Emerging investigator series: Disaggregating Residential Sector High-Resolution Smart Water Meter Data into Appliance End-Uses with Unsupervised Machine Learning

Journal:	*Environmental Science: Water Research & Technology*
Manuscript ID	EW-ART-08-2020-000724.R1
Article Type:	Paper

Environmental Science: Water Research and Technology

Water Impact Statement for:

Disaggregating Residential Sector High-Resolution Smart Water Meter Data into Appliance End-Uses with Unsupervised Machine Learning

Gabrielle M. Bethke[1], Abigail R. Cohen[1], and Ashlynn S. Stillwell[1]

[1]Civil and Environmental Engineering, University of Illinois at Urbana-Champaign

Understanding how water is used within the home can improve conservation and efficiency. Smart water meters measure residential water consumption, but current water event classification techniques require extensive training data—a significant barrier to larger scale implementation. We lower this barrier using unsupervised machine learning techniques that reduce training data requirements, and create customized, data-informed residential water sustainability recommendations.

Emerging investigator series: Disaggregating Residential Sector High-Resolution Smart Water Meter Data into Appliance End-Uses with Unsupervised Machine Learning

Gabrielle M. Bethke[1], Abigail R. Cohen[1], and Ashlynn S. Stillwell[1]

[1]Civil and Environmental Engineering, University of Illinois at Urbana-Champaign

Abstract

The residential sector accounts for a significant amount of water consumption in the United States. Understanding this water consumption behavior provides opportunity for water savings, which are important for sustaining freshwater resources. In this study, we analyzed 1-second resolution smart water meter data from a 4-person household over one year as demonstration. We disaggregated the data using derivative signals of the influent water flow rate at the water supply point of the home to identify start and end times of water events. *k*-means clustering, an unsupervised machine learning method, then categorized these water events based on information collected from the appliance/fixture end uses. The use of unsupervised learning reduces the training data requirements and lowers the barrier of implementation for the model. Using the water use profiles, we determined peak demand times and identified seasonal, weekly, and daily trends. These results provide insight into specific water conservation and efficiency opportunities within the household (e.g., reduced shower durations), including the reduction of water consumption during peak demand hours. The widespread implementation of this type of smart water metering and disaggregation system could improve water conservation and efficiency on a larger scale and reduce stress on local infrastructure systems and water resources.

Introduction

Freshwater is a necessary resource for human life. Unsustainable depletion of fresh surface water and groundwater resources, along with the depletion of water quality through contamination, are growing issues that threaten widespread water security (1). While the management of water resources is imperative for ensuring water security in the future, there is a general lack of understanding of residential water consumption on a day-to-day basis. While monthly total volumetric utility data are available, finer resolution data reveal possibilities for identifying specific, customized opportunities for residential water conservation and efficiency.

Implementing smart water meters in the residential sector is a way to improve the temporal resolution of residential water consumption data, and further analysis of the data from these meters provides useful insight into residential water end uses. In this study, water main flow rate data were collected on 1-second resolution from a smart water meter system at a 4-person residence in central Illinois from February 2018 through January 2019, as a demonstration of a residential smart water metering system in a temperate-humid climate. This work explored two research questions: 1) What water consumption trends exist at the whole-home and appliance levels?, 2) Can 1-second resolution whole-home water meter data be used to identify appliance end uses with limited training data?

We analyzed water consumption on daily and hourly timescales to identify household-level trends. By quantifying average daily water consumption for each month, we revealed seasonal variation in consumption behavior. On the weekly scale, we compared average water consumption values for each day of the week to identify differences in weekday and weekend consumption. We implemented additional time-of-day analysis to find peak demand times in the morning and evening. Timing of peak demand can be useful for residential demand prediction and plumbing design (2), and widespread implementation of smart water meters can inform conservation and efficiency recommendations.

In addition to analyzing the water consumption trends of the overall household, we aimed to estimate water end uses within the home as a demonstration for further implementation of smart water metering systems. We created a model to disaggregate and categorize water events into appliance/fixture end uses. With access to 1-second resolution data, we used this approach to pinpoint the instances in which water valves opened and closed by analyzing the derivative signals from the flow rate, providing water event information of average flow rate and duration. We classified these water end-use events using *k*-means clustering, an unsupervised machine learning method. This unsupervised machine learning approach greatly reduces the amount of input data necessary, lowering barriers in the widespread implementation of residential water use studies.

These results can be used to identify customized conservation and efficiency recommendations for household water consumption, which could promote positive water consumption behavior change through feedback mechanisms. Overall, this study serves as a basis for extracting water demand information from smart meter data on multiple timescales, reduces the barrier of implementation of smart water meter disaggregation and appliance/fixture identification models, and explores the implications of how this information could be used for improved conservation and efficiency.

Background

Quantitative residential water consumption data have traditionally only been available through utilities at the monthly level. While several studies have utilized surveys to gain further insight into residential water use behavior (3-5), these surveys are associated with a high level of uncertainty. Higher resolution water use data are needed to more accurately analyze how water is used within the residential sector. Smart water metering is an emerging technology that can be used at the household level to aid in water conservation and efficiency. These meters have advanced to track flow rates on the order of seconds (6). This type of information is a vast improvement from the traditional data collected by utilities, and these data can provide a more comprehensive quantification of water consumption behavior than previously possible, reduce the uncertainty of how water is consumed, and provide more detailed insight into household water consumption.

There are several different technologies available that can be used to measure the flow of water through a pipe. Consequently, there are a variety of smart water meters on the market, including pressure sensors, accelerometers, mechanical meters, and magnetic meters. Pressure sensors use Poisuelle’s law to estimate flow rates based on pressure changes within the pipe as the valve is open or closed (7). Accelerometers track vibrations in the pipe from turbulence of moving water (8), while ultrasonic sensors transmit ultrasonic beams and measure the difference in time between these beams in flowing water (9). Finally, mechanical meters calculate the flow rate based on the movement of a disk within the meter (6), and magnetic meters measure the flow via the voltage induced across the fluid within a magnetic field (10). The development of these technologies not only improves the temporal resolution of available flow rate data, but allows for accurate water flow rate measurements as low as 0.02 L/min (11). Pressure sensors and mechanical or magnetic flow meters have been used as residential smart water meters due to the cost of ultrasonic sensors and the intensive calibration required for accelerometers.

Smart water meter information can be beneficial on a city-wide scale if data are captured remotely, and can be used for urban water planning (12) and demand forecasting (13). Information from smart water meters can be used by both residents and water utilities in widespread water supply planning, and water conservation and efficiency can benefit an area's entire water infrastructure system (14). Monthly utility water consumption data are often used to predict water demand at the district level (15), but access to daily or hourly data opens new potential for these water demand predictions. Insight into how water is used during the day, including overall demand and peak demand times (2), is important for water utility operation and for the construction and maintenance of infrastructure. Water infrastructure is typically designed based on average and peak daily and hourly demands, and these demands have a direct impact on infrastructure cost (16). This infrastructure experiences the most stress during peak demand times, so reducing peak demand can reduce the need for expensive water distribution network augmentations (17). The analysis of smart water meter data can assist in identifying opportunities for water conservation, and ultimately assist in peak demand reduction (18-20).

Improved availability of smart meters can also reduce city-wide non-revenue water, through physical losses and metering inaccuracies. Mukheibir et al. (21) found that non-registration and under-registration of flow rates in residential water meters contributed to non-revenue water and errors in the city-scale water balance. Meters can become less reliable with age and usage, but are not always replaced, which can lead to large volumes of unaccounted-for-water (22-23). Implementation of new meters can mitigate these issues, and an improvement in temporal resolution of the data could lead to fewer water losses, as quick feedback from smart meter data can also be used to identify and repair leaking pipes (24), saving water, energy, and money (25-27).

Water demand also varies seasonally, and data from residential smart water meters have been used to understand these trends in water consumption, which are important for regional water supply planning. A study of Albuquerque, New Mexico, showed that overall demand increased in the summer due to additional outdoor use (28). Other studies have shown this type of outdoor water consumption variability (29-30), and there is evidence that indoor residential water use may also vary seasonally. Rathnayaka et al. (31) showed that shower duration is negatively correlated with outdoor temperature, while other water-consuming appliance use behavior remains steady throughout the year. This type of seasonal demand information can be useful for conservation and efficiency efforts on behalf of the consumer, as well as for widespread water supply planning and management. To analyze appliance-level water consumption, however, coarse temporal resolution data are no longer sufficient, as most appliances operate on a sub-minute timescale. Therefore, sub-minute data are necessary for understanding residential water end uses.

Non-intrusive load monitoring (NILM) is a concept originally developed for smart electricity meters to determine the energy consumption of each appliance based on the household current and voltage load. The first NILM system was developed by Hart (32) and several studies have performed this type analysis on buildings using smart electricity meters (33-35). However, these methods cannot be exactly replicated and applied to smart water meter data, which consists primarily of fixtures that may not be used at the maximum flow rate (e.g., faucets) and are more subject to human control of the signal during consumption. The concept of disaggregating overall household data into appliance end-uses, however, has been applied to smart water meter data using a variety of techniques.

There are several software packages available to disaggregate smart water meter data, but the accuracy and ease of implementation of these packages is highly variable based on the resolution of data, specific appliance characteristics, model calibration requirements, and dependence on human understanding of the appliance flow signatures (11). Most existing appliance identification algorithms utilize supervised learning, a type of machine learning approach that requires a ground truth (36). In the case of residential appliance/fixture identification from smart water meter data, the "ground truth" would be prior knowledge of what the output appliances/fixtures should be, and the model would need to be calibrated using this knowledge. Consequently, extensive training data or sub-metered measurements specific to the household are required to calibrate the model using supervised machine learning.

Many studies have applied supervised machine learning techniques for appliance identification. Trace Wizard (37) and Identiflow (10) utilize a decision tree algorithm based on boundary conditions from the physical water consumption characteristics of the appliances (e.g., flow rate, volume, and duration). HydroSense, developed by Froelich et al. (7), utilizes probabilistic-basic classification that uses pressure changes from valves opening and closing. This method requires the installation of many additional pressure sensors throughout the household, presenting challenges to potential implementation on a larger scale. Beal et al. (38) developed the South East Queensland Residential End Use Study (SEQREUS) approach in a widespread study in Australia that combines Hidden Markov Models (HMMs) and Dynamic Time Warping techniques, a technique that shows over 80 percent accuracy but requires manual classification of inconclusive and combined events, in addition to extensive training data (39).

While supervised learning is a technique that has shown success in residential water use studies, the end-use monitoring systems (40) and training data required to calibrate the models can be tedious to collect, and the calibration can be labor and computationally intensive (41). Applying unsupervised machine learning techniques to residential water meter disaggregation studies reduces this barrier of implementation (42). Unsupervised learning is an approach in which the goal is to infer results from the output data; that is, the model does not rely on labeled data like models that apply supervised learning. Recent studies have applied unsupervised learning techniques such as *k*-means clustering (43-45), which classifies data into a set number of clusters. When applied to residential water use, automated appliances operate similarly each time and should appear within the same cluster, and these water events can be classified by inferring the clusters based on the physical characteristics of the appliances as found in the water audit.

Separating and identifying overlapping, or concurrent, events is a significant challenge in residential water use studies, and the accuracy of existing smart meter disaggregation models decreases significantly when encountering concurrent events. Nguyen et al. (39) developed a method of separating concurrent events by calculating the vector gradients of the flow rate data to identify start and end times of events that overlap. Once these events are separated, they can be treated as single appliance/fixture water events in the event classification process. In this study, we utilized a vector gradient method to disaggregate concurrent water events from high-resolution household meter data. By applying *k*-means clustering, this study explored classification of water events using limited training data.

Methods

Setup and Data Collection

To advance understanding of residential water end uses, we installed a smart water metering system at a 4-person residence in central Illinois, collecting 1-second resolution flow rate data from February 2018 through January 2019. During installation of the water meter, we completed a water audit (shown in Figure SI1) to document the model and brand of water appliances and fixtures throughout the home, along with characteristics about the dwelling and residents (e.g., home size and age, number and age of occupants, etc.). This data collection included only factual data such that this work was determined not to meet the definition of human subjects research and, therefore, did not require Institutional Review Board (IRB) approval. Documentation of this IRB decision is available upon request.

Our methodology consisted of three general steps: 1) smart water meter flow rate monitoring, 2) data logging and formatting, and 3) data analysis. We collected the data from a smart water meter installed on the main water supply pipe into the residents' home. This meter is a custom ally® water meter provided by Sensus and monitors flow rate data at a 1-second resolution in units of gallons per minute (gpm); the analysis results are reported in both L/min and gpm. The meter can read a positive flow rate above 0.11 L/min (0.03 gpm), and the data are reported to 0.04 L/min (0.01 gpm) resolution. The water meter wrote data to a computer running a script to parse the raw data into a suitable format for further analysis. This data acquisition computer logged the data as a comma separated values file for each day. These files contained the water flow rate [gpm], temperature [K], pressure [psi], and volume [gal] timestamped at each second. With a focus on water consumption trends in this analysis, we only considered the flow rate and timestamp; additional analyses using the pressure and/or temperature data are reserved as future work. The data were further cleaned by adjusting the timestamp to Central time, converting flow to metric units, and removing any duplicate recordings or blank entries. Additionally, the meter often read values slightly above or below zero when water was not flowing. Recorded values between -0.11 L/min (-0.03 gpm) and +0.11 L/min (+0.03 gpm) were adjusted to 0.00 L/min (0.00 gpm), based on the reported flow rate resolution of the meter.

We analyzed the household's water use at both the whole-household and sub-household, or appliance/fixture, levels. The first portion of this analysis focuses on household level water use to determine temporal patterns on weekly and monthly bases, as well as a half-hourly basis to determine the peak times of use throughout the day. Through this approach, we estimated how residential water consumption varies seasonally and throughout the week (workdays and weekends) in terms of volume of water used per day and peak times of use.

Household Level Water Use

The first step in the household-level analysis was to identify total daily water consumption trends on the weekly and monthly levels. Using the trapezoidal rule, as shown in Equation 1, we estimated the overall volume of water consumed at the residence on each day and averaged these values over each month and day of the week. Days with water consumption less than 38 L (10 gal), suggesting no occupancy, were excluded from the larger-scale trend analysis.

$$Volume = \frac{1}{60}\int_a^b f(t)dt \approx \frac{1}{60}\left[\frac{f(a)+f(b)}{2} + \sum_{k=1}^{n-1}\left(a + k\frac{b-a}{n}\right)\right] \qquad \text{Equation 1}$$

Where a is the interval start time [s], b is the interval end time [s], f(t) [gpm or L/min] is the flow rate at time t [s], and n is the number of seconds in the interval

For longer time scales, we conducted a *t*-test (Equation 2) to determine whether the monthly and weekly mean consumption trends were statistically different ($p < 0.05$).

$$t = \frac{(\bar{x}_1 - \bar{x}_2)}{\sqrt{\frac{s_1^2}{n_1} - \frac{s_2^2}{n_2}}} \qquad \text{Equation 2}$$

where t is the t statistic, $\bar{x}$ is the mean of the sample, and s is the variance of the sample, and n is the number of subjects in the sample, conducting the following hypothesis test:

$$H_0: \mu_1 - \mu_2 = 0$$

$$H_A: \mu_1 - \mu_2 \neq 0$$

Appliance Level Water Use

The next portion of the study focused on the end use of water that enters the home through the water main. The goal was to provide a more quantitative method of creating a water-use profile for the specific test house. We first collected information from each of the water-using appliances and fixtures within the household through a water audit. We created a limited training dataset by running each water-using appliance and fixture in isolation to learn more about the characteristics of each end use and to compare against manufacturers' ratings. These appliances and fixtures were categorized into two groups: automatic end uses and human-controlled end uses.

We defined automatic end uses as those that function approximately the same during each water use event. The toilets and dishwasher are examples of automatic end uses that use water in approximately the same manner with each event (i.e., flushing, dishwashing cycle). While the clothes washing machine is also an automatic end use, it is programmed to fill with water based on sensing the amount of laundry in the washing machine.

Human-controlled end uses vary greatly with each use. Examples of these end uses are the sink, bathtub, and shower faucets. The user of these fixtures controls when the event starts and ends. Additionally, the flow rate of these events often varies because the user is in control of the amount that the valve opens during the event. These events are more difficult to identify because they do not have a single event "signature" like the automatic end uses.

We next analyzed how water was used within the household throughout the day. We separated the measured flow rate data at the water main into water events, defined as instances in which the flow rate was positive (i.e., greater than 0.11 L/min (0.03 gpm)) for at least 3 consecutive seconds. In this study, we isolated these water events from the overall dataset to support data disaggregation and ultimately identify the appliances and fixtures associated with these water events. A depiction of such water use events is shown in Figure 1

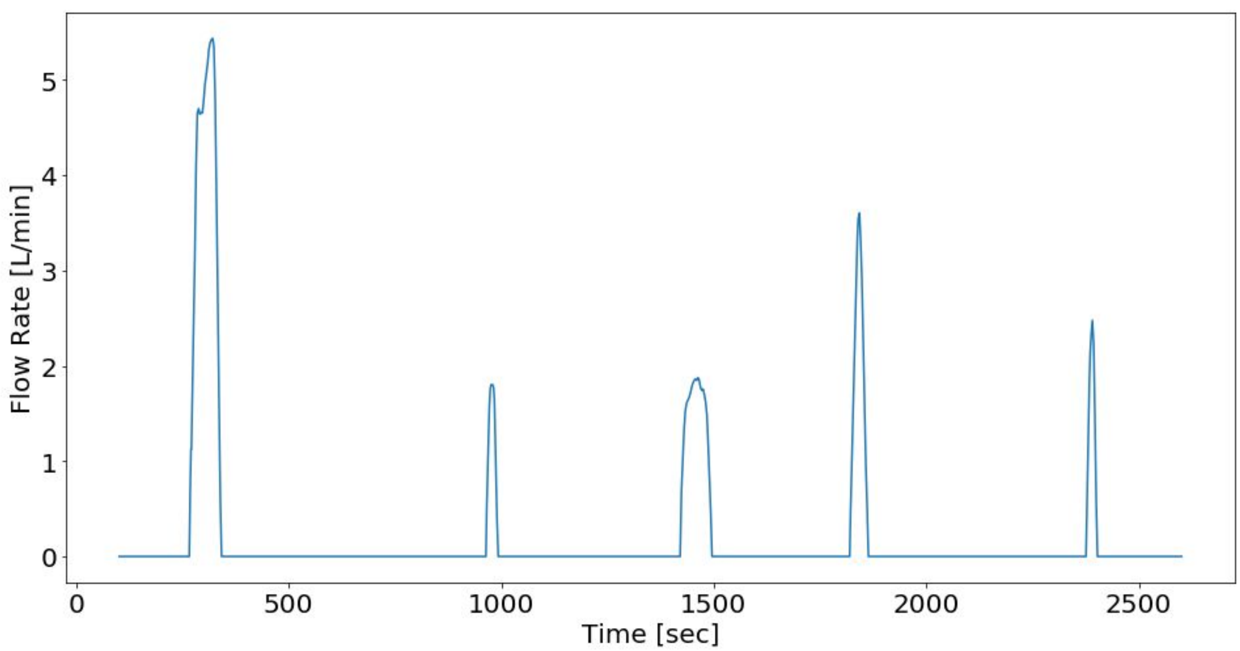

Figure 1: Water use events, or instances in which the flow rate is consecutively non-zero, were collected for further analysis.

Our disaggregation algorithm initially separates the data into instances of consecutive seconds with the flow rate is greater than zero (i.e., greater than 0.11 L/min (0.03 gpm) using the filtered data). These water events, along with time and flow rate information, were stored separately from the remainder "non-use" data, greatly decreasing the amount of data under consideration for disaggregation.

Separating Concurrent Events

Separating and identifying overlapping, or concurrent, water use events is a significant challenge in residential water studies reported in the literature, and the accuracy of existing smart meter disaggregation models decreases significantly when these types of events are encountered (11). Concurrent events occur fairly often, especially during longer duration events such as showers, and disaggregating concurrent events from one another is important for the purpose of creating a comprehensive water profile for the household. Nguyen et al. (39) developed a method of separating concurrent events by calculating the vector gradients of the flow rate data to identify start and end times of events that overlap. Once these events are separated, they can be treated as single-appliance water events in the event classification process. These vector gradients, g_i, are calculated using Equation 3.

$$g_i = \frac{a_{i+1} - a_i}{dt} \qquad \text{Equation 3}$$

where a_i is the water flow rate at time t_j, and dt is equal to the time step.

We applied Equation 3 to each collected water event, providing a vector gradient for each second as the derivative values associated with the flow rate. Because the water meter was capable of reading data at a resolution of 0.11 L/min (0.03 gpm), vector gradients with a value between -0.66 and +0.66 L/min^2, or a change of 0.11 L/min with a time differential of one second, were considered as zero so that only significant changes in flow rate were identified. These significant increases and decreases in flow rate within the event signified when water valves were opened or closed. For a single event, the major increase in flow signified the start of the water event when the valve was opened, and the water begins flowing. A significant decrease in the flow rate associated with this event signified the ending of the event when the valve is closed. The disaggregation process for concurrent events was based on the assumption that the absolute value of the major increase and decrease of a single event should be similar, a suitable assumption for an incompressible fluid.

To collect these positive and negative derivative signals, we created and applied a similar algorithm to the water event collection process to isolate the major increases and decreases from the vector gradients (derivatives) within each water use event. For each observed and isolated event, this disaggregation algorithm iterated through the vector gradients, creating lists of consecutive non-zero gradients. These lists of consecutive non-zero gradients were then summed and multiplied by the length of the list in seconds to obtain the overall increase or decrease in flow rate in L/min. The positive values were added to a list of positive signals, while the negative values were added to a separate list for later comparison.

Single events typically contained only one major increase and one major decrease, as shown in Figure 2. There were two exceptions to this assumption, which were both accounted for within the algorithm. First, human-controlled water events can have multiple increases or decreases if the valve is not opened or closed in a continuous manner. Secondly, shower events were treated slightly differently in the disaggregation process because both the bathtub and showerhead faucets were used during this single event. The bathtub was typically turned on first before switching to the showerhead, which has a lower flow rate. As displayed in Figure 2, shower events had one major increase and two major decreases that summed to a similar magnitude of the flow rate of the bathtub.

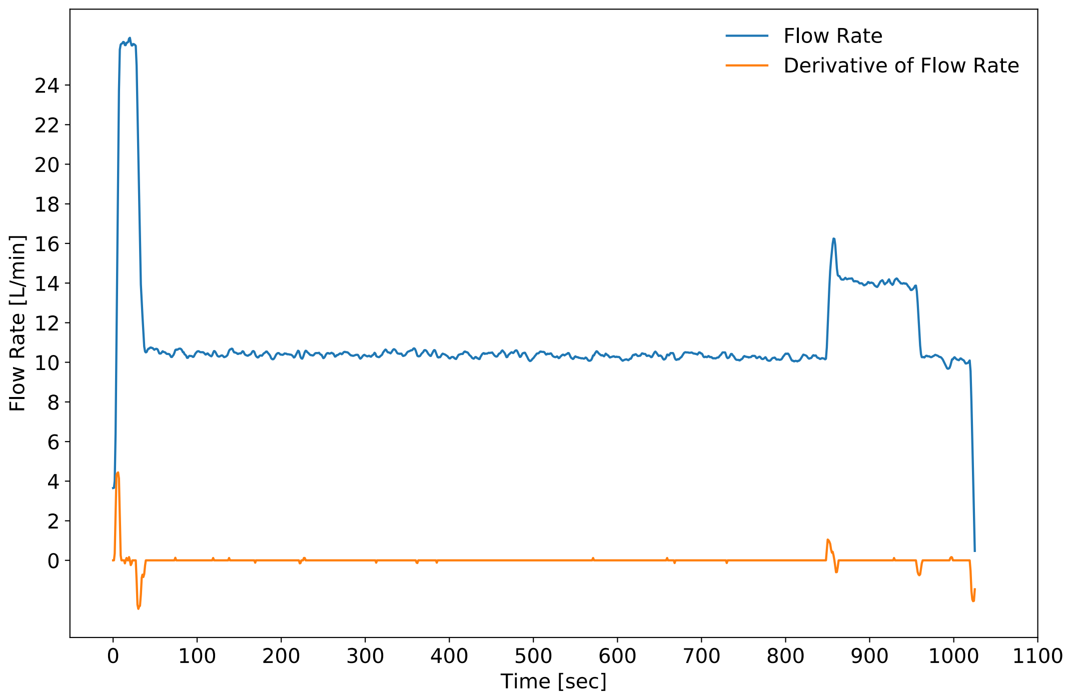

Figure 2: Shower events typically began with the bathtub valve opening before switching to the shower fixture, providing one major increase and two major decreases in the flow rate (blue). The vector gradient (derivative), shown in orange, was used to identify single and concurrent water use events.

We used the number of significant positive and negative derivative signals within a water use event to determine if the event was a single appliance/fixture event, or a concurrent water use event requiring further disaggregation. Figure 3 outlines the process used in our model. If the event had only one increase or only one decrease, it was treated as a single event. This condition held true for shower events and human-controlled water use events, as long as the valve was either opened or closed in a continuous manner. If the event had more than one increase in addition to more than one decrease, the event was determined to contain multiple appliances or fixtures operating simultaneously and was categorized as a concurrent event. We then applied our disaggregation algorithm to separate and classify the water use events.

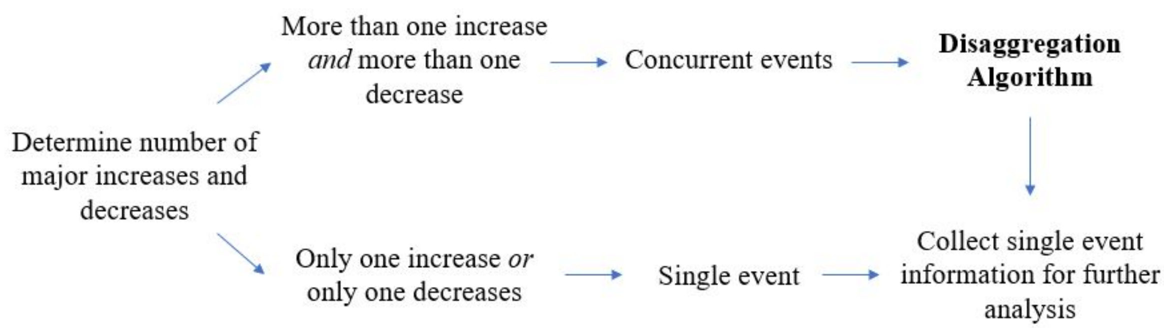

Figure 3: The process for determining concurrent events was based on the number of major increases and decreases within a water use event.

The concurrent event disaggregation process was based the assumption that, in instances other than shower events, the positive derivative signal of the event should be similar in magnitude to the negative derivative signal. This condition means that, in most cases, each positive derivative signal collected from the concurrent event should be matched to a negative signal. The first step in this process was to iterate through the list of major increases and match each item in this list to the major decrease closest in magnitude. The matched increase and decrease pairs were then removed from their respective lists and stored separately as start and end times of a single event. Because each positive signal could only be matched with one negative signal, the result depended on the order in which the positive and negative signals are matched. To optimize the amount of water use events correctly accounted for in the disaggregation process, the list of major increases was ordered from largest to smallest. Using this method, the event with the highest flow rate was matched first. This algorithm assured that the largest major increase was matched with the largest major decrease, improving the number of events accounted for by the model.

Because shower events contained two negative derivative signals, the algorithm responded differently to these types of events. It was important to include this exception because concurrent events are most likely to occur during long duration events, particularly showers. We implemented a test within the algorithm to determine if a shower event was likely to exist during a concurrent event. The bathtub had one of the largest flow rates of the water fixtures throughout the house. The only other fixtures within the home that operated at a flow rate above 15.2 L/min (4 gpm) included the washing machine and outdoor hose. Based on the training dataset, the washing machine did not use water continuously for greater than 300 seconds during normal wash cycles; it was a reasonable assumption that a shower might last for this duration or longer. Based on these assumptions, if the magnitude of the major increase was greater than 15.2 L/min, and the duration of the event was greater than 300 seconds, the event was determined to contain a shower. This checking-system test is outlined in Figure 4.

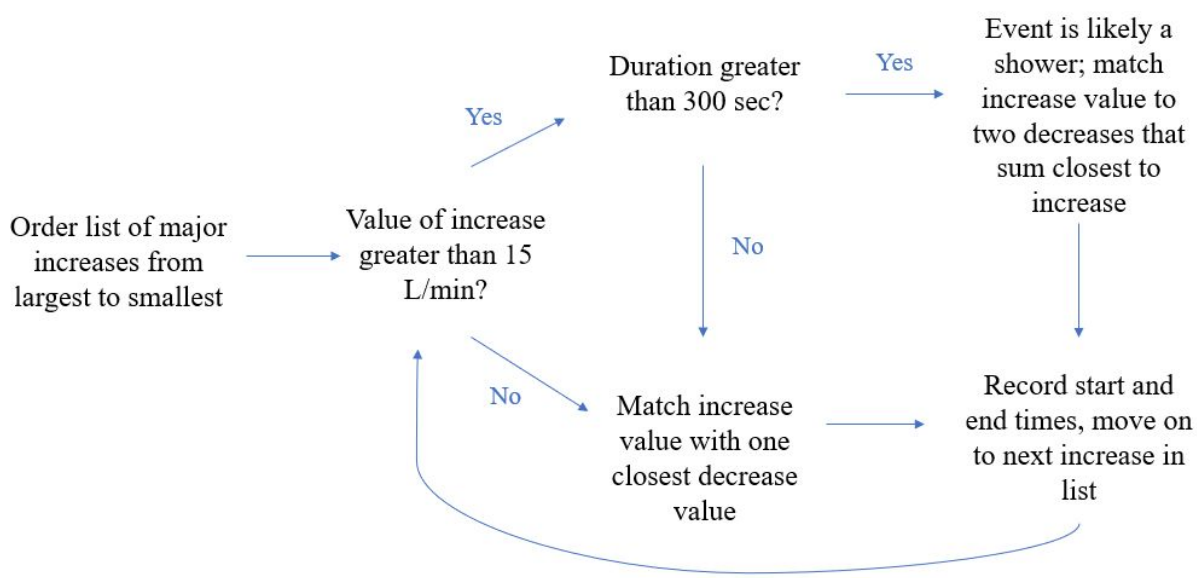

Figure 4: Matching the positive consecutive sums to the negative gradient signals required additional considerations. The algorithm behaved differently for shower events, due to the switch from the bathtub faucet to the showerhead during the event.

The algorithm iterated through the list of positive derivative signals until each item was matched with a negative signal closest in value. The start and end times were then recorded, based on when these increases and decreases occurred. Each isolated water use event was then treated as an individual event in the event classification process.

Appliance/Fixture Identification

We determined appliance/fixture end uses from the disaggregated water events using *k*-means clustering. The *k*-means clustering algorithm allocated each point (water use event) in the dataset to one of the predetermined clusters, while minimizing the size of the clusters. The average flow rate and duration were used as identifying factors for the appliance/fixture identification criteria. While higher dimensional *k*-means clustering was considered by including the start time, end time, and/or volume, the start and end times were determined to have little significance, while the volume was directly related to the average flow and duration. Therefore, we implemented 2-dimensional *k*-means clustering for the analysis, using the average flow rate and duration of the water use events as criteria.

Before performing the *k*-means clustering on the water event data, we removed events that likely represented leaks. By visualizing the average flow rate as a function of the event duration, it was clear that several events were low-flow, high-duration events, as shown in Figure 5. While leak detection is a valuable application of smart water meter data analysis, these events were removed before the clustering process for simplification purposes. The flow rate cutoff was adjusted from 0.11 L/min (0.03 gpm) to 0.38 L/min (0.1 gpm) to remove these low-flow, high-duration events, which were assigned to the "other" category to account for the water losses. To further process the data for clustering, we scaled the input average flow rate and duration data by dividing by the respective standard deviations, making the data unitless. This data scaling was an important step to ensure that the results of the clustering algorithm were not changed when analyzing the data in metric (liters) or U.S. customary (gallons) units.

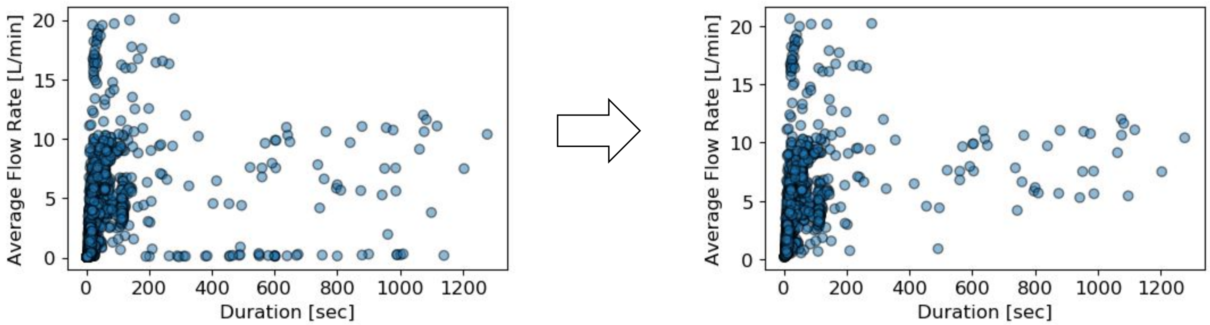

Figure 5: Low-flow, high-duration events, present at left, were removed for further analysis, shown right.

Results

Household-level trends

Figure 6 shows the daily average volume of water consumed for each month, representing average water consumption values only for days with consumption >38 L (>10 gal) as an indicator of occupancy. Results of the *t*-test for difference in mean household-level water use between months suggested no strong seasonal trend in the residents' water consumption ($p = 0.19$); that is, the statistical decision is to fail to reject the null hypothesis of no difference between means.

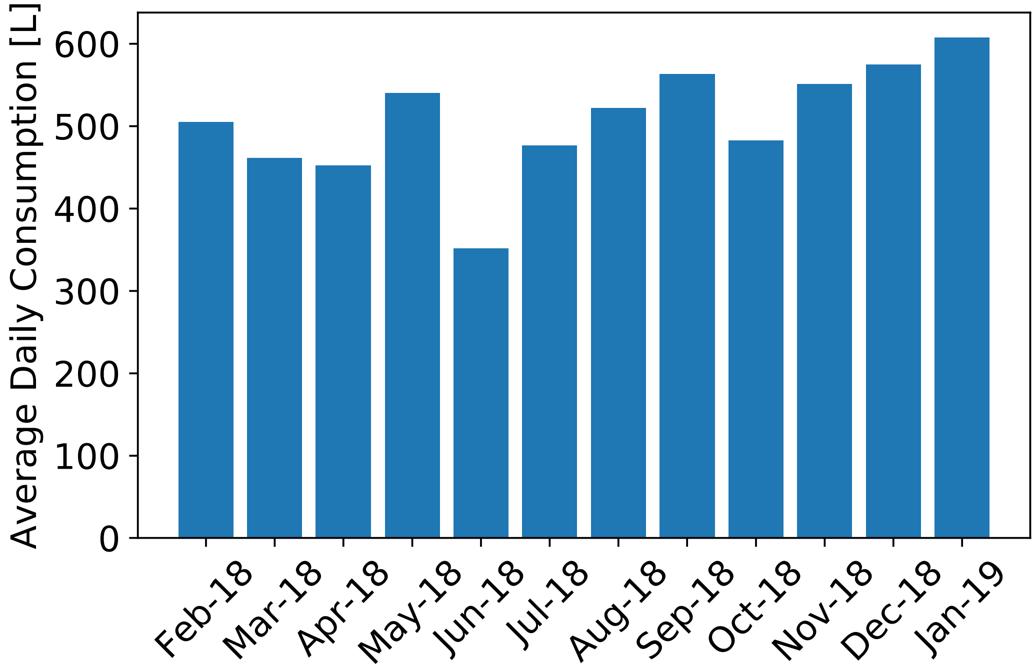

Figure 6: The study home did not show a strong seasonal water consumption trend. Data values represent average consumption for days in the month indicating occupancy (water consumption >38 L/day (>10 gal/day)).

While water consumption typically increases during the summer months due to increased outdoor use (29-30), this study household does not show such a trend. The residents of this house report using little outdoor water over the course of the analysis period. Based on the recorded data, a substantial drop in average consumption is observed during June 2018, which was statistically significant based on the *t*-test ($p < 0.01$); that is, the statistical decision is to reject the null hypothesis (in Equation 2) in favor of the alternate hypothesis of a statistically significant difference between means. This decrease, however, can likely be attributed to travel; while days with water consumption <38 L (<10 gal) were excluded, several days were recorded in June 2018 in which the residents consumed 110-190 L (30-50 gal), less than half of the typical daily average. Most of these days occurred just before or just after a period of several days of zero water consumption, suggesting these were the days that the residents departed or returned from travel and were only home and consuming water for part of the day. The rest of the monthly averages were not statistically different from each other based on a *t*-test (p = 0.19), suggesting no statistical evidence of a seasonal trend.

While there was little difference in water consumption behavior on the monthly level, daily averages for each day of the week suggest that the residents' water consumption varies on a weekly scale, as shown in Figure 7.

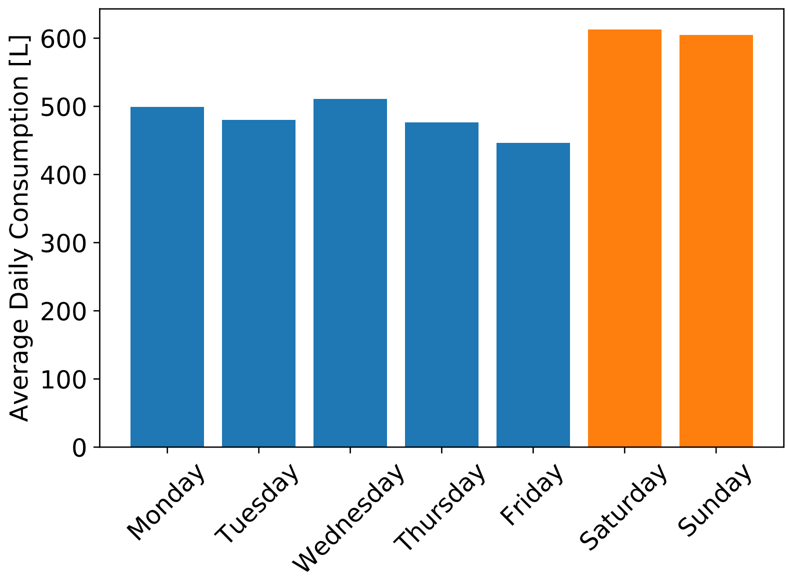

Figure 7: More water is consumed on average on weekend days than on weekdays.

The data reveal that more water is consumed on weekends than weekdays. The average daily consumption for all weekdays is 485 L (128 gal), with a standard deviation of 227 L (60 gal). The residents consumed 609 L (161 gal) on average on the weekend, with a slightly smaller standard deviation of 212 L (56 gal). A *t*-test determined this difference in means to be statistically significant ($p < 0.0001$). The differences in mean consumption for each weekday do not show statistical relevance (p = 0.19), nor do the Saturday and Sunday means (p = 0.47).

The residents not only consumed different amounts of water on weekends and weekdays, but the consumption profiles over the course of the day also differ. Daily time-of-use patterns were visualized based on 30-minute binning of water consumption data for the household, calculated using Equation 1. Peak time intervals were quantified as time periods in which the largest volume of water was consumed during the day. Residential water demand often has two peaks: one in the morning and another in the evening (2,46). To quantify these peak times for the study home, half-hour intervals were identified with the maximum volume of water consumed for both the morning (before 12:00 PM) and the evening (after 12:00 PM) time periods.

Figure 8 shows the volume of water consumed during each 30-minute interval throughout the day on a typical weekday and weekend day. These data are displayed for a typical weekday and typical weekend day instead of an average over each category because there was much more variance in how water is used over the weekends than weekdays. The weekday and weekend day shown in Figure 8 exhibited total water consumption values close to the respective averages, with the same morning and evening peak times as the respective medians.

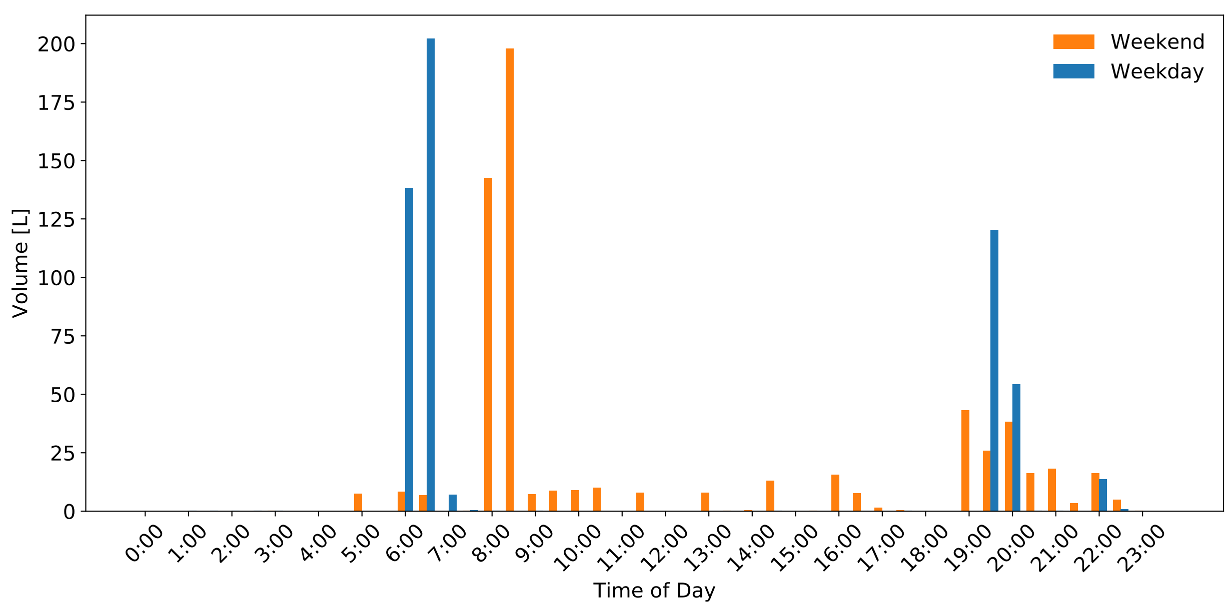

Figure 8: Water consumption exhibits different sub-daily temporal patterns on weekends and weekdays. The weekday peak time of use occurs in the morning, with the median time peak for weekdays between 6:30 AM and 7:00 AM throughout the year. Another smaller peak occurs at night, with no water used

during the day. The median evening peak for weekdays was between 7:30 PM and 8:00 PM (19:30 and 20:00). This sub-daily temporal pattern reflects a family that works during typical business hours and is home in the mornings and evenings.

The weekend day morning peak is about the same size as the weekday morning peak; however, the weekend day peak occurs about 2 hours later. The median peak morning interval over all of the weekends within the analysis period was between 8:30 AM and 9:00 AM. There was also another peak at night in which less volume was consumed than the evening weekday peak, but both evening peaks occur around the same time. The median evening weekend peak interval was between 7:00 PM and 7:30 PM (19:00 and 19:30). Additionally, water was used throughout the day on the weekends during the hours that the home was typically unoccupied during the week.

Appliance identification results

We applied the *k*-means clustering algorithm to the disaggregated data to produce four initial groups, as shown in Figure 9, based on the limited training data. Appliances and fixtures that displayed flow rate and duration data within the range of each group were identified to determine if further clustering of each of these initial groups was necessary. Group 1 was the lowest-volume, shortest-duration group of events. There were no identifiable appliances within the range of this cluster for average flow rate and duration, and these Group 1 events were determined to be human-controlled low-flow faucets. The Group 2 cluster was characterized by a larger average flow rate and longer duration than Group 1. This cluster contained events within the range of the three household toilets and the dishwasher. Group 3 contained the group of events with the highest average flow rates, and the appliances with characteristics consistent with this cluster included the washing machine and the outdoor hose. The long-duration events fell into Group 4, which primarily consisted of showers.

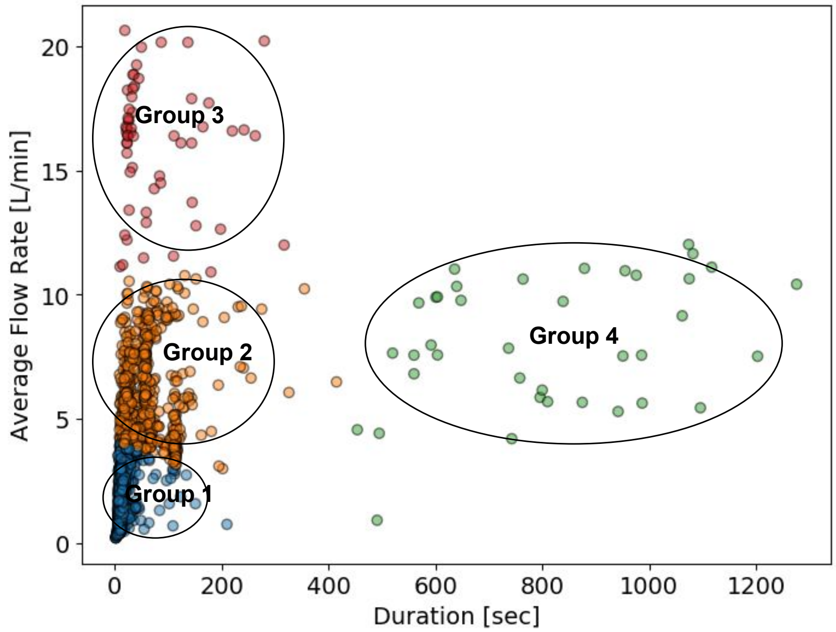

Figure 9: The data were divided into four initial groups: 1) low-flow faucets, 2) toilets and dishwasher, 3) washing machine and outdoors, and 4) showers.

The clustering analysis, which was performed separately on each month of data, classified the data into similar initial groups. This similar grouping suggests that the household consumed water in a similar manner throughout the year and is consistent with the lack of seasonal trend observed in overall water consumption. The cluster centers of Groups 1 and 2 showed the most consistency throughout the year, with small variances in duration and average flow. Group 3 showed the most variation in average flow rate with a standard deviation of 31.4 L (8.31 gal), which suggests variation in washing machine cycles and outdoor water usage but could also be due to disaggregation errors. Group 4, which consisted of showers, showed the most variation in duration with a standard deviation of 73 seconds due to the human-controlled nature of this type of event. The cluster centers for each group over the course of the analysis period are shown in Table 1.

Table 1: Cluster centers for initial groups using *k*-means clustering were fairly consistent for each month from February 2018 to January 2019.

	Group 1		**Group 2**		**Group 3**		**Group 4**	
	Duration [s]	Average flow [L/min]	Duration [s]	Average flow [L/min]	Duration [s]	Average flow [L/min]	Duration [s]	Average flow [L/min]
Feb	16.9	1.58	59.7	6.24	89.7	7.18	719	7.18
Mar	14.5	1.51	62.7	5.82	76.8	16.1	715	7.48
Apr	15.6	1.54	61.3	5.75	69.8	15.6	820	7.98
May	13.7	1.43	57.8	4.84	72.5	10.5	849	8.13
Jun	12.8	1.47	54.3	4.95	79.5	11.2	703	7.48
Jul	14.9	1.73	60.4	5.82	79.6	15.9	809	8.05
Aug	12.9	1.36	58.0	5.03	89.1	12.1	847	8.20
Sep	13.5	1.43	47.3	4.76	87.5	10.7	796	8.39
Oct	14.0	1.51	52.3	4.84	77.3	11.9	788	7.41
Nov	13.9	1.43	55.4	4.61	64.2	9.71	971	8.28
Dec	14.1	1.47	51.8	4.61	70.6	9.45	885	8.39
Jan	11.8	1.17	53.5	4.50	89.8	9.71	844	7.94

We further applied the *k*-means clustering algorithm to each initial group to distinguish sub-clusters in appliance/fixture end uses. The sub-clusters and appliance/fixture classifications associated with these clusters are shown in Figure 10, and appliances/fixtures categorized from within each group are shown in Table 2 based on limited training data. All events in Group 1 were classified as low-flow faucets, so no additional clustering was necessary for this group of events. The three household toilets and the dishwasher were within the range of Group 2; additional *k*-means clustering into 9 clusters was applied to this group. In response, there existed clusters with centers quite similar to the average flow rate and duration from the training data, and the events in these clusters were classified into the appropriate appliance/fixture category. The events in the other 5 clusters in this group were classified as human-controlled faucet events, further labeled as medium-flow faucets to differentiate from the low-flow faucets in Group 1. Next, we applied the clustering algorithm to Group 3, dividing this group into 3 clusters. One of these clusters was assigned as the washing machine based on both the cluster center and range of the data within this cluster. It should be noted that the particular model of washing machine in the home is programmed to fill with the appropriate amount of water based on the size of the load of laundry, and the duration and flow rate of each cycle was highly variable in the test dataset. However, most cycles had durations and average flow rates within the range of the cluster. The remaining cluster, shown in Figure 10, was classified as the "Other" category, including possible outdoor water events and disaggregation errors. The third cluster did not have characteristics comparable to any of the appliances and these events were classified as medium-flow faucets. Finally, no additional sub-clusters were applied to Group 4; all of these events were classified as showers based on the long duration and flow characteristics. The final appliance/fixture end uses found in each group are summarized in Table 2.

Table 2: Different appliances and fixtures were represented in the water end uses found in each group from *k*-means clustering.

Cluster Group	End-Uses
Group 1	Low-flow faucets
Group 2	Downstairs toilet, Upstairs toilet 1, Upstairs toilet 2, Dishwasher, Medium-flow faucets
Group 3	Washing Machine, Other
Group 4	Showers

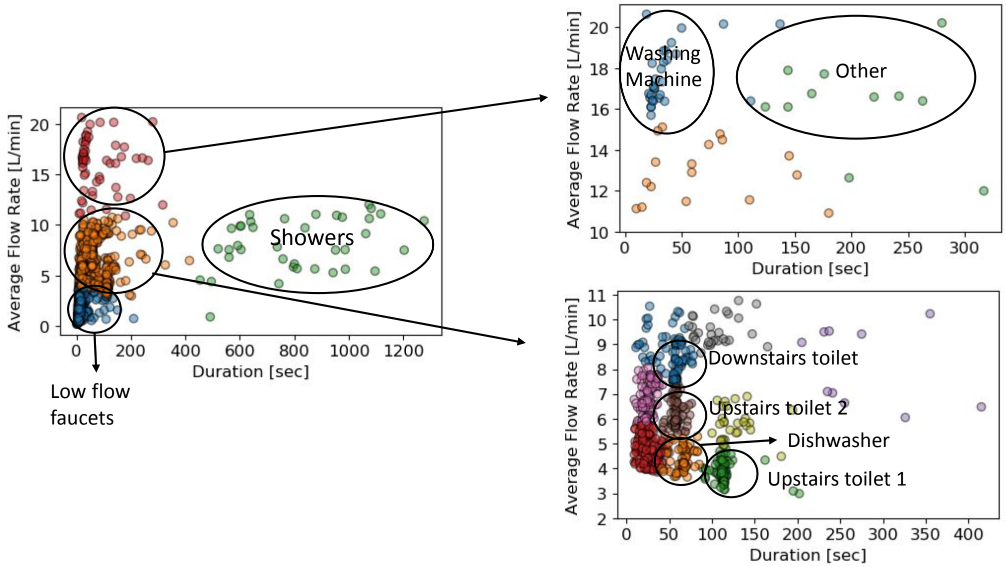

Figure 10: Most of the appliance/fixture end uses were found in sub-clusters of the initial clustering groups.

The water end-use profile of the home is shown in Figure 11. Most of the water in the home was consumed through human-controlled events, specifically the shower, which accounted for an estimated 39 percent of the overall household water consumption, at an average of 4,950 L (1,300 gal) per month. Cominola et al. (42) reported similar shower-dominated water end-use profiles based on data from Australia. Medium-flow and low-flow faucets were the second-highest form of water consumption with an estimated 32 percent contribution to the overall water footprint. Toilets were the next-largest consumer of water, accounting for 895 L (236 gal), or 7.1 percent of overall consumption per month on average. The washing machine and dishwasher were the lowest water consumers, accounting for 2.7 percent and 1.2 percent of water use, respectively; however, it is important to note that there is additional uncertainty associated with these appliances that have multiple cycles per use. The “Other” category consisted of pipe leaks, disaggregation errors, and outdoor uses; these categories consumed an average of 1,030 L (272 gal) per month.

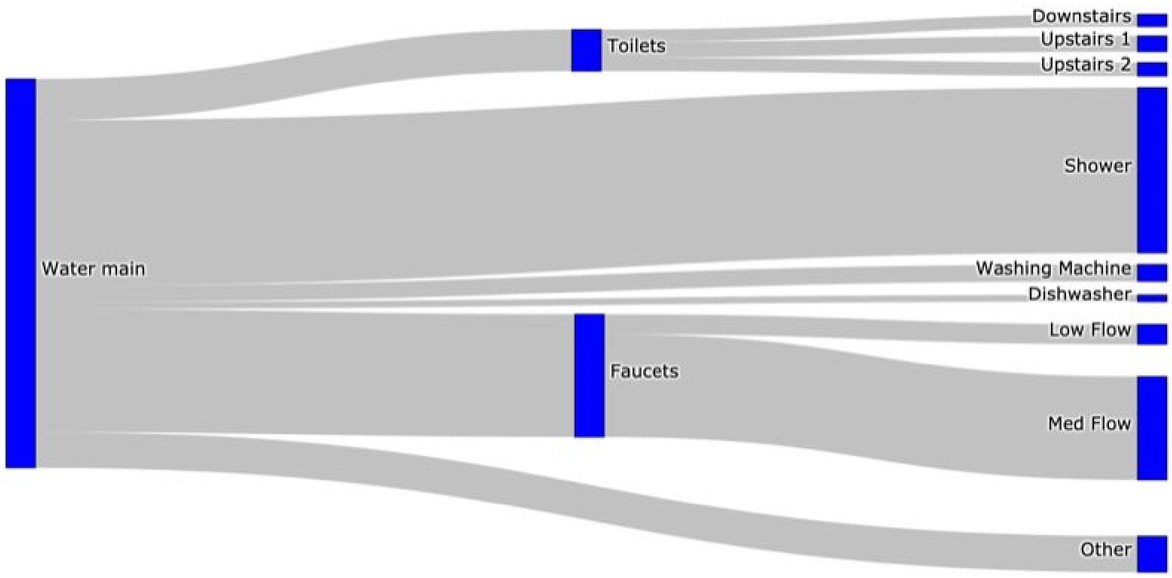

Figure 11: Most of the water in the study home was consumed through human-controlled events, specifically showers and other faucets.

Existing literature suggests that outdoor water consumption varies seasonally, peaking in the summer largely due to outdoor uses (29-30). The fact that the occupants of this study home report consuming little water for outdoor uses is supported by the lack of seasonal trend in water consumption behavior. Literature also suggests that while indoor water use shows little seasonal variation (47), others have shown that shower duration increases as temperature decreases (31), so indoor water consumption might increase during colder months. While seasonal differences in total water consumption were not statistically significant for this study home (p = 0.19), there was an increase in average daily consumption during the months of January and February, when temperatures were coldest. These results are displayed in Figure 12. Based on the water consumption profile for each month, the showers did account for less water in the warmer summer months, with the lowest shower consumption occurring during June 2018.

The colder months generally showed more water consumption from the shower, except for January 2019. This result could be due to other factors beyond the scope of this study, such as travel and data gaps.

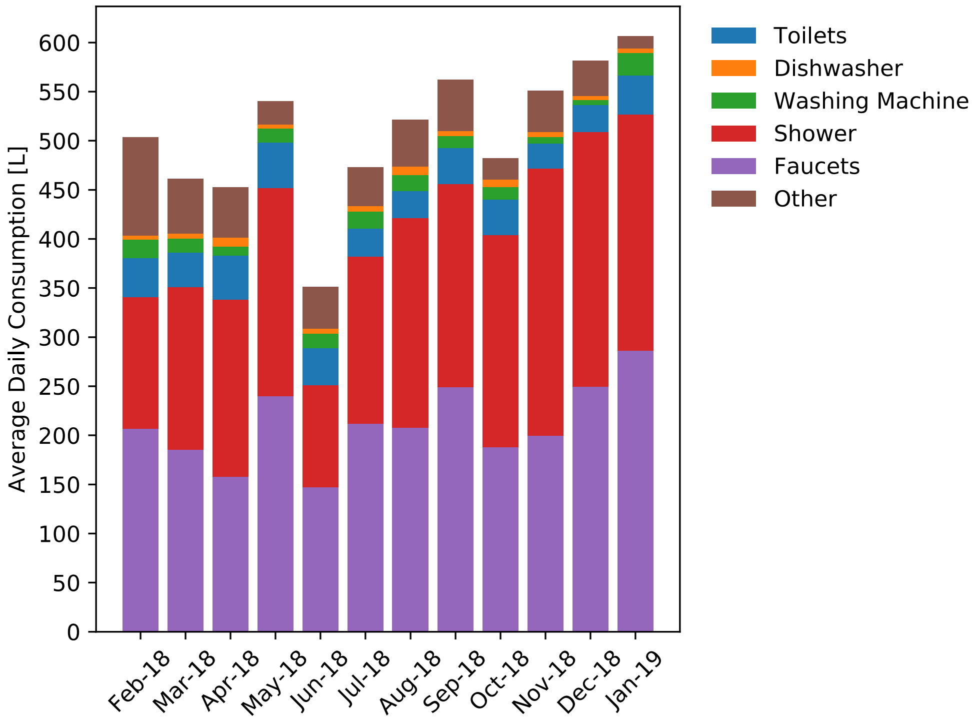

Figure 12: The water end-use profiles for each month show seasonal trends, particularly in shower use.

End-use disaggregation metrics

While training data were not required for this clustering analysis, information on how the appliances/fixtures operate was necessary for identifying appliances associated with observed water end use events. A training dataset was necessary to understand what appliances/fixtures were installed in the home and how each appliance/fixture operates. Table 3 shows the sample training data compared to the average cluster centers across all months for each appliance/fixture. It is also important to note special characteristics associated with each appliance/fixture. While the dishwasher and toilets are automatic appliances that consistently operate the same way, the washing machine is different in that it is programmed to fill with the appropriate amount of water based on the size of the load of laundry. Additionally, the washing machine has numerous cycles, each of which differ by load. By sampling the dataset, most of the washing machine cycles were within the range of the assigned cluster, but there were likely some cycles that were not accounted for in the process.

Table 3: Training data and disaggregated data characteristics show strong performance of the disaggregation algorithm.

Appliance/Fixture	**Training Data Appliance Characteristics**		**Disaggregated Data Coherence Average Cluster Centers**		**Disaggregated Data Frequency**
	Duration [s]	Average Flow [L/min]	Duration [s]	Average Flow [L/min]	Average Percent of Analysis Days Present in Cluster
Upstairs Toilet 1	105-108	3.82	113 ± 1.3	3.89 ± 0.09	76.0
Upstairs Toilet 2	55-56	6.08	58.0 ± 3.4	5.98 ± 0.02	70.4
Downstairs Toilet	54	7.67	58.8 ± 1.8	7.52 ± 0.09	80.9
Dishwasher	66.9	3.78	64.8 ± 3.8	4.17 ± 0.04	53.8
Washing Machine	23-116	1.89-20.0	36.2 ± 9.1	16.5 ± 0.15	74.1
Shower	128-133	7.18-9.83	854 ± 131	7.92 ± 0.11	83.6

Disaggregating water use data into end uses introduces uncertainty in the absence of sub-metering measurements, which can be intrusive for household occupants (40) and can change occupant behavior (48). Large-scale smart water metering systems often lack sub-metered data and instead use pre-trained disaggregation algorithms (42). Since unsupervised learning methods do not require labeled data, specific qualitative and quantitative metrics have been proposed to assess results (49). In Table 3, we quantify the metrics of coherence (spread of the examples in a cluster) and frequency (percent of days covered by an activity) (49). Accounting for the unknown uses in the "Other" category, 91.8 percent of the incoming water to the home was categorized into appliance/fixture end uses through this process, representing the significance metric proposed by Cardell-Oliver (49). This amount of categorization is consistent with other models, which typically have a success rate of around 80 percent for appliance/fixture identification (11).

There were some additional uncertainties associated with this analysis that came from utilizing *k*-means clustering to classify water use events. It is important to note that because the *k*-means clustering algorithm aims to minimize the size of the cluster centroids, it assigns points to the closest centroid. This approach is inexact, and it is not always possible to obtain clusters containing points belonging exclusively to one appliance. It is also possible for a human-controlled faucet event to have the same average flow rate and duration of an automatic appliance, especially the appliances in Group 2, where most of the events were medium-flow faucets and likely overlap with the toilets and dishwasher clusters. It is also impossible to know the optimal number of clusters in the data by using this algorithm, so human labor is required to manually determine the clusters that exist, especially when human-controlled events such as faucets often do not appear in clear clusters. These uncertainties, however, are a tradeoff of using an unsupervised learning method in which training data are not required. In some reported instances in literature, unsupervised classifiers performed better than supervised classifiers (50).

The frequency metrics for the individual appliances give some insight into the accuracy of the clustering algorithm. It is likely that the toilets and shower are underestimated using this method since these events are expected to occur every day in a four-person household. The toilet clusters are found within Group 2 of initial clusters, which consistently has the largest number of water events, and the inter-cluster density of the sub-clusters in this group can decrease the accuracy of the *k*-clustering (51). Additionally, toilet flushes are commonly concurrent events. Because toilets are in the medium flow category, the disaggregation algorithm is more likely to miss the derivative signal if it is combined with a lower flow event, which could change the average flow and duration enough to move the event out of the appropriate cluster. Because the shower cluster is classified as all of Group 4, the discrepancy occurs during the initial clustering of the four groups and the unaccounted showers are likely categorized within the medium-sized clusters as well based on the flow rate of these events. This likely misclassification is especially true for showers shorter in duration, as these points might lie closer to the Group 2 centroid, particularly if other longer showers occurred within the same month.

Broader implications

Water conservation and efficiency implications

The overall daily average water consumption in the study home over the period of analysis (February 2018 to January 2019) was 519 L (137 gal) per day for the 4-person household. The USGS (2015) estimates per person water consumption to average 303-379 L (80-100 gal) of water per day within the household, so the residents in this study consume less than expected. This result could be due to a few factors. Sharing appliances such as dishwashers and washing machines among four people can increase efficiency and reduce per-person consumption. Cooking for multiple people is also more water efficient than cooking individual meals. Additionally, this particular household reports they do not attribute much water to outdoor uses, which is estimated to be the highest water consumer in the residential sector in the United States (6,52).

Much of the residents' water consumption occurs during typical morning and evening utility-scale peak demand hours (2,46). Decreasing consumption during this time on a widespread scale could contribute to lowering overall peak demand for the local utility and reduce pressure on existing water infrastructure. Supplying feedback to residents that consume most of their water during peak demand times might encourage behavior change and ultimately lower peak demand. Launching this type of study in more homes in the area could provide more insight into possible measures to reduce peak water demand.

While the household's overall per-person water consumption is below the U.S. average, the end-use analysis suggests that there are still opportunities for improved water conservation and efficiency. Most of the water in the home was consumed through human-controlled events, specifically showers and medium-flow faucets. Because human-controlled events depend heavily on the user's behavior, these end uses present the most opportunity for behavior change. Showers are the primary-identifiable end use and the largest area of potential conservation in this home. The average shower duration over the course of the analysis period is 855 seconds, consuming a volume of 112 L (29.8 gal). This result suggests that reducing the duration of showers by a few minutes could save 14,000 L (3,700 gal) over the course of the year, based on current showerhead flow rates. Similarly, lower-flow showerheads could be retrofitted into the study home as a focused water efficiency approach. It should also be noted that showers are large energy consumers because of the energy required to heat the water (53-55), so water conservation and efficiency associated with showers is also an energy-saving measure.

While changes in behavior can lead to significant water savings, the water efficiency of the appliances/fixtures is also relevant. Using low-flow fixtures and upgrading to water-efficient appliances is a method of water efficiency that does not require behavior change. To address how the appliances/fixtures in the home performed in terms of efficiency, the performance of each appliance based on the meter data was compared to the manufacturers' ratings. The appliances/fixtures throughout the home were documented during the initial meter installation regarding manufacturer and model, and this information was used to find the flow rates or volume of water consumed as stated by the manufacturer. Some of the fixtures had the water use printed on them, including the toilets and the showerheads; however, most rating data were found in the respective appliance/fixture manuals online. The manufacturers' ratings were compared to the measured water use from the training data and disaggregated data from the smart water meter in Table SI1 in the Supporting Information.

Insight into how appliances and fixtures actually function can inform water demand management strategies. The deployment of smart water meters has changed these strategies because now specific household appliances/fixtures as well as human behaviors can be monitored with fine resolution data (11,42). When comparing the actual use to the manufacturers' rating, it is easier to identify where appliances are functioning incorrectly, or rather not up to standards. For example, if a toilet is using more water per flush than the manufacturer reports it should, a homeowner can target this toilet to be replaced or fixed. This finding also signifies that home occupants cannot rely on manufacturer-provided ratings alone to determine their water use. In this study home, for example, a toilet rated at 6.06 Lpf (1.60 gpf) uses 7.04 Lpf (1.86 gpf) in practice, introducing questions about the difference in operating conditions in manufacturers' rating tests versus in-home conditions. With more homes installing smart meters, additional data can be collected in support of better understanding in-home conditions.

While some companies give ranges for how much water their appliances use, it is interesting to see where the actual measurements fall when in use in a home. In the case of the study home's dishwasher, the measured readings were below the range that the company reported. The training data volume of 11.6 L (3.07 gal) per dishwasher cycle was outside the manufacturer-reported range of 11.7-27.2 L (3.1-7.2 gal). For the washing machine, the measured water use fell around the center of the manufacturer-reported range, shown in Table SI1 in the Supporting Information. These results motivate more widespread study regarding where most appliances/fixtures typically fall within manufacturer-provided water use ranges.

Water-use feedback and data concerns

Disaggregating and classifying water events obtained from residential smart water meter data reveals detailed information about how water is consumed within the home. Understanding the overall water consumption profile of the home presents opportunities for improved residential water conservation and efficiency and long-term water resource sustainability, since many individuals underestimate water use (53-54). Several studies have used smart water meter data to pinpoint opportunities for improved efficiency within the household. Cardell-Oliver et al. (56) used smart meters to identify trends in high-magnitude water consumption behaviors for conservation implications. Schultz et al. (57) used water meters to provide personalized feedback on normative behavior to promote water conservation. In a study by Sønderlund et al. (25), 45% of participants stated that the availability of high-resolution water use data would encourage them to conserve water, suggesting that the widespread implementation of smart water meters to provide such information could effectively reduce water consumption in the residential sector.

While access to detailed residential water consumption information has been successfully used for conservation and efficiency purposes in small-scale studies, the effectiveness of this feedback is uncertain and large-scale implementation could present additional challenges. Factors such as lifestyle, social practices, household values, and socio-demographic factors also play a role in water consumption behavior such that behavior changes based on feedback from water meters are also likely to vary. Aitken et al. (58) found that water conservation depends highly on the values of the household, but education and targeted information promoting water conservation can effectively change individual behavior. Social practices and lifestyle also influence water consumption and may affect a user's willingness to adjust their behavior (59). Socio-demographic factors affect water consumption, with evidence of correlation between household income and outdoor water use, as well as property size and overall water consumption (6,60). Willis et al. (61) found that household income also has an effect on appliance end uses, as well as the feasibility of switching to water-efficient appliances.

In addition to uncertainty of the effectiveness of feedback from smart water meters, privacy concerns can also inhibit the widespread implementation of these water meters. Beckel et al. (62) found that information such as employment status can be revealed through smart meter data as a result of access to high temporal resolution data that reveal water use by time of day. Analysis of the data from our test home revealed likely work schedules and travel periods. Behavioral analysis with smart meter data could introduce privacy concerns, and there would need to be measures taken to protect consumer privacy if smart meters were implemented on a larger scale; however, effective communication strategies such as transparent public engagement can enhance the success of smart metering at a larger scale (59). Allowing others to access water use information can pose privacy and possibly safety concerns. These concerns among the public could be a significant barrier to large scale smart water meter implementation. Cities and utilities aiming to implement smart water meters to improve water conservation and efficiency should take action to protect the privacy of the households. Taking steps to protect consumer data and openly communicate the purpose of data collection and analysis could help gain public support. While anonymizing the data might make it difficult to provide direct feedback to the residents, such approaches might be an option for water utilities interested in using smart meters for infrastructure planning or water supply planning.

Conclusion

Our residential water use results suggest that the study residence is below the expected water consumption for a 4-person household in the United States. Outdoor water use is a large contributor to residential water consumption and might be the largest opportunity for improvement in terms of water conservation in this sector in general. However, for our study home in particular, showers present the largest opportunity for water conservation and efficiency.

Analysis of 1-second resolution smart water meter data provided detailed insight into the water consumption patterns of the study home. Seasonal trends for this household indicate low outdoor water usage, which is likely a large factor in the household's lower-than-average water consumption. The residents consume more water on average over the weekends, but this consumption is more spread out throughout the day and the peak household demand is typically lower on these days. Widespread measurement of water consumption could be beneficial for predicting peak residential water demand (2), motivating additional smart water metering demonstrations.

Our smart water meter data disaggregation estimated appliance/fixture end uses based on derivative signals of the flow rate, which occur within a few seconds, to pinpoint when water events began and ended. We used *k*-means clustering, an unsupervised machine learning method, to classify the disaggregated water events into appliance/fixture end uses using a small training dataset. Most available algorithms require extensive training data to calibrate the model, a significant barrier to the implementation of these studies on a broader scale. By using water meter data with more refined temporal resolution, our approach classified about 92% of the end uses to appliance/fixture events.

The results of this study can be used to provide feedback to the residents of the home to encourage water conservation and efficiency by pinpointing end-uses that consume large quantities of water. Feedback based on these results has the potential to alter the behavior of the residents in terms of water consumption. Installing more water-efficient appliances/fixtures also has the potential to reduce the household's overall water consumption. If the availability of smart water meters were more widespread in the residential sector, the actual performance of appliances/fixtures might be better understood to advance water sustainability at the household level.

Data Availability

Example smart water meter data used in this analysis are available as .csv files at https://stillwell.cee.illinois.edu/data/.

Conflicts of Interest

There are no conflicts to declare.

Acknowledgments

Thanks to Micah Stickling, former Undergraduate Research Assistant, for installing the smart water meter and conducting the initial data visualization. The custom ally® water meter was provided by Sensus. This work was supported by the National Science Foundation, grant CBET-1847404; the opinions, findings, and conclusions or recommendations expressed here at those of the authors and do not necessarily reflect the views of the National Science Foundation. Additional support was provided by the Research Experience for Undergraduates program in Civil and Environmental Engineering at the University of Illinois at Urbana-Champaign.

Author contributions: G.M.B. created the disaggregation methodology and conducted the analysis; A.R.C. collected manufacturers' appliance/fixture data and assisted with the analysis; G.M.B. and A.S.S. collected the smart water meter data and formulated the study; A.S.S. supervised the study; all authors contributed to writing the manuscript.

References

(1) Jepson WE, Wutich A, Collins SM, Boateng GO, Young SL. Progress in household water insecurity metrics: a cross-disciplinary approach. *Wiley Interdisciplinary Reviews: Water*. 2017;4(3): e1214.

(2) Omaghomi T, Buchberger S, Cole D, Hewitt J, Wolfe T. Probability of Water Fixture Use during Peak Hour in Residential Buildings. *Journal of Water Resources Planning and Management*. 2020;146(5): 04020027.

(3) Grafton RQ, Ward MB, To H, Kompas T. Determinants of residential water consumption: Evidence and analysis from a 10-country household survey. *Water Resources Research*. 2011;47(8): W08537.

(4) Basani M, Isham J, Reilly B. The Determinants of Water Connection and Water Consumption: Empirical Evidence from a Cambodian Survey. *World Development*. 2008;36(5): 953-958.

(5) Shan Y, Yang L, Perren K, Zhang Y. Household water consumption: insight from a survey in Greece and Poland. *Procedia Engineering*. 2015;119: 1409-1418.

(6) Mayer PW, DeOreo WB. *Residential End Uses of Water*. Boulder, CO: Aquacraft, Inc. Water Engineering and Management; 1999.

(7) Froehlich J, Larson E, Saba E, Campbell T, Atlas L, Fogarty J, Patel S. A longitudinal study of pressure sensing to infer real-world water usage events in the home. *International conference on pervasive computing*. Springer, Berlin, Heidelberg; 2011;50-69.

(8) Evans RP, Blotter JD, Stephens AG. Flow rate measurements using flow-induced pipe vibration. *J. Fluids Eng*. 2004;126(2): 280-285.

(9) Mori M, Tezuka K, Tezuki H. *U.S. Patent Application No. 10/584,318*; 2007.

(10) Kowalski M, Marshallsay D. Using measured microcomponent data to model the impact of water conservation strategies on the diurnal consumption profile. *Water Science and Technology: Water Supply*. 2005;5(3-4): 145-150.

(11) Cominola A, Giuliani M, Piga D, Castelletti A, Rizzoli AE. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. *Environmental Modelling & Software*. 2015;72: 198-214.

(12) Stewart RA, Willis R, Giurco D, Panuwatwanich K, Capati G. Web-based knowledge management system: linking smart metering to the future of urban water planning. *Australian Planner*. 2010;47(2): 66-74. Available from: doi:10.1080/07293681003767769.

(13) Makki AA, Stewart RA, Beal CD, Panuwatwanich K. Novel bottom-up water demand forecasting model: Revealing the determinants, drivers and predictors of residential indoor end-use consumption. *Resources, Conservation and Recycling*. 2015;95:15-37. Available from: doi:10.1016/j.resconrec.2014.11.009.

(14) Willis RM, Stewart RA, Giurco DP, Talebpourm MR, Mousavinejad A. End use water consumption in households: impact of socio-demographic factors and efficient devices. *Journal of Cleaner Production*. 2013;60: 107-115.

(15) House-Peters LA, Chang H. Modeling the impact of land use and climate change on neighborhood-scale evaporation and nighttime cooling: A surface energy balance approach. *Landscape and urban planning*. 2011;103(2): 139-155.

(16) Lucas SA, Coombes PJ, Sharma AK, 2010. The impact of diurnal water use patterns, demand management and rainwater tanks on water supply network design. *Water Supply*. 2010;10(1): 69-80.

(17) Beal C, Stewart RA. Identifying residential water end-uses underpinning peak day and peak hour demand. *Journal of Water Resources Planning Management.* 2014;140(7): 04014008.

(18) Gurnung TR, Stewart RA, Sharma AK, Beal CD. Smart meters for enhanced water supply network modelling and infrastructure planning. *Resources, Conservation and Recycling*. 2014;90: 34-50.

(19) Beal CD, Gurung TR, Stewart RA. Demand-side management for supply-side efficiency: Modeling tailored strategies for reducing peak residential water demand. *Sustainable Production and Consumption*. 2016;6:1-11. Available from: doi:10.1016/j.spc.2015.11.005.

(20) Pesantez JE, Berglund EZ, Kaza N. Smart meters data for modeling and forecasting water demand at the user-level. *Environmental Modelling & Software*. 2020;125: 104633.

(21) Mukheibir P, Stewart RA, Giurco D, O'Halloran K. Understanding non-registration in domestic water meters: Implications for meter replacement strategies. *Water*. 2012;32(8): 95-100.

(22) Arregui F, Cabrera E, Cobacho R, García-Serra J. Key factors affecting water meter accuracy. *IWA Water Loss Conference*. Halifax, Canada; 2005.

(23) Arregui FJ, Cabrera E, Cobacho R, Garcia-Serra J. (2006). Reducing apparent losses caused by meters inaccuracies. *Water Practice and Technology*. 2006;1(4): wpt2006093.

(24) Britton TC, Stewart RA, O'Halloran KR. Smart metering: enabler for rapid and effective post meter leakage identification and water loss management. *Journal of Cleaner Production.* 2013;54: 166-176.

(25) Sønderlund AL, Smith JR, Hutton CJ, Kapelan J, Savic D. Effectiveness of Smart Meter-Based Consumption Feedback in Curbing Household Water Use: Knowns and Unknowns. *Journal of Water Resources Planning and Management.* 2016;142(12): 04016060.

(26) Savić D, Vamvakeridou-Lyroudia L, Kapelan Z. Smart Meters, Smart Water, Smart Societies: The iWIDGET Project. *Procedia Engineering*. 2014;89:1105-1112. Available from: doi:10.1016/j.proeng.2014.11.231.

(27) Hastie AG, Chini CM, Stillwell AS. (2020). A mass balance approach to urban water analysis using multi-resolution data. *Journal of Industrial Ecology*. Early view. Available from: doi:10.1111/jiec.12995.

(28) Gutzler DS, Nims JS. Interannual Variability of Water Demand and Summer Climate in Albuquerque, New Mexico. *Journal of Applied Meteorology*. 2015;44: 1777-1787.

(29) Duncan HP, Mitchell VG. A stochastic demand generator for domestic water use. *Proceedings of Water Down Under 2008*. 725; 2008.

(30) Balling RC, Gober P. Climate Variability and Residential Water Use in the City of Phoenix, Arizona. *Journal of Applied Meteorology and Climatology*. 2007;46(7): 1130-1137.

(31) Rathnayaka K, Malano H, Maheepala S, Gerorge B, Nawarathna B, Arora M, Roberts P. Seasonal Demand Dynamics of Residential Water End-Uses. *Water*. 2015;7: 202-216.

(32) Hart GW. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*. 1992;80(12): 1870-1891.

(33) Leeb SB, Shaw SR, Kirtley JL. Transient event detection in spectral envelope estimates for nonintrusive load monitoring. *IEEE Transactions on Power Delivery*. 1995;10(3): 1200-1210.

(34) Parson O, Ghosh S, Weal M, Rogers A. Non-intrusive load monitoring using prior models of general appliance types. *Twenty-Sixth AAAI Conference on Artificial Intelligence*; 2012.

(35) Batra N, Kelly J, Parson O, Dutta H, Knottenbelt W, Rogers A, Srivastava M. NILMTK: an open source toolkit for non-intrusive load monitoring. *Proceedings of the 5th international conference on Future energy systems*. 2014; 265-276.

(36) Love BC. Comparing supervised and unsupervised category learning. *Psychonomic bulletin & review*. 2002;9(4): 829-835.

(37) DeOreo WB, Heaney JP, Mayer PW. Flow trace analysis to access water use. *Journal-American Water Works Association*. 1996;88(1): 79-90.

(38) Beal C, Stewart R, Huang T, Rey E. Applying smart metering technology to disaggregate residential water end uses in South-East Queensland. *Water: Journal of the Australian Water Association*. 2011;38(1): 80-84.

(39) Nguyen KA, Stewart RA, Zhang H. An intelligent pattern recognition model to automate the categorisation of residential water end-use events. *Environmental Modelling & Software*. 2013;47: 108-127.

(40) Di Mauro A, Di Nardo A, Francesco Santonastaso G, Venticinque S. Development of an IoT System for the Generation of a Database of Residential Water End-Use Consumption Time Series. *Environmental Science Proceedings*. 2020;2: 20.

(41) Cominola A, Giuliani M, Piga D, Castelletti A, Rizzoli AE. A Hybrid Signature-based iterative disaggregation algorithm for Non-Intrusive Load Monitoring. *Applied Energy*. 2017;185(1): 331-344.

(42) Cominola A, Nguyen K, Giuliani M, Stewart RA, Maier HR, Castelletti A. Data Mining to Uncover Heterogeneous Water Use Behaviors from Smart Meter Data. *Water Resources Research*. 2019;55: 9315-9333.

(43) Laspidou C, Papageorgiou E, Kokkinos K, Sahu S, Gupta A, Tassiulas L. Exploring patterns in water consumption by clustering. *Procedia Engineering*. 2015;119: 1439-1446.

(44) Cheifetz N, Noumir Z, Samé A, Sandraz AC, Féliers C, Heim V. Modeling and clustering water demand patterns from real-world smart meter data. *Drinking Water Engineering and Science*. 2017;10: 75-82.

(45) Mounce SR, Furnass WR, Goya E, Hawkins M, Boxall JB. Clustering and classification of aggregated smart meter data to better understand how demand patterns relate to customer type. *Proceedings of Computing and Control for the Water Industry (CCWI 2016)*; 2016.

(46) Buchberger S, Omaghomi T, Wolfe T, Hewitt J, Cole D. *Peak Water Demand Study*. International Association of Plumbing and Mechanical Officials; 2015. Available from: https://www.iapmo.org/media/3857/peak-water-demand-study-executive-summary.pdf.

(47) Abdallah AM, Rosenberg DE. Heterogeneous residential water and energy linkages and implications for conservation and management. *Journal of Water Resources Planning and Management*. 2014;140(3): 288-297.

(48) Schwartz D, Fischhoff B, Krishnamurti T, Sowell F. The Hawthorne effect and energy awareness. *Proceedings of the National Academy of Sciences of the United States of America*. 2013;110(38): 15242-15246.

(49) Cardell-Oliver R. Discovering Water Use Activities for Smart Metering. *IEEE 8th International Conference on Intelligent Sensors, Sensor Networks and Information Processing.* 2013;171-176.

(50) Becker V, Kleiminger W. Exploring zero-training algorithms for occupancy detection based on smart meter measurements. *Computer Science – Research and Development*. 2018;33: 25-36.

(51) Wu S, Chow TWS. Clustering of the self-organizing map using a clustering validity index based on inter-cluster and intra-cluster density. Pattern Recognition. 2004;37(2): 175-188.

(52) U.S. Geologic Survey (USGS). *Domestic Water Use*. Available from https://www.usgs.gov/mission-areas/water-resources/science/domestic-water-use [Accessed August 14, 2020].

(53) Attari SZ. Perceptions of water use. *Proceedings of the National Academy of Sciences*. 2014;111(14): 5129-5134.

(54) Inskeep BD, Attari SZ. The water short list: The most effective actions US households can take to curb water use. *Environment: Science and policy for sustainable development*. 2014;56(4): 4-15.

(55) Chini CM, Schreiber KL, Barker ZA, Stillwell AS. Quantifying Energy and Water Savings in the U.S. Residential Sector. *Environmental Science & Technology*. 2016;50(17): 9003-9012.

(56) Cardell-Oliver R, Wang J, Gigney H. Smart Meter Analytics to Pinpoint Opportunities for Reducing Household Water Use. *Journal of Water Resources Planning and Management.* 2016;142(6): 04016007.

(57) Schultz WP, Messina A, Tronu G. Personalized Normative Feedback and the Moderating Role of Personal Norms: A Field Experiment to Reduce Residential Water Consumption. *Environment and Behavior.* 2014;48(5): 686-710.

(58) Aitken CK. Residential Water Use: Predicting and Reducing Consumption. *Journal of Applied Social Psychology*. 1994;24(2): 136-158.

(59) Michalec A, Hayes H, Longhurst J, Tudgey D. Enhancing the communication potential of smart metering for energy and water. *Utilities Policy*. 2019;56: 33-40.

(60) Loh M, Coghlan P. *Domestic Water Study*. Perth, WA: Water Corporation; 2003. Available from: https://www.water.wa.gov.au/__data/assets/pdf_file/0016/5029/42338.pdf.

(61) Willis RM, Stewart RA, Panuwatwanich K, Williams PR, Hollingsworth AL. Quantifying the influence of environmental and water conservation attitudes on household end use water consumption. *Journal of Environmental Management*. 2011;92(8): 1996-2009.

(62) Beckel C, Sadamori L, Thorsten S, Santini S. Revealing household characteristics from smart meter data. *Energy*. 2014;78: 397-410.