Using Rule Mining to Understand Appliance Energy Consumption Patterns

Sami Rollins
Department of Computer Science
University of San Francisco
Email: srollins@cs.usfca.edu

Nilanjan Banerjee
Department of Computer Science
and Electrical Engineering
University of Maryland, Baltimore County
Email: nilanb@umbc.edu

Abstract—**Managing energy in the home is key to creating a sustainable future for our society. More tools are increasingly available to measure home energy usage, however these tools provide little insight into questions such as why an appliance consumes more energy than normal or what kinds of behavioral changes might be most likely to reduce energy usage in the home. To answer these questions, a deeper understanding of the causal factors that influence energy usage is necessary. In this work, we conduct a broad study of factors that influence energy consumption of individual devices in the home. Our first contribution is collection of a context-rich data set from six homes across the United States. The second contribution of this work is a set of insights into key factors influencing energy usage derived by the novel application of a rule mining algorithm to identify significant associations between energy usage and four key features: hour of the day, day of the week, use of other appliances in the home, and user-supplied annotations of activities such as working or cooking. Our analysis confirms our hypothesis that, though most devices show a regular pattern of daily or weekly use, this is not true for all devices. Associations that relate use of two different devices in the same home are often stronger, and are observed for nearly 25% of device uses. Overall, we observe that the associations derived from the first five weeks of data in our data set are sufficient to explain nearly 70% of the device uses in the subsequent five weeks of data, and over 90% of the associations identified during the first five weeks recur in the latter portion of the data set. The associations identified by our approach may be used to to aid in end-user applications that heighten awareness and encourage energy savings, improve energy disaggregation algorithms, or even detect anomalous uses that may signal problems in aging-in-place homes.**

I. INTRODUCTION

A key component of creating a sustainable world is managing energy in the home. Pervasive technologies such as smart meters and home energy measurement tools are poised to transform energy management by identifying energy waste and optimizing usage. Presently, these tools provide a plethora of information, though users must still draw their own insights about usage based on analysis of raw data such as power consumption readings. Ad hoc analysis by users combined with application of static recommendations, such as reducing the temperature of a home heating system, can reduce energy consumption. Providing a less burdensome mechanism for a user to understand *why* consumption is high or low at a given time, however, is necessary to fully realize the potential of the home energy management domain.

In this work, we undertake a study of the causal factors that impact energy usage at the granularity of individual appliances

in the home. Understanding how the energy consumption of an appliance is related to the activity that the user is performing, the time at which it is used, or other devices that may be used concurrently is a key element of many applications that seek to help users reduce or better understand energy usage. It has been demonstrated, for example, that users better understand context-annotated energy data in comparison to raw energy consumption data [1], [2]. Augmenting an interface such as FigureEnergy [1]—a web interface that helps users to better understand energy usage—with contextual cues can lead to heightened awareness and future changes to energy usage patterns. As another example, it has been shown that having apriori knowledge of devices that are used concurrently can improve the accuracy of energy disaggregation algorithms [3]. Automatically deriving causal ties between devices such as a washer and dryer, therefore, is a promising approach to improve non-intrusive load monitoring. Finally, connecting use of a device to the time when it is normally used can help a user to conceptualize how she could use the appliance less, either fewer hours per day or fewer times per week or month. Moreover, the ability to automatically derive information about the normal usage of a device can also enable more intelligent systems that alert users of anomalous behavior. Such systems can also be applied to aging-in-place and assistive healthcare facilities, where it is critical to automatically infer abnormal patient behavior [4], [5].

The first contribution of this work is the collection of a broad and context-rich data set that provides a novel set of features used for our analysis and evaluation. We have instrumented six homes, located in three different areas of the United States, with appliance-level energy meters that collect raw power readings and report them to a server for postprocessing. From this raw data set we are able to extract features such as time of use and day of use. Because time-based features are not applicable to all appliances, we also propose the use of two additional, user-centric, contextual features: we explore the relationship between use of different appliances in the home and we also collect in situ activity annotations from users via a smart phone application. These annotations connect usage of appliances to the activities, for example cooking or working, that are accomplished through use of the appliance.

The second contribution of this work is set of key insights into the factors affecting energy consumption derived by applying a rule mining approach in this domain. Our algorithm uses the information-theoretic metric of JMeasure [6] to identify
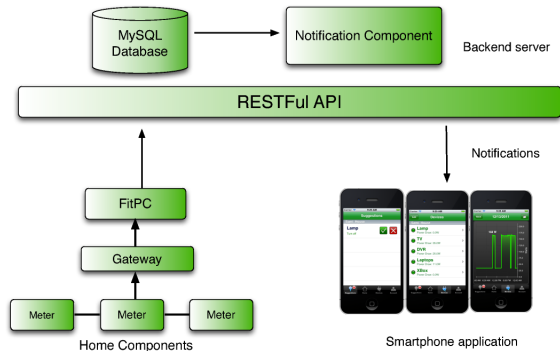
Fig. 1. Green Homes architecture.

rules describing associations between use of a device and the following features: hour of day, day of week, use of other devices in the home, and activities performed by the user. Our analysis confirms our hypothesis that, though most devices show a regular pattern of daily or weekly use, this is not true for all devices. Associations that relate use of two different devices in the same home are often stronger, and are observed for nearly 25% of device uses. Associations with user activities are infrequently observed given our limited data set, however it is clear that devices that are used irregularly demonstrate this type of association more frequently than other devices. Overall, we observe that the associations derived from the first five weeks of data in our data set are sufficient to explain nearly 70% of the device uses in the subsequent five weeks of data. Moreover, over 90% of the associations identified during the first five weeks recur in the latter portion of the data set demonstrating that the JMeasure-based approach is able to extract meaningful associations. Using examples from two homes, we discuss the types of rules inferred by our system, and their possible applications in understanding energy consumption in the home, detecting anomalous behavior in assistive care homes, and energy disaggregation for non-intrusive load monitoring.

II. BACKGROUND AND DATA COLLECTION

The primary goal of this work is to identify causal factors that influence energy consumption of individual devices in the home. We hypothesize that time-based factors, including the hour of the day in which a device is used or the day of the week on which a device is used, are strongly associated with many devices. A living room lamp, for example, may consistently be used each day during a few hours in the evening. Initial analysis, however, confirms that many devices are used more irregularly. A clothes washer may be used only once per week and not consistently on the same days. We, therefore, explore two additional features that influence device usage. First, we consider whether two devices in the same home, for example a clothes washer and a clothes dryer, are frequently used together. Second, we consider whether there is an association between use of a device and the self-reported activity that a user accomplishes when using the device, for instance *cooking dinner* may be strongly associated with use of the oven.

To test our hypotheses, we have deployed a measurement infrastructure that collects a broad and contextually-rich data set from six homes across the United States. Our home energy measurement system, Green Homes [7], [8], collects raw energy consumption data from several devices in each home, and uses a novel *in situ* approach for soliciting user annotations to describe activities performed. The result is a collection of traces that describe extremely varied usage. In this section, we present an overview of the data collection system and features of our data set as they relate to the hypotheses tested in this work. Further details regarding the underlying system architecture and deployment can be found in [7], [8].

Figure 1 illustrates the architecture of the Green Homes system. In each home, we deploy between five and ten off-the-shelf energy meters that communicate with a dual-radio gateway [9]. Also in each home is a client component that polls the energy meters every 30 seconds and reports the raw power draw readings for each device to a centralized server. Data on the server, including graphs of past usage and current device status, may be accessed by a web or smartphone application. Our current deployment supports both iOS and Android applications.

Raw power draw readings for each device, along with the timestamp at which each reading was taken, are sufficient to explore how the time-based factors like hour of the day and day of the week influence energy consumption in each home. Using this raw data we can also examine whether use of certain devices in the home is strongly correlated with use of other devices in the home. These data, however, do not provide us with enough information to determine what the user may have been doing when the energy usage of a particular device increased. To capture this information, we use a novel annotation collection component [10] that sends a notification to a user's phone when it determines that the user may be performing a new activity. The user may respond by providing an annotation describing her current activity. In addition to asking the user to provide free-form text describing his or her current activity, we also require users to provide a category to represent the activity from the following set of categories: `cooking`, `entertainment`, `work`, `chores`, and `other`. Logging activity categories in addition to raw annotations has two advantages. First, it helps us collect more uniform data across subjects as raw annotations can vary widely across users. Secondly, the raw annotations logged were sometimes obvious and not useful. For instance, one of our subjects logged "Watching TV" when he switched on his television—an event that triggered the notification and can be inferred without any user annotation. We chose these specific categories since home users understand them well [1].

The notification component uses an aggressive power profiling algorithm to identify when energy usage has increased indicating a possible user context change. We chose not to ask users to annotate decreases in energy consumption to minimize the intrusiveness of the system, and because the primary uses of the annotations are centered around understanding periods of high energy consumption. The system is designed to collect annotations at as fine a granularity as possible, therefore the goal is to determine when a device is in use. For some devices, for example lamps, it is straightforward to use a basic threshold to determine that the device has transitioned from a 0 watt off state to a 20 watt on state. Other devices have less clear power states, for instance in one home we have a meter

attached to a power strip that powers several devices including a television and a cable box. When the television is on, the power consumption of the strip can increase by more than 20 watts when the user changes channels on his television.

To address this challenge, our power profiling algorithm uses the DBSCAN clustering algorithm [11] to produce a unique profile for each device in our study. DBSCAN is a density-based algorithm that identifies clusters while excluding noise. DBSCAN is an ideal choice for this application as it does not require the number of clusters to be provided as input and it can be implemented very efficiently, particularly for one-dimensional data such as ours. The algorithm takes as input two parameters—*eps* specifies the neighborhood of a point, which in our case represents the minimum number of watts separating two distinct power states. We experimentally determined that 2 watts yields the best results. The second parameter is *minpts*, which represents the minimum number of points required to form a cluster. We use 10 in our algorithm.

The power profiles generated describe from 2–5 different states for the devices in the study. Lamps, for example, often have two states: 0, and a second state with a range of about 10 watts from the low end to the high end of the state. The television and cable box setup described above has five states: 0, 45–65, 81–90, 146–264, and 274–296 watts. This likely identifies idle and active states for each of the devices connected to the power strip. Anecdotally, this approach is able to capture significant changes in state while excluding fluctuations that occur during the time a device is active and in use.

Once the profile is generated, we execute a final step of the algorithm to evaluate whether a device is non-interactive and represents only background load for the user. In the case of a refrigerator, for example, the transition into a higher power state likely does not indicate a change in user context. A refrigerator might run twice per hour every hour of the day, hence it is useful to exclude a non-interactive device from notifications. To identify background loads, we apply a heuristic that will classify a device as non-interactive if in more than 80% of the hours for which data was reported for the device there was a change in power state for the device. Effectively, if a device transitions between power states in more than 19 hours of the day then the device is likely not manually controlled by the user and will not trigger notifications. The notification component also determines whether it is necessary to retrain to produce an updated power profile. Retraining will occur if the number of power draw readings that fall outside of the states identified in a device's current power profile exceeds a threshold, in our current implementation 25, or if a state in the power profile should be eliminated because it has not been visited for two weeks.

Though our data collection is ongoing, we have extracted traces from a ten-week period from November 1, 2012 through February 9, 2013 to use for the analysis in this work. Table I provides an overview of the features of this data set, and we note that subjects 4–6 are researchers involved with this project. However, since the goal of this paper is to perform a broad study of contextual cues and appliance energy consumption, we believe that data collected from subjects involved in the project minimally impacts our conclusions. We collect data from a broad set of devices including lamps, which are

likely to have regular usage patterns, DVRs, which have very consistent usage patterns, refrigerators, which are regular and out of the user's control, and clothes washers, which are used relatively infrequently and with varying consistency. The mean number of daily notifications per home varies and is impacted both by the number and type of devices measured in each home. In effect, the number of daily notifications indicates that we are able to identify between three and eight important energy consumption events each day. We present the mean daily annotations for the first five weeks of the data collection period separate from the last five weeks as our analysis uses the first half of the data to train our rule mining algorithm, described in the next section, and the remaining five weeks to evaluate the algorithm. We note that while the number of annotations provided by subjects during the test weeks is consistent with other approaches for collecting annotations [1], the number of annotations declined over time. This suggests that this approach for collecting annotations is effective for a short period, but may need to be augmented, for example with a web-based interface to provide after-the-fact annotations, in order to encourage ongoing use. Even so, we are able to effectively use the set of annotations collected to demonstrate the feasibility of using context to understand appliance usage patterns.

III. Using Rule Mining to Understand Device Usage

In this section, we first describe our data preprocessing methods and then describe our rule mining approach for extracting causal rules between context and device usage. For our analysis, we consider associations between device usage the following contextual features: hour of the day, day of the week, other devices used concurrently, and user activities collected using the annotation system described in §II.

A. Data preprocessing and context extraction

We collect appliance-level data at a fine granularity of once every 30 seconds using our measurement infrastructure, however this resolution is too fine grained for extracting association rules. Moreover, our raw data contain several instances of erroneous readings uploaded by malfunctioning energy meters. We, therefore, first perform a data cleaning and data chunking step illustrated in Figure 2 (a). In this step, we filter out instances when data points were missing, and only consider contiguous data chunks for our analysis. Additionally, we filter any negative power consumption values and any power values above 2KW. We have confirmed that none of the devices in our deployment are rated above 2KW, thus any readings in this range are erroneous.

Using only the valid and contiguous raw power readings, we next chunk data into one hour windows and create a profile for each device that specifies whether it is *on* or *off* during each window. To determine whether a device is *on* during a time slot we apply a simple heuristic: in the given hour, if the mean power draw of a device exceeds the overall hourly mean by 5% the device is considered *on*, otherwise it is considered *off*. Note that even if a device is used for only a small portion of the one-hour window this approach does identify the *on* state, provided the device is not used for a small portion of every one-hour window in the data set.

Subject ID	Measured Devices	Mean Daily Notifications All Days	Mean Daily Annotations Train	Test
1	space heater, PCs, laptop, treadmill, fridge, lamp, TV	8	3	.13
2	laptop, microwave, coffee maker, router, PC, monitor, fridge	8	1.5	.6
3	toaster, dishwasher, TVs, lamps, microwave, clothes washer	5.8	1.2	1
4	microwave, washer/dryer, TV/electronics, lamp, fridge	3.8	2.2	.7
5	fridge, tv/electronics, space heater, router, washer, desk	4.8	2.9	.5
6	lamp, TV, DVR, laptops, XBox	3.4	1.6	.5

TABLE I. THE TABLE SHOWS THE DEVICES MEASURED IN EACH HOME ALONG WITH THE THE MEAN DAILY NOTIFICATIONS AND MEAN DAILY ANNOTATIONS DURING THE FIRST AND SECOND FIVE WEEKS OF DATA COLLECTION.
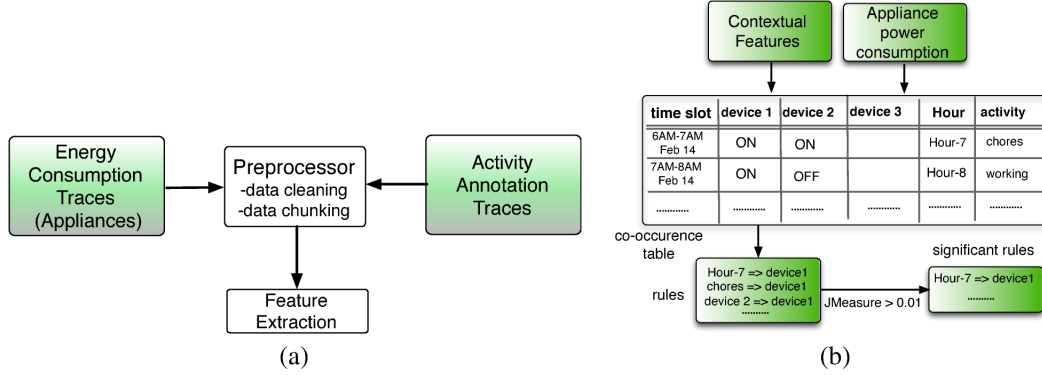

Fig. 2. (a) The activity and appliance energy consumption traces are cleaned for anomalous and missing data and then chunked into one hour slots. The preprocessor also determines whether an appliance was ON during each time slot. (b) The contextual data as well as the raw energy consumption data are grouped into a co-occurrence table, and rules are extracted from the table. We use an information-theoretic measure, called JMeasure, to determine rules that are statistically significant.

There are two reasons we choose one hour as our window size. The first reason is simply because we use the hour of the day as a relevant feature. Secondly, our goal is to identify the smallest window possible to ensure that we do not capture associations that do not really exist, or falsely inflate the probability of an appliance being used if it is only used for a small portion of the block we consider. On the other hand, we want the block size to be large enough to capture associations that might exist between two appliances that are used in sequence, like a washer and dryer. Initial experiments confirmed that blocks of 30 minutes do, in fact, miss associations between a washer and dryer in our study suggesting that 30 minutes is too short. Conversely, 2-hour blocks do not capture associations with activity annotations that happen infrequently suggesting that 2 hours is too large. One-hour blocks provide a reasonable compromise and are able to capture both the associations missed by the other time blocks.

After the preprocessing step, we extract relevant features from the data (see Figure 2 (a)). Time of the day and day of the week features are both common and easy to extract. Recall that we also consider two novel features. Devices used concurrently can be determined easily from the trace of on and off periods for both devices. Finally, user activities are extracted by determining whether an annotation occurred during a particular one-hour time slot.

B. Rule mining algorithm

The goal of our rule mining algorithm is to derive meaningful insights by identifying the most significant associations between context and device usage. Though rule mining has

been used to identify associations in a variety of domains, for example network trace mining [12], our work takes an in-depth look at its utility specifically in the domain of home energy usage. In this context, the algorithm must be capable of capturing associations between observations that occur infrequently, for example use of a clothes washer and dryer that occurs only once per week. It must also be capable of excluding associations between long-running appliances, for instance a space heater, and other appliances that happen to run concurrently but have no causal relationship.

Identifying candidate rules: Once feature extraction is complete, the first step of the rule mining algorithm is to identify candidate rules. Figure 2 (b) illustrates the co-occurrence matrix we derive by identifying all observations in each possible time slot. A row of the matrix corresponds to a one-hour block. The columns of the matrix store the state of appliances in the home, any activity that the user has annotated in that block, the hour of the day, and day of the week, for that block. The plausible set of rules is extracted from the matrix in the form $X \implies Y$ where X is a feature such as time of the day and Y is the state of the appliance. For instance, `chores` $\implies$ `device 1` is a valid candidate rule in Figure 2 (b) because on Feb 14 from 6AM-7AM `chores` was provided as an annotation and `device 1` was on.

Determining rule significance: The full set of rules derived is large and only a small subset of them are likely to indicate meaningful associations. The next step of the algorithm, therefore, is to identify which rules are significant. Information-theoretic measures have been successfully used to determine significance of association rules in previous work [13] because they measure the *content* contained in a rule. In this work,

we employ JMeasure [6], an information-theoretic significance metric. JMeasure for a rule of the form $X \implies Y$ quantifies the mutual information $I(Y; X = 1)$ expressed by the following equation:

$$\text{JMeasure } (X \implies Y) = P(X \wedge Y) \cdot \log(\frac{P(Y|X)}{P(Y)}) +$$
$$P(X \wedge Y') \cdot \log(\frac{P(Y'|X)}{P(Y')})$$

There are two reasons JMeasure is appropriate for our application. First, the value of JMeasure depends on the direction of the rule. Since our goal is to understand what context has a strong association with device usage, the rule direction is an important consideration. Second, and more importantly, JMeasure for a rule $X \implies Y$ does not solely depend on the frequency of co-occurrence of X and Y or the conditional probability $P(X|Y)$. In fact, it is a product of the frequency of co-occurrence ($P(X \wedge Y)$ and $P(X \wedge Y')$) and a measure of whether the occurrence of X improves the chance of occurrence of Y ($\log(\frac{P(Y|X)}{P(Y)})$ and $\log(\frac{P(Y'|X)}{P(Y')})$), an entropy measure. Hence, it accounts for cases where X and Y do not occur often, but when they do occur they occur together, like the washer/dryer example described above. It also accounts for cases when the conditional probability is low but frequency of co-occurrence is high. We use a default threshold of 0.01 to filter significant rules from non-significant ones. The intuition behind using the cut-off is similar to that used in previous network rule mining approaches [12], however, we evaluate the inferred rules for different JMeasure thresholds in §IV.

IV. Associations Impacting Energy Consumption

In this section, we present the rich and diverse set of associations we have discovered by applying our rule mining approach across the data collected from the homes in our study.

A. Experimental Setup

The experiments in this section use the ten weeks of data collected from six homes as discussed in §II. The first five weeks of raw power measurements and user annotations are used as the training set from which a set of associations are identified. The last five weeks are then used as a test set to verify that the associations discovered are, in fact, associations that recur. For this analysis, we discard any devices for which there are not at least ten uses during the testing period or ten uses during the training period. In most cases, these devices simply reported a power draw of 0 during the entire experiment period, though this also filtered out a DVR that is always on but has a power draw that is always within 5% of the mean.

During the test period, if a device X is identified as *on* we consult the rules to determine whether any of the known associations apply. We first look for hour of day and day of week associations. If there is no association $h \implies X$ and no association $d \implies X$ for the given hour h and day d, then we consider whether an annotation rule applies. If there is a rule $A \implies X$ and the user has provided the annotation A during the test period, then we apply the annotation rule. If no annotation rule applies then we consider device rules $Y \implies X$ for all devices Y.
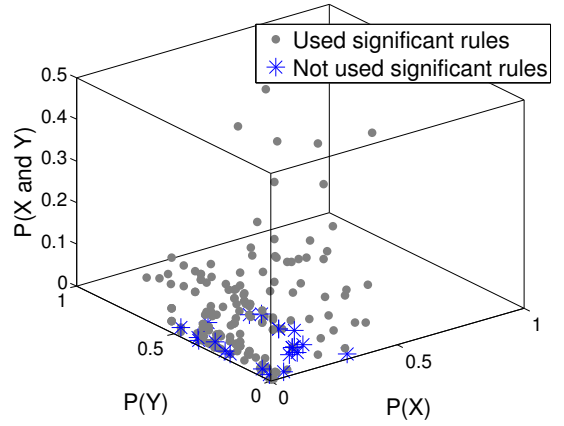

Fig. 3. For all significant rules (JMeasure > 0.01) of the form $X \implies Y$, the figure plots the joint probability of X and Y, probability of occurrence of X, and probability of occurrence of Y. The grey dots correspond to rules that were applied at least once in the testing data and blue stars correspond to rules that were not applied. Several of the applied rules correspond to low values of $P(X \wedge Y)$ and low values of $P(X|Y) = \frac{P(X \wedge Y)}{P(Y)}$, which implies that metrics that just consider frequency of co-occurrence or conditional probabilities will miss important associations.

B. Results

We first consider whether the rule mining approach we employ is able to identify useful associations for the devices in our study. Figure 3 presents, for all rules $X \implies Y$ where JMeasure > 0.01, $P(X)$, $P(Y)$, and $P(X \wedge Y)$. A grey dot corresponds to an association that was identified during the training period and was also seen during the testing period. Our primary finding from this figure is that *over 90% of the associations we identify as important during the training period are subsequently observed during the testing period, demonstrating that the JMeasure-based approach is able to extract meaningful and recurring associations.* Moreover, a large portion of the used rules have $P(Y) < .5$ and $P(X)$ and $P(X \wedge Y)$ significantly less than .5, demonstrating that our approach is able to identify associations between events even when the probability that either event will occur is low. As a more specific example, a microwave in our study has a probability of use of only 8% and the annotation `cooking` has a probability of occurrence of only 2% in the home where the microwave is located. The probability that they will occur together is only 1%, however our algorithm identifies the association between the microwave and `cooking` and observes this association twice during the testing period.

Subject	Time	Annotation	Device
1	100	14.3	57
2	100	0	80
3	100	43	100
4	57	43	86
5	83	33	33
6	100	100	100

TABLE II.   THE TABLE SHOWS THE PERCENTAGE OF DEVICES THAT YIELD ASSOCIATIONS OF EACH TYPE FOR EACH HOME.

Table II provides a more comprehensive view of the types associations identified for the devices in each home. The table shows, for each subject, the percentage of devices in the home that have associations with each of the three features considered: Time, Annotation, and Device. Note that

we have grouped the day of the week and hour of the day features into the Time column. The data demonstrates that *a wide range of associations are discovered and there is variation across homes with respect to the types of associations found.* Unsurprisingly, time associations are the most common. Almost all devices in the study show a regular daily or weekly usage pattern. Subject 4, however, demonstrates much less predictable time-based patterns than other subjects with only 57% of devices having an association with either a day of week or hour of day. Associations with annotations are the least common, though this is likely impacted by the limited number of annotations that we were able to collect in our pilot study. Finally, a significant portion of devices in the study are frequently used in concert with other devices in the same home, demonstrating that our approach may be used to identify ensembles of devices useful for better understanding energy requirements and waste, as well as aiding in problems such as energy disaggregation.
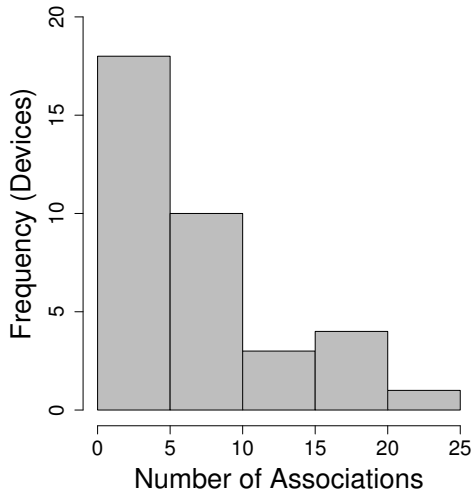

Fig. 4.   Histogram of the number of associations identified for the devices across all homes.

While Table II shows that a significant portion of devices in each home have associations of all three types, Figure 4 provides a closer look at the total number of associations identified for each device in the study. The figure shows a histogram of the number of devices that have a given number of associations. Our first observation is that *most devices have a small number of associations.* Commonly, devices have fewer than five associations, reiterating the finding that *our approach extracts only the most meaningful associations.* Devices that have large numbers of associations are those such as the PC of subject 1; in this case the device has 15 hour associations, three day associations, two annotation associations, and two device associations. We are currently exploring extensions to our approach that will further prioritize the associations identified for these kinds of devices.

Our final experiment considers which types of associations are observed most frequently. Recall that over 90% of associations identified during the training period are also observed during the testing period. In Figure 5 we report, for each use during the test period, the type of association identified for the use. The leftmost bar, for example, shows that nearly 70% of uses during the testing period exhibit an association identified during the training period. *The majority of device uses show time-based associations,* however *associations with use of other devices are also frequently observed.* Recall that our algorithm first looks for a time-based association, therefore if a device use exhibits both a time-based and a device-based association only the time-based association will be reported here. Figure 5 also explores the impact of the JMeasure threshold. As we would expect, *as the JMeasure threshold increases, fewer associations are discovered by our algorithm.* The time-based associations decrease much more dramatically, though the device-based associations actually show an increase for some values of the JMeasure threshold. In these cases, the device associations existed but were not reported for lower values of JMeasure threshold because a time-based association existed. This demonstrates that *associations that relate use of two devices are are much stronger than time-based associations.*

We note that Figure 5 shows that annotation associations are rarely observed. This is not surprising given the low number of annotations in our data set, however we hypothesize that annotation associations are most likely to be present for devices that are used irregularly. To test this hypothesis, Figure 6 compares the associations observed during testing for all devices and for infrequently used devices—those that have a probability of use less than 10% during the training period. The results confirm our hypothesis: *for devices that are used infrequently, annotation-based associations occur more often.*

C. Case Study

To more clearly illustrate the types of associations extracted by our approach, Figures 7 and 8 provide a graphical representation of the extracted associations for subjects 1 and 4. Each figure shows the devices measured in the subjects' home and, for each device, any associations identified. A directional arrow from X to Y corresponds to an association $X \implies Y$. The text next to each device shows any annotation associations identified. Finally, blue shaded boxes illustrate associations between use of the device and the given day or hour. For example, a blue shaded box in the first square of the Hour row indicates an association between the hour midnight to 1AM, and the use of the device. Though we have included associations from only two homes, other homes show similar patterns.

As ongoing work, we are looking at how to apply these associations to applications including energy disaggregation, energy awareness and conservation, and detection of abnormal energy consumption. First, prior work has demonstrated that energy visualizations [1], particularly those that integrate contextual information [1], [2], heighten awareness and understanding of home energy consumption, which can lead to decreased overall usage. As a simple example, we can use an association such as `chores` $\implies$ `washer` and `dryer` to help the user to better understand how much energy is typically consumed by this *activity* rather than showing the underlying devices used to accomplish the task. Similarly, an association between `cooking` and microwave can be used to demonstrate, for instance, that energy used for the `cooking` activity is
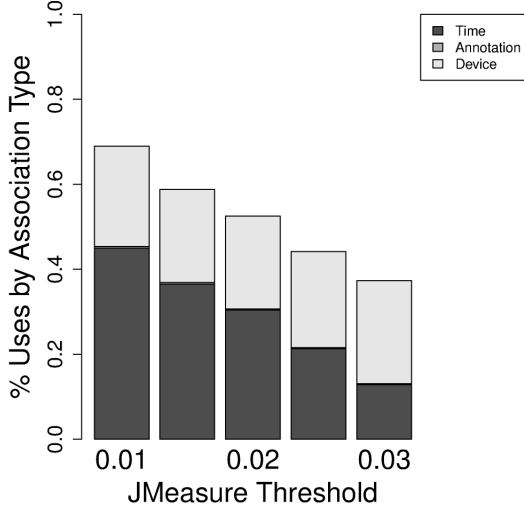
Fig. 5. For JMeasure threshold values .01 to .03, the percentage of uses that demonstrate the given association.
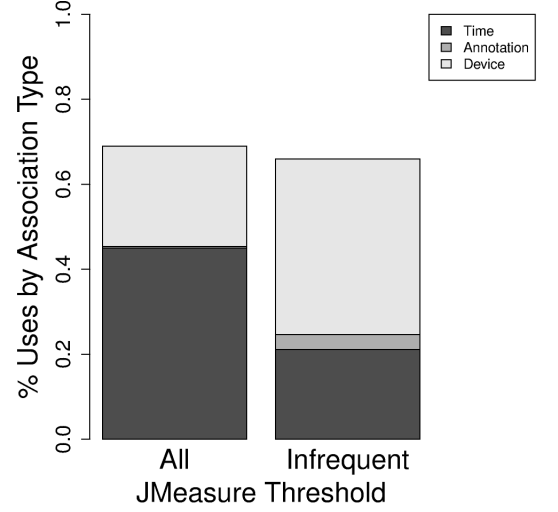

Fig. 6. The percentage of uses that demonstrate the given association for devices with probability of use lower than .1 during training.
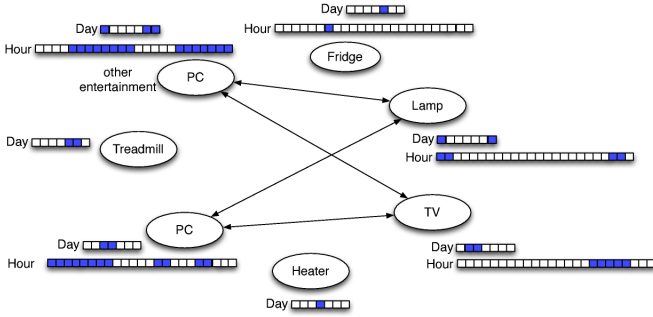

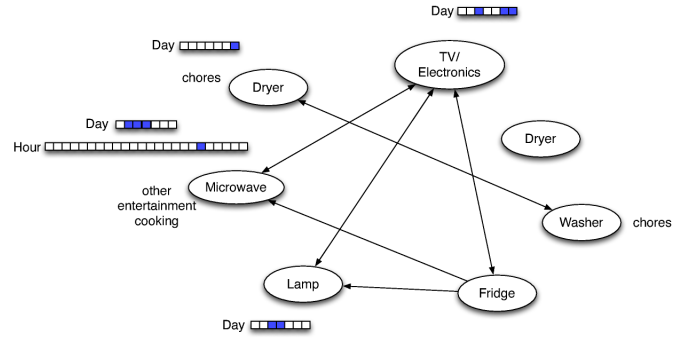Fig. 7. The associations identified for subject 1.

Fig. 8. The associations identified for subject 4.

lower in some instances because the microwave is used for a shorter period of time, or a lower-power appliance is used in favor of the microwave.

Prior work has also shown that the accuracy of energy disaggregation algorithms can be substantially improved by integrating dependencies between appliances [3]. Given a signal collected using a non-intrusive single point monitoring device [14], extracting the signal generated by an ensemble of devices is often a more tractable problem that identifying each individual device. Our approach reveals several ensembles that may be used for this purpose, for example the PCs in the home of subject 1 are both frequently used in concert with the lamp and the TV. Similarly, subject 4 typically uses his washer and dryer in concert. Though the latter example is not surprising, we note that our approach automatically derives these associations, even for unlabeled devices.

Finally, detecting anomalous appliance usage, for example identifying an appliance that was left on accidentally, is useful with respect to energy savings, and also has applications to domains such as aging-in-place homes or remote care for

patients with Alzheimer's [15], [4]. Associations with hour of the day and day of the week are a clear starting point for such an application. Appliances such as subject 1's TV and lamp have clear hour-of-day patterns. Suppose a TV is normally on at 10am. If it is discovered off during that hour it may signal that the user has failed to get out of bed and, thus, may be ill and need assistance. Similarly, if a lamp is normally turned off by 11pm but remains on well into the night this may indicate that the user has either forgotten to turn off the device or may need help from a family member or care giver.

V. RELATED WORK

Our paper builds on previous work on context-aware energy management, rule mining techniques, methods for annotating energy data, and demand-response systems in the home.

Context-aware energy management: Augmenting energy data with contextual information can help in several applications including demand prediction of home appliances and whole home energy consumption [16], [17], [18], [19], [20]. Most of the focus in this area has been on collecting data on

environmental context such as temperature, humidity, weather conditions and ambient light. Collecting these additional dimensions of data require deploying additional sensors [21], [22], [23], [24], [25]. Activities such as cooking, entertainment, and work have also been shown to be directly related to home energy consumption and provide additional usage context [26]. Users can relate to and understand such activities [1] and energy conservation recommendations made akin to these activities have a higher chance of being adopted. Our goal in this work is, therefore, to collect a rich context-annotated dataset and use a rule mining approach that derives insights into causal links between such context and appliance usage.

Energy annotations: Collecting annotated energy traces and providing energy usage feedback are active research areas in the human-computer interaction community. Eco-feedback systems use visualization techniques to engage the user in the conservation process [27], [28], [29]. A primary goal is to make the user cognizant of potential energy bottlenecks. Costanza et al. [1] propose a time series-based web interface called *FigureEnergy* where users can label their activities. They also provide a visualization interface where users can understand the impact of performing certain activities on the total home energy consumption. Unfortunately, the annotations are based on the user remembering the activities he performed during the day. Our system uses monitoring and in situ notification to prompt the user to log activities. Additionally, our primary goal is to use the activities and other context to automatically mine associations. These rules can then augment a visualization system like *FigureEnergy* to show users where and why they might have used energy.

Rule mining: Rule mining is a common data mining technique that is used to find associations between two or more random variables. It has been used extensively in relational databases [30], semantic web applications [31], and network trace mining [12]. There are also efforts to integrate the rule mining approach with classification algorithms [32], and design privacy preserving rule mining techniques [33]. Our work applies the rule mining approach to the area of home energy management, similar to the preliminary work of Ong, Berges, and Noh [34]. We, however, derive significant underlying associations between several contextual factors including hour of the day, day of the week, and user activities with normal usage of appliances and show a broad set of insights based on our context-rich data set. Finally, similar to the network mining systems [12], we illustrate that JMeasure [6] is a good measure of significance for the derived rules in the home energy management domain.

Demand-response in homes: A primary goal of the Green Homes project is to devise techniques that balance energy supply with demand, especially in renewable energy driven homes. To this end, deriving associations between user context and appliance energy consumption can provide better ways to understand energy demand and provide timely energy saving recommendations to home users. Hence, our proposed approach is complementary to several demand-response systems in homes [35], [36] and can be used to improve their performance. These include systems that flatten peak energy consumption [37], methods to predict energy generation and consumption [38], balance energy demand with supply [7], and minimize whole building energy consumption.

VI. Future Work and Conclusion

This paper presents a novel approach for understanding the causal factors that influence energy consumption in the home. By applying a rule mining algorithm to a context-rich data set collected from six homes across the United States we are able to extract a broad set of associations between device usage and hour of the day, day of the week, activities performed by the user, and use of other devices in the home. Our analysis confirms that while time-based associations are observed most frequently, associations between devices are common and often stronger than time-based associations. Moreover, associations with activities performed by the user are most common for devices that have a sporadic usage pattern. The associations derived by our approach can aid in energy disaggregation, end-user energy saving systems, and even aging-in-place applications.

We are currently pursuing several areas related to this work. First, we are exploring whether useful insights may be derived by applying alternate machine learning algorithms. In particular, we have considered whether it is feasible to predict hours when appliances will be used based on observation of past use. Early results indicate that complex machine learning algorithms such as Support Vector Machines suffer from poor accuracy and simpler algorithms such as KNN show similarly poor results—in the range of less than 30% for some appliances such as refrigerators. These early results suggest that using a simple metric such as JMeasure, which captures relationships even in case the probability of both observations is low, is a superior approach.

We are also planning to collect data from additional homes, as well as expand our data collection in each home. We anticipate that in homes where a large number of devices are measured we will see similar time-based and annotation-based relationships, however we may need to refine our algorithm to extract appropriate device-to-device relationships. We are exploring the impact of varying the JMeasure threshold used in identifying device-to-device associations, and we are also exploring how to identify device groups, for example by location in the home, that may assist with narrowing the scope of which devices may show associations.

References

[1] Enrico Costanza, Sarvapali D. Ramchurn, and Nicholas R. Jennings. Understanding domestic energy consumption through interactive visualisation: a field study. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, pages 216–225, New York, NY, USA, 2012. ACM.

[2] Gwendolyn Brandon and Alan Lewis. Reducing Household Energy Consumption: A Qualitative and Quantitative Field Study. *Journal of Environmental Psychology*, 19(1):75 – 85, 1999.

[3] Hyungsul Kim, Manish Marwah, Martin F. Arlitt, Geoff Lyon, and Jiawei Han. Unsupervised disaggregation of low frequency power measurements. In *SDM*, pages 747–758. SIAM / Omnipress, 2011.

[4] Sweta Sneha and Upkar Varshney. Enabling ubiquitous patient monitoring: Model, decision protocols, opportunities and challenges. *Decision Support Systems*, 46(3):606 – 619, 2009. Wireless in the Healthcare.

[5] Diane J. Cook and Sajal K. Das. How smart are our environments? an updated look at the state of the art. *Pervasive and Mobile Computing*, 3(2):53 – 73, 2007. Design and Use of Smart Environments.

[6] Gregory Piateski and William Frawley. *Knowledge Discovery in Databases*. MIT Press, Cambridge, MA, USA, 1991.

[7] Nilanjan Banerjee, Sami Rollins, and Kevin Moran. Automating Energy Management in Green Homes. In *ACM Sigcomm Workshop on Home Networks*, August 2011.

[8] David Lachut, Simon Piel, Lazeeb Choudhury, Yucheng Xiong, Sami Rollins, Kevin Moran, and Nilanjan Banerjee. Minimizing intrusiveness in home energy measurement. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '12, pages 56–63, New York, NY, USA, 2012. ACM.

[9] http://micasaverde.com/vera.php. Micasaverde vera 2.

[10] Sami Rollins, Nilanjan Banerjee, Lazeeb Choudhury, and David Lachut. Poster Abstract: An In Situ System for Annotation of Home Energy Data. In *ACM Workshop On Embedded Systems For Energy-Efficiency In Buildings*, November 2013.

[11] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc. of 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231, 1996.

[12] Srikanth Kandula, Ranveer Chandra, and Dina Katabi. What's going on?: learning communication rules in edge networks. In *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, SIGCOMM '08, pages 87–98, New York, NY, USA, 2008. ACM.

[13] Julien Blanchard, Fabrice Guillet, Regis Gras, and Henri Briand. Using information-theoretic measures to assess association rule interestingness. In *Proceedings of the Fifth IEEE International Conference on Data Mining*, ICDM '05, pages 66–73, Washington, DC, USA, 2005. IEEE Computer Society.

[14] J. Froehlich, E. Larson, S. Gupta, G. Cohn, M. Reynolds, and S. Patel. Disaggregated end-use energy sensing for the smart grid. *Pervasive Computing, IEEE*, 10(1):28 –39, jan.-march 2011.

[15] Georgina Corte Franco, Floriane Gallay, Marc Berenguer, Christine Mourrain, and Pascal Couturier. Non-invasive monitoring of the activities of daily living of elderly people at home a pilot study of the usage of domestic appliances. *Journal of Telemedicine and Telecare*, 14(5):231–235, 2008.

[16] Corinna Fischer. Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency*, 1:79–104, 2008.

[17] Geraldine Fitzpatrick and Greg Smith. Technology-enabled feedback on domestic energy consumption: Articulating a set of design concerns. *IEEE Pervasive Computing*, 8(1):37–44, January 2009.

[18] Tom Hargreaves, Michael Nye, and Jacquelin Burgess. Making energy visible: A qualitative field study of how householders interact with feedback from smart energy monitors. *Energy Policy*, 38(10):6111 – 6119, 2010.

[19] James Pierce, Chloe Fan, Derek Lomas, Gabriela Marcu, and Eric Paulos. Some consideration on the (in)effectiveness of residential energy feedback systems. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, DIS '10, pages 244–247, New York, NY, USA, 2010. ACM.

[20] Yolande A.A. Strengers. Designing eco-feedback systems for everyday life. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2135–2144, New York, NY, USA, 2011. ACM.

[21] Thomas Schmid, David Culler, and Prabal Dutta. Meter any wire, anywhere by virtualizing the voltage channel. In *BuildSys*, 2010.

[22] Jay Taneja, David Culler, and Prabal Dutta. Towards cooperative grids: Sensor/actuator networks for promoting renewables. In *SmarGridComm*, 2010.

[23] Vijay Srinivasan, John Stankovic, and Kamin Whitehouse. WaterSense: Water Flow Disaggregation using Motion Sensors. In *BuildSys*, 2011.

[24] Shwetak N. Patel, Sidhant Gupta, and Matthew S. Reynolds. The design and evaluation of an end-user-deployable, whole house, contactless power consumption sensor. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, pages 2471–2480, New York, NY, USA, 2010. ACM.

[25] Timothy W. Hnat, Vijay Srinivasan, Jiakang Lu, Tamim I. Sookoor, Raymond Dawson, John Stankovic, and Kamin Whitehouse. The hitchhiker's guide to successful residential sensing deployments. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, SenSys '11, pages 232–245, New York, NY, USA, 2011. ACM.

[26] S. Szewcyzk, K. Dwan, B. Minor, B. Swedlove, and D. Cook. Annotating smart environment sensor data for activity learning. *Technol. Health Care*, 17(3):161–169, August 2009.

[27] Jon Froehlich, Leah Findlater, Marilyn Ostergren, Solai Ramanathan, Josh Peterson, Inness Wragg, Eric Larson, Fabia Fu, Mazhengmin Bai, Shwetak Patel, and James A. Landay. The design and evaluation of prototype eco-feedback displays for fixture-level water usage data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 2367–2376, New York, NY, USA, 2012. ACM.

[28] Jon Froehlich, Leah Findlater, and James Landay. The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1999–2008, New York, NY, USA, 2010. ACM.

[29] Yolande A.A. Strengers. Designing eco-feedback systems for everyday life. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2135–2144, New York, NY, USA, 2011. ACM.

[30] Sunita Sarawagi, Shiby Thomas, and Rakesh Agrawal. Integrating association rule mining with relational database systems: alternatives and implications. In *Proceedings of the 1998 ACM SIGMOD international conference on Management of data*, SIGMOD '98, pages 343–354, New York, NY, USA, 1998. ACM.

[31] Xin Jin, Yanzan Zhou, and Bamshad Mobasher. Web usage mining based on probabilistic latent semantic analysis. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '04, pages 197–205, New York, NY, USA, 2004. ACM.

[32] Albrecht Zimmermann and Luc Raedt. Corclass: Correlated association rule mining for classification. In Einoshin Suzuki and Setsuo Arikawa, editors, *Discovery Science*, volume 3245 of *Lecture Notes in Computer Science*, pages 60–72. Springer Berlin Heidelberg, 2004.

[33] Y. Saygin, V.S. Verykios, and A.K. Elmagarmid. Privacy preserving association rule mining. In *Research Issues in Data Engineering: Engineering E-Commerce/E-Business Systems, 2002. RIDE-2EC 2002. Proceedings. Twelfth International Workshop on*, pages 151–158.

[34] Leneve Ong, Mario Bergés, and Hae Young Noh. Exploring sequential and association rule mining for pattern-based energy demand characterization. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 25:1–25:2, New York, NY, USA, 2013. ACM.

[35] Tanuja Bapat, Neha Sengupta, Sunil K. Ghai, Vijay Arya, Yedendra B. Shrinivasan, and Deva Seetharam. User-sensitive scheduling of home appliances. In *GreenNets*, 2011.

[36] Thomas Weng, Bharathan Balaji, Seemanta Dutta, Rajesh Gupta, and Yuvraj Agarwal. Managing plug-loads for demand response within buildings. In *In Proceedings of the ACM Workshop on Embedded Sensing Systems For Energy-Efficiency*, Seattle, October 2011.

[37] Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht. SmartCap: Flattening Peak Electricity Demand in Smart Homes. In *IEEE International Conference on Pervasive Computing and Communications*, March, 2012.

[38] Ting Zhu, Aditya Mishra, David Irwin, Navin Sharma, Prashant Shenoy, and Don Towsley. The case for efficient renewable energy management for smart homes. In *BuildSys*, 2011.