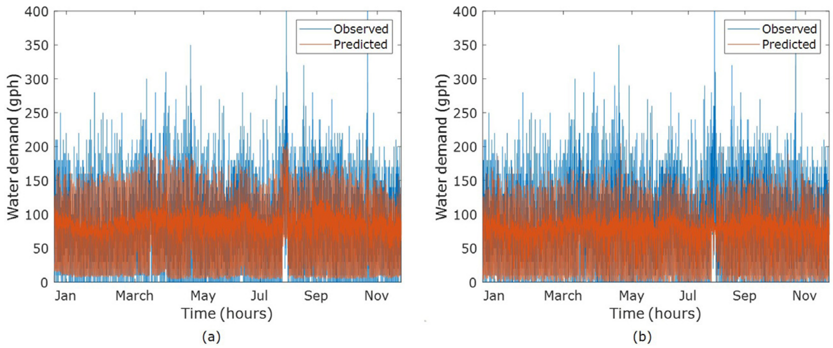

Fig. 5. Hourly demand reported at one meter using the (a) best Individual model and (b) the best Group model during 2017. Refer to the online version of the paper to view the figure in color.

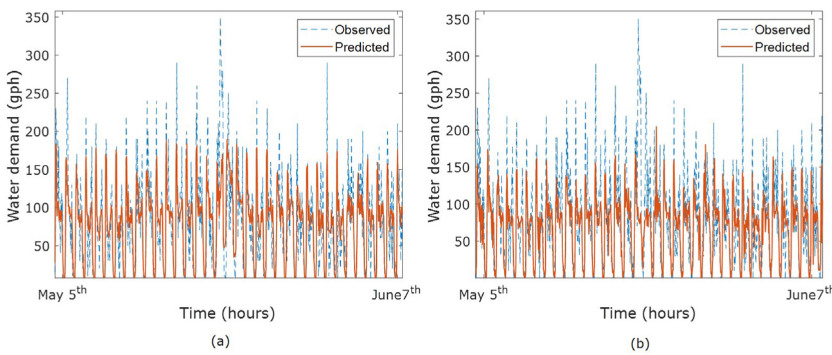

Fig. 6. Hourly demand reported at one meter using the (a) best Individual model and (b) the best Group model from May 5th to June 7th of 2017. Refer to the online version of the paper to view the figure in color.

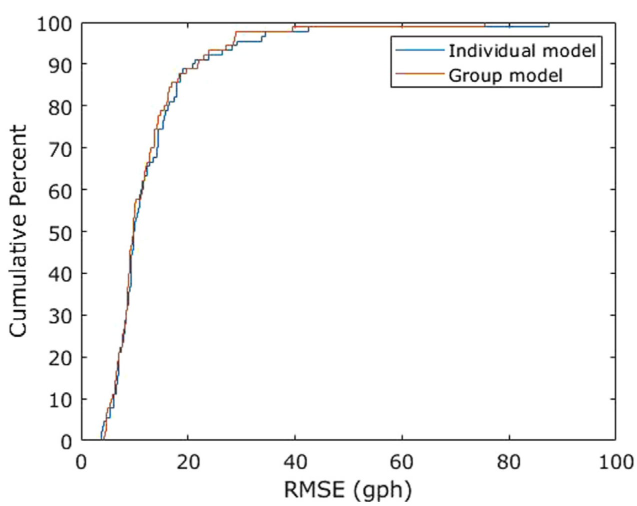

Fig. 7. Cumulative Distribution of RMSE values for the best Individual and Group models applied to each smart meter. Results are shown for test data.

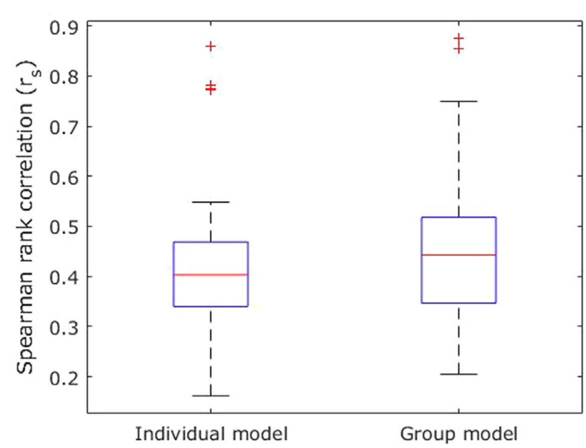

Fig. 8. Spearman's rank correlation for the best Individual and Group models applied to each smart meter. Results are shown for test data.

reject the null hypothesis that the r_s values come from populations with the same distributions. These results demonstrate that there is no statistical difference in performance of the best Individual and Group models based on RMSE or r_s at the 5% significance level for five new meters.

6. Discussion

This research tests the application of three ML methods for forecasting water demands on an hourly basis at individual accounts. Forecasting demands at individual accounts one hour ahead can enable a utility to identify abnormal consumption when comparing modeled and observed values with errors of ±10 gph. Forecasting water demand with one-hour resolution data is a challenging modeling task, and the precision of forecasting models may be improved through further research; however, this margin of error may be sufficient to identify a running toilet, for example, which can account for a loss of up to 140 gph (US EPA Water Sense).

To compare the performance of this work with previously conducted research reported by Walker et al. (2015), we calculate the Pearson's correlation coefficient (R) for each of the 85 meters using the best Group model. Our results show a range of R values from 0.25 to 0.80, which is a value higher than the range obtained by Walker et al. (2015). They report R values ranging from to 0.30 to 0.65 for

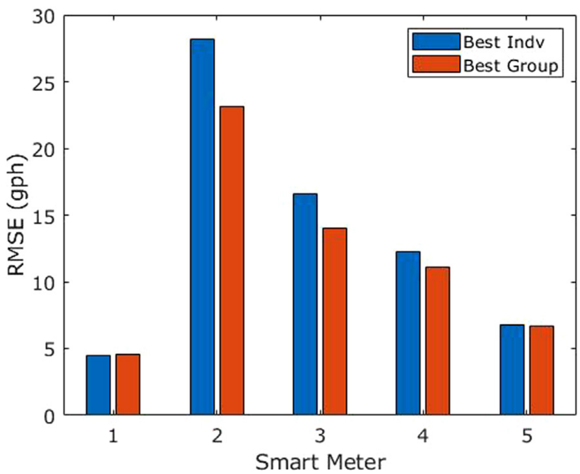

Fig. 9. Performance of the best Individual model and best Group model for 5 new meters.

predicting hourly water demands at nine account-level meters. We focus on the use of r_s in our analysis above to evaluate the monotonic relationship between observed and modeled data, instead of a linear relationship, which is represented by the calculation of R. Due to the volumetric resolution of the smart meters (10 gallons), the observed data set includes many zeros, and the MAPE cannot be calculated to compare these models with results reported by Candelieri (2017). Previous research that explored the use of SVR used only the time of day as a predictor (Candelieri et al., 2015), while this research explores alternative representations of previous water consumption, weather variables, and property characteristics as predictors. The methodology proposed here includes weekends in the analysis and explores the effect of demands that change during the weekend. Previous work in water demand using smart meters removed weekends from the data set to reduce noisiness in the data (McKenna et al., 2014). McKenna et al. (2014) do not report performance of their models for comparison purposes.

The approach presented here does not rely on information about the type of user account (i.e., residential vs. non-residential), which gives the process general applicability. The models perform similarly well for single-family residential meters and other meters. The data set is limited in representing types of other meters, and further research can explore the effects of user types on model performance through additional data. Clustering the data set to differentiate weekdays and weekends was, in general, effective in finding models with reduced errors. The performance of the best Individual model for weekdays and weekends differs by only 1.8%, demonstrating that the model can perform similarly well for both weekdays and weekends with slightly better outcomes for weekdays, in terms of RMSE. The models that were developed in this research had limited capabilities to accurately predict demand peaks (Fig. 5) because most peaks do not follow a periodic or predictable pattern. The hourly volumetric resolution of the meters (10 gallons) produces a zero-inflated time series, which creates difficulties in developing forecasting models. The number of meters allows us to test the capabilities of the models on new data sets using only five smart meters. Future work can use larger sets of smart meters to develop models. Training models for big data is computationally expensive, and optimization techniques were not applied when training the Group models due to the impracticability of the required computational time.

ANN is one of the most used techniques in the urban water demand field, as shown in previous works (Adamowski, 2008; Romano and Kapelan, 2014; Walker et al., 2015). This research found that RF performs similar to or better than ANN, as suggested by Herrera et al. (2010). This study did not explore the effects of different types of normalization to train a neural network and only applied the Min–Max scaling. Future work may explore the outcome of applying different normalization techniques.

The application of RF presents advantages for a better understanding of the model in terms of feature importance, as shown by Villarin and Rodriguez-Galiano (2019). Computing time is also an important criterion in the selection of models. RF, ANN, and SVR take around 30 s to train an Individual model using a PC with an i7 processor and 16.0 GB of RAM. When working with a Group model, ANN takes thirty minutes to train, whereas RF and SVR take four and six hours, respectively. While RF generates models with higher correlations, the time required for training models is much longer than the time required for ANNs.

Deployment strategies vary among utilities in the temporal resolution at which data are collected, which affects not only the cost of maintaining an AMI system, but also the type of analysis and modeling that can be conducted (McKenna et al., 2012). Hourly data, which may be considered as the upper limit of high temporal resolution, is too coarse to identify end-uses or support water quality modeling, but the data can be analyzed to detect anomalies in post meter water consumption (Britton et al., 2013) and to forecast water demands (Candelieri et al., 2015). The models developed here can be applied to forecast anomalies and send alerts to consumers. In the data set explored in this research, water consumption is reported in increments of 10 gph. This level of resolution can affect the performance of the ML methods explored here, which report continuous values as model output. Classifier approaches may provide a better performance for application to this data set. The precision that can be achieved using a classifier approach will affect the time required to train the classifier, and exploration of classifier methods was outside of the scope of the work presented here. The ML methods described here are broadly applicable across smart meter data sets of varying resolution and frequency. As technology for power or battery life and communication improves, data sets collected at smart meters may continue to increase in resolution and frequency, which may increase the variability of demands. Trade-offs between precision and predictive capabilities for increasingly high resolution data sets should be explored in future research.

7. Conclusions

This research develops a set of models to forecast water demand using data reported by smart meters installed at the user-level. The models are developed to forecast water demand at the subsequent time step. The input data used as predictors consist of lagged or previously observed water demand, weather variables, and characteristics of the households. Models were trained for individual meters and for the data set as a whole. Individual and Group models were compared using test data, which was randomly selected from the time series. Each Individual model was trained specifically for a meter and was able to continue to predict demands at that meter only marginally better than the Group model for all the predictor variables and clusters analyzed. When comparing the best Individual model with the best Group model using the test data set of each smart meter, the results in terms of RMSE are not different at the 5% significance level. However, for the Group model, the correlation between modeled and observed data, as measured by the Spearman's rank correlation (r_s), is higher than the r_s of the Individual model at the 5% significance level. A third comparison was performed to test the best Individual model and the best Group model for a new set of five meters. Based on that comparison, the Group model performed marginally better than the Individual model. The Group model used a much larger data set in training and was able to predict demands at new meters better than an Individual model, which was trained using a limited data set. However, when evaluating the results using a two-sample Kolmogrov Smirnov test, the RMSE values reported by Individual and Group models are not different at the 5% significance level.

Three ML methods were applied to forecast water demand based on regression: RFs, ANNs, and SVR. Despite fine-tuning the methods with similar combination of settings reported by previous works (Herrera et al., 2010; Mouatadid and Adamowski, 2017; Antunes et al., 2018), RF and ANN models outperformed SVR in all the applications. For the individual models, a Bayesian optimization of the hyperparameters was applied to SVR, and the RMSE values remained higher than those reported by RF and ANN models. This optimization technique was feasible only for the Individual models that have around 8700 data points. For the Group models (around 740,000 data points), the SVR settings were fixed, and the error followed the same pattern as the Individual models. This conclusion agrees with those reported by Brentan et al. (2018a), who applied SVR for water demand data.

The inclusion of exogenous variables, such as weather and property characteristics, only marginally improves the model performance. Effectively, most of the information from these ancillary variables is already captured in the consumption data. For Individual models, the performance of demand-driven models is not affected by the inclusion of weather variables. The results are in accordance with previous works that did not find a significant correlation between short-term (hourly) water demand and weather variables. Group models showed a slight improvement in performance due to the inclusion of weather variables and characteristics of the households (building area, lot area, building age, and property value) in the array of predictors. While an improvement was observed with the inclusion of these exogenous variables, obtaining this information may be impractical. Characteristics of the households are obtained from census data or surveys, and this data may not be public to protect the privacy of constituents.

Seasonality was analyzed by clustering data for weekdays and weekends and based on the time of the day. The performance of both Individual and Group models improved by clustering data, where no clustering resulted in the highest errors, and the highest level of clustering resulted in the lowest errors. Differentiating weekdays from weekends and clustering for the time of the day resulted in the lowest error for the Individual models, whereas only clustering for weekdays and weekends reported the lowest error for the Group model. In summary, clustering for seasonality improved Individual models more than Group models.

The median RMSE value reported by the models varies from 9.5 to 16 gph of water, giving some level of confidence for using models to alert customers of high water use anomalies that indicate potential post meter leaks when comparing actual and predicted consumption. The meter resolution fundamentally affects the performance of forecasting models. This variability is expected when working with noisy data, which is a characteristic of individual accounts at hourly temporal resolution. A higher resolution may lead to a better performance of forecasting models, and trade-offs may emerge in precision and variability in the data set, which is in accordance with recent works (Cominola et al., 2018a). Future work can explore the use of classifiers for forecasting data reported in discrete intervals. Real-time forecasting methods have a critical role in smart water management and provide a tool for identifying leaks, encouraging conservation, and shaving peak demands. We anticipate that further research will demonstrate the utility of these models in enhancing the performance of water distribution infrastructure.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was supported in part by the UNC Inter-institutional Planning Grant Program and the Ecuadorian government through the Secretaría Nacional de Educación Superior, Ciencia, Tecnología e Innovación (SENESCYT). The authors express their thanks to the Town of Cary for sharing data and collaborating in this research.

References

Adamowski, J.F., 2008. Peak daily water demand forecast modeling using artificial neural networks. J. Water Resour. Plann. Manag. 134 (2), 119–128. http://dx.doi.org/10.1061/(ASCE)0733-9496(2008)134:2(119), Retrieved from http://ascelibrary.org/doi/10.1061/(ASCE)0733-9496(2008)134:2(119).

Aitken, C., McMahon, T., Wearing, A., Finlayson, B., 1994. Residential water use: Predicting and reducing consumption. J. Appl. Soc. Psychol. 24 (2), 136–158.

Aksela, K., Aksela, M., 2011. Demand estimation with automated meter reading in a distribution network. J. Water Resour. Plann. Manag. 137 (5), 456–467. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000131, Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE29WR.1943-5452.0000131.

Alvisi, S., Ansaloni, N., Franchini, M., 2014. Generation of synthetic water demand time series at different temporal and spatial aggregation levels. Urban Water J. 11 (4), 297–310. http://dx.doi.org/10.1080/1573062X.2013.801499, Retrieved from http://www.tandfonline.com/doi/abs/10.1080/1573062X.2013.801499.

Antunes, A., Andrade-Campos, A., Sardinha-Lourenço, A., Oliveira, M., 2018. Short-term water demand forecasting using machine learning techniques. J. Hydroinform. (August), http://dx.doi.org/10.2166/hydro.2018.163.

Arandia-Perez, E., Uber, J.G., Boccelli, D.L., Janke, R., Hartman, D., Lee, Y., 2014. Modeling automatic meter reading water demands as nonhomogeneous point processes. J. Water Resour. Plann. Manag. 140 (1), 55–64. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000318, Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE29WR.1943-5452.0000318.

Arbelaitz, O., Gurrutxaga, I., Muguerza, J., Pérez, J.M., Perona, I., 2013. An extensive comparative study of cluster validity indices. Pattern Recognit. 46 (1), 243–256. http://dx.doi.org/10.1016/j.patcog.2012.07.021.

Arbues, F., Garcia-Valinas, M.A., Martinez-Espineira, R., 2003. Estimation of residential water demand: A state-of-the-art review. J. Socio-Econ. 32 (1), 81–102. http://dx.doi.org/10.1016/S1053-5357(03)00005-2.

Beal, C.D., Flynn, J., 2015. Toward the digital water age: Survey and case studies of Australian water utility smart-metering programs. Util. Policy 32, 29–37. http://dx.doi.org/10.1016/j.jup.2014.12.006.

Beal, C.D., Stewart, R.A., 2014. Identifying residential water end uses underpinning peak day and peak hour demand. J. Water Resour. Plann. Manag. 140 (7), 04014008. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000357, Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE29WR.1943-5452.0000357.

Berglund, E.Z., Monroe, J.G., Ahmed, I., Noghabaei, M., Do, J., Pesantez, J.E., …, Levis, J., 2020. Smart infrastructure: A vision for the role of the civil engineering profession in smart cities. J. Infrastruct. Syst. (in–press).

Blokker, E.J., Vreeburg, J.H., Buchberger, S.G., Van Dijk, J.C., 2008. Importance of demand modelling in network water quality models: A review. Drink. Water Eng. Sci. 1 (1), 27–38. http://dx.doi.org/10.5194/dwes-1-27-2008.

Blokker, E.J., Vreeburg, J.H., Van Dijk, J.C., 2009. Simulating residential water demand with a stochastic end-use model. J. Water Resour. Plann. Manag. 137 (February), 19–26. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000146, Retrieved from http://ascelibrary.org/doi/abs/10.1061/(ASCE)WR.1943-5452.0000002.

Bougadis, J., Adamowski, K., Diduch, R., 2005. Short-term municipal water demand forecasting. Hydrol. Process. 19 (1), 137–148. http://dx.doi.org/10.1002/hyp.5763.

Boyle, T., Giurco, D., Mukheibir, P., Liu, A., Moy, C., White, S., Stewart, R., 2013. Intelligent metering for urban water: A review. Water (Switzerland) 5 (3), 1052–1081. http://dx.doi.org/10.3390/w5031052.

Breiman, L., 1996. Bagging predictors - springer. Mach. Learn. 140, 123–140. http://dx.doi.org/10.1007/BF00058655, Retrieved from http://link.springer.com/article/10.1007/BF00058655%5Cnpapers2://publication/uuid/4331D2FC-33D0-45AB-80A0-B6FA074D5FE1.

Breiman, L., 2001. Random forests. Mach. Learn. 45 (1), 5–32. http://dx.doi.org/10.1023/A:1010933404324, Retrieved from http://link.springer.com/10.1023/A:1010933404324.

Breiman, L., Friedman, J., Olshen, R., Stone, C., 1984. Classification and regression trees, first ed. Taylor & Francis Group, Boca Raton, FL.

Brentan, B., Laurain, V., Aberkane, S., 2018a. How to infer prior knowledge in water distribution data-driven models?. In: WDSA - CCWI 2018 Joint Conference. OJS/PKP, Kingston, ON, Canada, pp. 91–99.

Brentan, B.M., Meirelles, G.L., Manzi, D., Luvizotto, E., 2018b. Water demand time series generation for distribution network modeling and water demand forecasting. Urban Water J. 15 (2), 150–158. http://dx.doi.org/10.1080/1573062X.2018.1424211, Retrieved from https://www.tandfonline.com/doi/full/10.1080/1573062X.2018.1424211.

Britton, T.C., Stewart, R.A., O'Halloran, K.R., 2013. Smart metering: enabler for rapid and effective post meter leakage identification and water loss management. J. Cleaner Prod. 54, 166–176. http://dx.doi.org/10.1016/j.jclepro.2013.05.018, Retrieved from http://www.sciencedirect.com/science/article/pii/S0959652613003272.

Buchberger, S., Omaghomi, T., Wolfe, T., Hewitt, J., Cole, D., 2017. Peak water demand study. probability estimates for efficient fixtures in single and multi-family residential buildings. In: Executive Summary. IAPMO, Online. https://www.iapmo.org/.

Buchberger, S.G., Wells, G.J., 1996. Intensity, duration, and frequency of residential water demands. J. Water Resour. Plann. Manag. 122 (1), 11–19. http://dx.doi.org/10.1061/(ASCE)0733-9496(1996)122:1(11), Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE290733-9496281996291223A1281129.

Candelieri, A., 2017. Clustering and support vector regression for water demand forecasting and anomaly detection. Water 9 (3), 224. http://dx.doi.org/10.3390/w9030224, Retrieved from http://www.mdpi.com/2073-4441/9/3/224.

Candelieri, A., Soldi, D., Archetti, F., 2015. Short-term forecasting of hourly water consumption by using automatic metering readers data. Procedia Eng. 119, 844–853.

Cardell-Oliver, R., 2013. Water use signature patterns for analyzing household consumption using medium resolution meter data. Water Resour. Res. 49 (12), 8589–8599. http://dx.doi.org/10.1002/2013WR014458.

Cardell-Oliver, R., Wang, J., Gigney, H., 2014. Smart meter analytics to pinpoint opportunities for reducing household water use. J. Water Resour. Plann. Manag. c (6), 1–9. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000634.

Chen, J., Boccelli, D.L., 2018. Forecasting hourly water demands with seasonal autoregressive models for real-time application. Water Resour. Res. http://dx.doi.org/10.1002/2017WR022007, Retrieved from http://doi.wiley.com/10.1002/2017WR022007.

Cominola, A., Giuliani, M., Castelletti, A., Rosenberg, D.E., Abdallah, A.M., 2018a. Implications of data sampling resolution on water use simulation, end-use disaggregation, and demand management. Environ. Model. Softw. 102, 199–212. http://dx.doi.org/10.1016/j.envsoft.2017.11.022.

Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A.E., 2015. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. Environ. Model. Softw. 72, 198–214. http://dx.doi.org/10.1016/j.envsoft.2015.07.012.

Cominola, A., Spang, E.S., Giuliani, M., Castelletti, A., Lund, J.R., Loge, F.J., 2018b. Segmentation analysis of residential water-electricity demand for customized demand-side management programs. J. Cleaner Prod. 172, 1607–1619. http://dx.doi.org/10.1016/j.jclepro.2017.10.203.

Creaco, E., Blokker, M., Buchberger, S., 2017a. Models for generating household water demand pulses: Literature review and comparison. J. Water Resour. Plann. Manag. 143 (6), 04017013. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000763, Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE29WR.1943-5452.0000763.

Creaco, E., Kossieris, P., Vamvakeridou-Lyroudia, L., Makropoulos, C., Kapelan, Z., Savic, D., 2016. Parameterizing residential water demand pulse models through smart meter readings. Environ. Model. Softw. 80, 33–40. http://dx.doi.org/10.1016/J.ENVSOFT.2016.02.019, Retrieved from https://www.sciencedirect.com/science/article/pii/S1364815216300391.

Creaco, E., Pezzinga, G., Savic, D., 2017b. On the choice of the demand and hydraulic modeling approach to WDN real-time simulation. Water Resour. Res. 53 (7), 6159–6177. http://dx.doi.org/10.1002/2016WR020104.

Cutore, P., Campisano, A., Kapelan, Z., Modica, C., Savic, D., 2008. Probabilistic prediction of urban water consumption using the scem-ua algorithm. Urban Water J. 5 (2), 125–132. http://dx.doi.org/10.1080/15730620701754434.

David, A., Vassilvitskii, S., 2007. K-means + +: The Advantages of Careful Seeding. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'07), pp. 1027–1035.

Donkor, E.A., Mazzuchi, T.A., Soyer, R., Alan Roberson, J., 2014. Urban water demand forecasting: Review of methods and models. J. Water Resour. Plann. Manag. 140 (2), 146–159. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000314, Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE29WR.1943-5452.0000314.

Duerr, I., Merrill, H.R., Wang, C., Bai, R., Boyer, M., Dukes, M.D., Bliznyuk, N., 2018. Forecasting urban household water demand with statistical and machine learning methods using large space-time data: A comparative study. Environ. Model. Softw. 102, 29–38. http://dx.doi.org/10.1016/J.ENVSOFT.2018.01.002, Retrieved from https://www.sciencedirect.com/science/article/pii/S1364815217305406.

Fan, R.-E., Chen, P.-H., Lin, C.-J., 2005. Working set selection using second order information for training support vector machines. J. Mach. Learn. Res. 6 (Dec), 1889–1918.

Gagliardi, F., Alvisi, S., Franchini, M., Guidorzi, M., 2017. A comparison between pattern based and neural network short-term water demand forecasting models. Water Sci. Technol.: Water Supply 2 (1), 1–10. http://dx.doi.org/10.2166/ws.2017.045.

Galelli, S., Castelletti, A., 2013. Tree-based iterative input variable selection for hydrological modeling. Water Resour. Res. 49 (7), 4295–4310. http://dx.doi.org/10.1002/wrcr.20339.

Gargano, R., Di Palma, F., de Marinis, G., Granata, F., Greco, R., 2016. A stochastic approach for the water demand of residential end users. Urban Water J. 13 (6), 569–582. http://dx.doi.org/10.1080/1573062X.2015.1011666, Retrieved from http://www.tandfonline.com/doi/full/10.1080/1573062X.2015.1011666.

Gelbart, M.A., Snoek, J., Adams, R.P., 2014. Bayesian optimization with unknown constraints. In: Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, pp. 250–259.

Giurco, D.P., White, S.B., Stewart, R.A., 2010. Smart metering and water end-use data: Conservation benefits and privacy risks. Water 2 (3), 461–467. http://dx.doi.org/10.3390/w2030461, Retrieved from http://www.mdpi.com/2073-4441/2/3/461/.

Guo, G., Liu, S., Wu, Y., Li, J., Zhou, R., Zhu, X., 2018. Short-term water demand forecast based on deep learning method. J. Water Resour. Plann. Manag. 144 (12), http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000992.

Gurung, T.R., Stewart, R.A., Beal, C.D., Sharma, A.K., 2015. Smart meter enabled water end-use demand data: Platform for the enhanced infrastructure planning of contemporary urban water supply networks. J. Cleaner Prod. 87 (1), 642–654. http://dx.doi.org/10.1016/j.jclepro.2014.09.054.

Gurung, T.R., Stewart, R.A., Beal, C.D., Sharma, A.K., 2017. Smart meter enabled informatics for economically efficient diversified water supply infrastructure planning. J. Cleaner Prod. 163, S138–S147. http://dx.doi.org/10.1016/j.jclepro.2017.05.140, Retrieved from http://dx.doi.org/10.1016/j.jclepro.2016.07.017.

Gurung, T.R., Stewart, R.A., Sharma, A.K., Beal, C.D., 2014. Smart meters for enhanced water supply network modelling and infrastructure planning. Resour. Conserv. Recycl. 90, 34–50. http://dx.doi.org/10.1016/j.resconrec.2014.06.005.

Haykin, S., 2008. Neural Networks and Learning Machines, third ed. PEARSON, ISBN: 978-0131471399.

Herrera, M., Torgo, L., Izquierdo, J., Perez-Garcia, R., 2010. Predictive models for forecasting hourly urban water demand. J. Hydrol. 387 (1–2), 141–150. http://dx.doi.org/10.1016/j.jhydrol.2010.04.005.

House-Peters, L., Pratt, B., Chang, H., 2010. Effects of urban spatial structure, sociodemographics, and climate on residential water consumption in hillsboro, oregon. J. Amer. Water Resour. Assoc. (JAWRA) 46 (3), 461–472, doi: 10.1111.

Hussien, W., Memon, F.A., Savic, D.A., 2016. Assessing and modelling the influence of household characteristics on per Capita water consumption. Water Resour. Manag. 1–25. http://dx.doi.org/10.1007/s11269-016-1314-x.

Jain, A., Ormsbee, L.E., 2002. Short-term water demand forecast modeling techniques - conventional methods versus AI. J. Amer. Water Works Assoc. 94 (7), 64–72. http://dx.doi.org/10.1002/j.1551-8833.2002.tb09507.x.

James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. An Introduction to Statistical Learning, Vol. 103, third ed. Springer US, New York, NY, USA, http://dx.doi.org/10.1007/978-1-4614-7138-7, Retrieved from http://link.springer.com/10.1007/978-1-4614-7138-7.

Kavousian, A., Rajagopal, R., Fischer, M., 2013. Determinants of residential electricity consumption: Using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants' behavior. Energy 55, 184–194. http://dx.doi.org/10.1016/j.energy.2013.03.086.

Kofinas, D.T., Spyropoulou, A., Laspidou, C.S., 2018. A methodology for synthetic household water consumption data generation. Environ. Model. Softw. 100 (2), 48–66. http://dx.doi.org/10.1016/J.ENVSOFT.2017.11.021, Retrieved from https://www.sciencedirect.com/science/article/pii/S1364815216310520?via%3Dihub.

Lloyd, S., 1982. Least squares quantization in PCM. IEEE Trans. Inf. Theory 28 (2), 129–137.

Luciani, C., Casellato, F., Alvisi, S., Franchini, M., 2019. Green smart technology for water (gst4water): Water loss identification at user level by using smart metering systems. Water (Switzerland) 11 (3), http://dx.doi.org/10.3390/w11030405.

Maas, A., Goemans, C., Manning, D., Kroll, S., Arabi, M., Rodriguez-Mcgoffin, M., 2017. Evaluating the effect of conservation motivations on residential water demand. J. Environ. Manag. 196, 394–401. http://dx.doi.org/10.1016/j.jenvman.2017.03.008.

March, H., Morote, A.F., Rico, A.M., Sauri, D., 2017. Household smart water metering in Spain: Insights from the experience of remote meter reading in alicante. Sustain. (Switzerland) 9 (4), http://dx.doi.org/10.3390/su9040582.

Marsaglia, G., Tsang, W.W., Wang, J., 2003. Evaluating Kolmogorov's distribution. J. Stat. Softw. 8 (18), 1–4.

McKenna, S., Fusco, F., Eck, B., 2014. Water demand pattern classification from smart meter data. Procedia Eng. 70, 1121–1130. http://dx.doi.org/10.1016/j.proeng.2014.02.124.

McKenna, E., Richardson, I., Thomson, M., 2012. Smart meter data: Balancing consumer privacy concerns with legitimate applications. Energy Policy 41, 807–814. http://dx.doi.org/10.1016/j.enpol.2011.11.049.

Monks, I., Stewart, R.A., Sahin, O., Keller, R., 2019. Revealing unreported benefits of digital water metering: Literature review and expert opinions. Water 11 (4), 838.

Mostafavi, N., Shojaei, H.R., Beheshtian, A., Hoque, S., 2018. Residential water consumption modeling in the integrated urban metabolism analysis tool (IUMAT). Resour. Conserv. Recycl. 131, 64–74. http://dx.doi.org/10.1016/J.RESCONREC.2017.12.019, Retrieved from https://www.sciencedirect.com/science/article/pii/S0921344917304500.

Mouatadid, S., Adamowski, J., 2017. Using extreme learning machines for short-term urban water demand forecasting. Urban Water J. 14 (6), 630–638. http://dx.doi.org/10.1080/1573062X.2016.1236133, Retrieved from https://www.tandfonline.com/doi/full/10.1080/1573062X.2016.1236133.

Myers, L., Sirois, M.J., 2004. Spearman correlation coefficients, differences between. Encyclopedia Stat. Sci. 12.

Nguyen, K.A., Stewart, R.A., Zhang, H., 2014. An autonomous and intelligent expert system for residential water end-use classification. Expert Syst. Appl. 41 (2), 342–356. http://dx.doi.org/10.1016/j.eswa.2013.07.049.

Odan, F.K., Fernanda, L., Reis, R., 2012. Hybrid water demand forecasting model associating artificial neural network with fourier series. J. Water Resour. Plann. Manag. 138 (JUNE), 245–256. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000177.

Praskievicz, S., Chang, H., 2009. Identifying the relationships between urban water consumption and weather variables in Seoul, Korea. Phys. Geogr. 30 (4), 324–337. http://dx.doi.org/10.2747/0272-3646.30.4.324, Retrieved from http://bellwether.metapress.com/openurl.asp?genre=article{&}id=doi:10.2747/0272-3646.30.4.324.

Price, J.I., Chermak, J.M., Felardo, J., 2014. Low-flow appliances and household water demand: An evaluation of demand-side management policy in albuquerque, new Mexico. J. Environ. Manag. 133, 37–44. http://dx.doi.org/10.1016/J.JENVMAN.2013.11.025, Retrieved from https://www.sciencedirect.com/science/article/pii/S0301479713007184.

Romano, M., Kapelan, Z., 2014. Adaptive water demand forecasting for near real-time management of smart water distribution systems. Environ. Model. Softw. 60, 265–276. http://dx.doi.org/10.1016/j.envsoft.2014.06.016.

Savic, D., Vamvakeridou-Lyroudia, L., Kapelan, Z., 2014. Smart meters, smart water, smart societies: The iwidget project. Procedia Eng. 89, 1105–1112. http://dx.doi.org/10.1016/j.proeng.2014.11.231.

Sebri, M., 2016. Forecasting urban water demand: A meta-regression analysis. J. Environ. Manag. 183, 777–785. http://dx.doi.org/10.1016/j.jenvman.2016.09.032.

Smola, A.J., Scholkopf, B., 2004. A tutorial on support vector regression. Stat. Comput. 14, 199–222. http://dx.doi.org/10.1023/B:STCO.0000035301.49549.88.

Sonderlund, A.L., Smith, J.R., Hutton, C.J., Kapelan, Z., Savic, D., 2016. Effectiveness of smart meter-based consumption feedback in curbing household water use: Knowns and unknowns. J. Water Resour. Plann. Manag. 142 (12), 04016060. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000703, Retrieved from http://ascelibrary.org/doi/10.1061/28ASCE29WR.1943-5452.0000703.

Stewart, R.A., Nguyen, K., Beal, C., Zhang, H., Sahin, O., Bertone, E., ., Kossieris, P., 2018. Integrated intelligent water-energy metering systems and informatics: visioning a digital multi-utility service provider. Environ. Model. Softw. 105, 94–117. http://dx.doi.org/10.1016/j.envsoft.2018.03.006.

Town of Cary, 2013. Town of Cary public information. https://pubfiles.townofcary.org/GIS/ [Online; accessed 8-October-2017].

Villarin, M.C., Rodriguez-Galiano, V.F., 2019. Machine learning for modeling water demand. J. Water Resour. Plann. Manag. 145 (5), 04019017. http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0001067, Retrieved from http://ascelibrary.org/doi/10.1061/%28ASCE%29WR.1943-5452.0001067.

Walker, D., Creaco, E., Vamvakeridou-Lyroudia, L., Farmani, R., Kapelan, Z., Savić, D., 2015. Forecasting domestic water consumption from smart meter readings using statistical methods and artificial neural networks. Procedia Eng. 119 (1), 1419–1428. http://dx.doi.org/10.1016/j.proeng.2015.08.1002.

Williams, P., Li, S., Feng, J., Wu, S., 2005. Scaling the kernel function to improve performance of the support vector machine. pp. 831–836, Retrieved from http://link.springer.com/10.1007/11427391{_}133 http://dx.doi.org/10.1007/11427391_133.

Xenochristou, M., Blokker, M., Vertommen, I., Urbanus, J., Kapelan, Z., 2018. Investigating the influence of weather on water consumption: A dutch Case study. In: WDSA - CCWI 2018 Joint Conference. OJS/PKP, Kingston, ON, Canada, pp. 26–35.

Xenochristou, M., Kapelan, Z., Hutton, C., 2019. Using smart demand-metering data and customer characteristics to investigate influence of weather on water consumption in the UK. J. Water Resour. Plann. Manag. 146 (2), 04019073.