Behavior-based Home Energy Prediction

Chao Chen and Diane J. Cook
School of Electrical Engineering and Computer Science
Washington State University
Pullman, USA
Email: {cchen, cook}@eecs.wsu.edu

Abstract—**In the effort to build a sustainable society, smart home research attention is being directed toward green technology and environmentally-friendly building designs. In this paper, we analyze the distribution of home energy consumption, and then present both linear and non-linear regression learning models for predicting energy usage given known human behavior and time-scale features. To guarantee the validity of our methods, two real-world data sets collected over three months are applied into training the models. Based upon our learning models, a web-based end-user system is developed for providing users feedback about behavior-based energy usage to promote energy efficiency and sustainability through behavior changes.**

***Keywords*-smart environments; behavior; energy prediction;**

I. Introduction

In 2010, the United States consumed 98,003 Quadrillion Btu of power energy. This is a 200 percent increase from 1949 [1]. The growth of energy usage is not entirely due to manufacturing plants and automobiles, as is often assumed. In fact, worldwide residential sector is responsible for 16-50% of energy consumption consumed by all sectors [2]. Society is becoming increasingly aware of the impact residential lifestyle choices make on energy usage and the environment. As a result, there is an urgent need to develop technologies that examine energy usage in homes to encourage energy efficient behaviors.

Although households and buildings are responsible for over 40 percent of energy usage in most countries [3], many residents still receive little or no detailed feedback about their personal energy usage. A power utility bill traditionally provides information about a month's total energy consumption and a total price to be paid, leaving homeowners to guess what might explain a higher or lower than usual bill. Earlier studies have shown that home residents reduce energy expenditure by 5-15% just as a response to acquiring and viewing raw energy usage [4]. Residential behavior, which varies widely, can influence energy usage significantly in a given house [5]. Clearly, the typical utility bill provides no information about the relationship between residential behavior and corresponding energy usage. Occupants behavior with regards to use and settings of appliances is difficult to be captured in normal homes. Since behavior-based energy information is capable of encouraging individuals to modify habits in ways that would be beneficial for both the household and community, it would be desirable to develop technologies that could extract the information from the home and communicate it to residents.

A smart home environment is one that acquires and applies knowledge about its residents and their physical surroundings in order to improve their experience in that setting. Such home environments, equipped with sensors for detecting motion, light level, temperature, and energy and water consumption, are ideal testbeds for investigating techniques of inducing behavior changes to reduce energy usage.

The long-term vision for this project is to enhance understanding of human resource consumption and to provide tools that promote resource efficiency in smart homes. We hypothesize that providing users with behavior-based knowledge of energy consumption, suggestions for energy reduction will result in more substantial decreases in overall consumption. This view is supported by an increasing body of work that links awareness of energy consumption and its impact to behavioral change [6]. Additionally, we hypothesize that energy consumption is correlated with human activities and can therefore be predicted using smart home-based sensor information. These hypotheses are validated by implementing algorithms to perform these steps and evaluating the algorithms using two real data sets collected in smart apartments. Finally, a discussion of how the results of this work can be used to give smart home residents feedback on their energy consumption is included.

II. Related Work

Using sensor technology combined with data mining and machine learning, many researchers are now working on smart environments which can discover and recognize residents' activities and respond to their needs in a context-aware way [7]. As household consumption of electricity has been growing dramatically, the need to develop technologies that improve energy efficiency and monitor energy usage in a household is emerging as a critical research area.

Technologies to address this need are beginning to emerge. Non-intrusive appliance load monitoring [8] have been designed to detect the turning on and off of individual appliances in an electrical circuit. Several academic studies focused on this topic to estimate residential energy levels based on appliance usage [9][10]. With respect to energy conservation, some industrial products focus on providing energy information services and saving tips to residents. Google PowerMeter [11] is a free energy monitoring tool for saving energy by providing energy information via smart meters. The companies, such as Microsoft Hohm [12] and Opower [13], apply statistical methods and data mining algorithms to analyze a home's raw utility data and give customers usable energy saving tips. However, these works are orthogonal to this paper, in which we provide users with feedback that is actually related to resident behavior in the home. Likewise, several studies exist that predict building energy consumption at a highly aggregated level for a large collection of buildings [14], but these studies also differ from our work, as we consider human behaviors in an individual building as primary features for predicting energy usage.

In comparison to other past works, several different contributions are offered by this paper. The data sets we use are capable of representing realistic resident patterns over the monitoring period, since our monitoring did not affect the residents daily routine. We also analyze the specific distribution of power usage in home environments while residents performed these daily routines; we use machine learning techniques to build predictive models, and finally we present a web-based end-user interface for providing users with the feedback of their energy consumption together with a visualization of their behavior patterns. The result of this work can be used to improve energy efficiency in homes by behavior changes.

III. CASAS Smart Home Environment

The CASAS smart home environment testbeds used to analyze energy usage are apartments located on the Washington State University campus. As shown in Figure 1, the Kyoto smart home apartment testbed consists of three bedrooms, one bathroom, a kitchen, and a living/dining room. Only the first floor of the Tulum apartment is monitored, which consists of a living room, dining room, kitchen, and bathroom. To track peoples location, we use motion sensors placed on the ceilings. The circles in the figure stand for the positions of motion sensors. They facilitate the tracking of residents as they move through the space. In addition, the testbeds also include temperature sensors as well as custom-built analog sensors to provide temperature readings and hot water, cold water and stove burner use. A power meter records the amount of instantaneous power usage and the total amount of power consumption. An in-house sensor network captures all sensor events and stores them in a SQL database for later analysis.

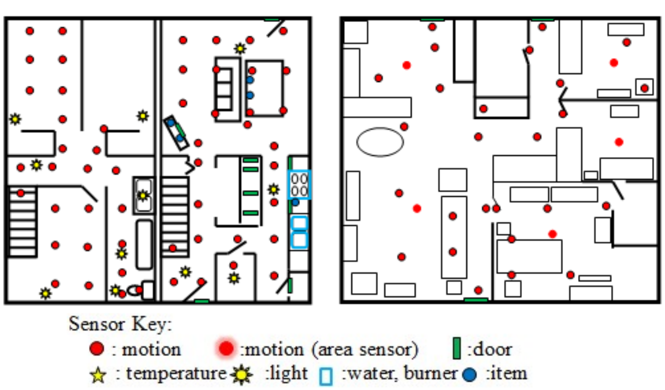

Fig. 1. CASAS smart apartment testbeds: Kyoto (left) and Tulum (right).

The sensor data gathered for our SQL database is expressed by several features, summarized in Table I. These four fields (Date, Time, Sensor ID and Message) are generated by the CASAS data collection system automatically.

TABLE I
Raw data from sensors. Sensor IDs beginning with "M" refer to motion sensors, "T" refers to temperature sensors, and "P" refers to power readings.

Date	Time	Senosr ID	Message
2009-02-06	17:17:36	M45	ON
2009-02-06	17:17:40	M45	OFF
2009-02-06	11:13:26	T004	21.5°C
2009-02-06	11:18:37	P001	930W

To provide real training data, data was collected from each apartment while two sets of two volunteers were living in each space. The training data was gathered during two months in Kyoto and three months in Tulum respectively. More than 100,000 sensor events were generated in each site during this time. All of the experimental data [1] are produced by the day-to-day lives of these residents, which ensure that the results of this analysis are applicable to other real-world environments.

IV. Home Energy Data Distribution

A preliminary analysis of energy data is used to explore the distribution of energy usage, which our prediction models are developed based upon in the following section. Figure 2 shows density functions of energy consumption per hour for our two smart environments. The figure illustrates that the energy usage per hour in smart environments roughly follows a log-normal distribution skewed to the left. This means that energy consumption in the house varies multiplicatively between different hours: most of the time, energy usage stays at a relatively low level, but the hours in the high percent of energy usage consume several times as much energy as the hours in the low percent. The result that hourly energy usage in an individual house approximately follows a log-normal distribution is interesting. In fact, there are various branches of phenomena in economics and science that roughly follow log-normal distributions [15], including some features that may influence energy consumption, such as income and age of marriage. In particular, energy consumption of different buildings is also verified to roughly follow a log-normal distribution [14]. Since there may be some noise or abnormalities existing in raw energy usage, we only consider log-scale energy usage falling within three standard deviations of the mean, which means that 99.7% of the instances can be applied toward future energy prediction.

To predict energy consumption, we assume that energy consumption has a measurable relationship with residents routine activities, which can be detected using motion sensors installed in the celling. These activities are either directly or indirectly associated with a number of electrical appliances and thus have a unique pattern of power consumption. To illustrate this relationship between behavioral patterns and

[1] http://ailab.wsu.edu/casas/datasets/index.html

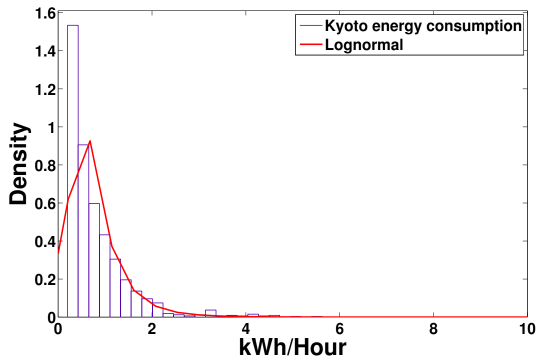

(a) Kyoto

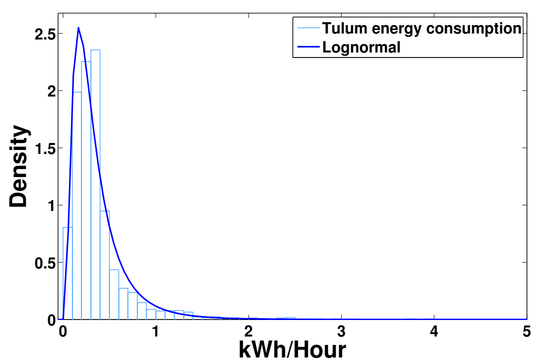

(b) Tulum

Fig. 2. Lognormal distributions of energy consumption per hour for both Kyoto and Tulum smart environments.

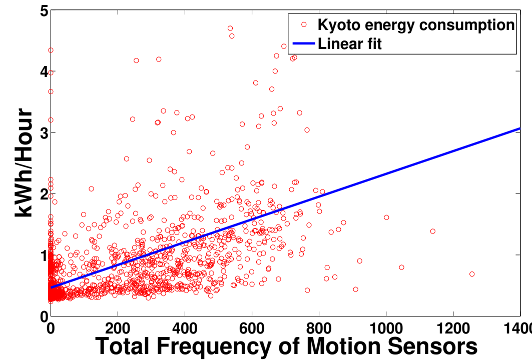

(a) Kyoto

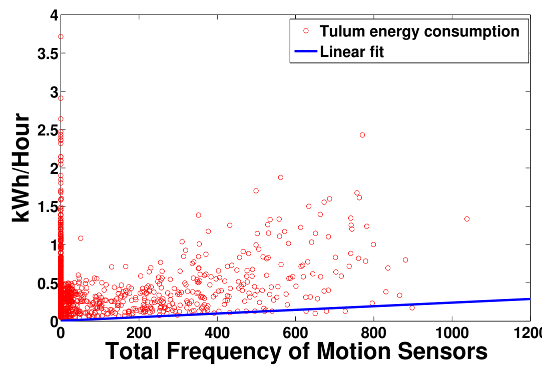

(b) Tulum

Fig. 3. Plots of total frequency of motion sensors versus energy consumption per hour (Kyoto (a), correlation coefficient: 0.31; Tulum (b), correlation coefficient: 0.24). The line shows the least-squares linear fit.

energy consumption, Figure 3 plots the total frequency of motion sensors against energy consumption per hour. Although there is, not unexpectedly, plenty of noise, there is also an obviously linear relationship between motion sensor frequency and power consumption, indicating that the levels of energy consumption growth in direct proportion to users activities at home. We argue that energy consumption can be measured more accurately given greater sensor density and diversity. The next section will be further to exploit different known features used to derive predictive models of energy consumption.

V. Data Features and Feature Selection

Before illustrating the predictive techniques we use for our data sets, we summarize the specific features we extract from raw data as well as the feature selection techniques we use to identify the maximum relevance features for this paper. The features used to predict energy are shown in Table II.

Since a large number of features are generated during feature extraction, it is necessary to determine which features are the most important factors to determine energy prediction. To identify these features, we employ a heuristic minimum-Redundancy-Maximum-Relevance (mRMR) [16] selection framework, which selects features mutually far away from each other that still maintain high relevance to the final target. In mRMR, Max-Relevance is used to determine the mean value of all mutual information values between an individual feature x_i and class c:

$$I(x_i, y) = \iint p(x_i, y) log \frac{p(x_i, y)}{p(x_i)p(y)} \, \mathrm{d}x \, \mathrm{d}y \tag{1}$$

$$maxD(S, c), \;\; D = \frac{1}{|S|} \sum_{x_i \in S} I(x_i, c) \tag{2}$$

Since the Max-Relevance features may have a very high possibility of redundancy, a minimal redundancy condition can be added to select mutually exclusive features:

$$minR(S), \;\; R = \frac{1}{|S|^2} \sum_{x_i, x_j \in S} I(x_i, x_j) \tag{3}$$

Thus, the operator $\phi(D, R)$ is the mRMR criterion combining the above two constraints:

$$max\phi(D, R), \;\; \phi = D - R \tag{4}$$

By applying the mRMR, the most relevant and least-redundant features among the candidate features are selected. Table III lists the top six important features from Kyoto and Tulum, respectively, which would be expected to have a large impact on energy consumption. Inspecting the results of feature selection, some specific motion sensor features and time-related features have been selected.

TABLE II
DATA FEATURES FOR REGRESSION MODELS.

Feature Name	Description
Length of day	This feature represents the time length since midnight when one instance happens (in seconds).
Day of week	This feature shows the current day of week (Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday).
Weekday/Weekend	This feature is a binary variable to determine whether the current day is a weekday or weekend.
Time of day	This feature represents different time slots (morning, noon, afternoon, evening, night, and late night).
Times of individual sensors	This feature represents times of different activated motion sensors that were during the time window.
Number of kinds of motion sensors	This feature represents the number of motion sensors that were triggered in various rooms.
Total number of times of motion sensors	This feature represents the total number of motion sensor events that were triggered during the time window.

TABLE III
SELECTED FEATURES USING MRMR

Max-Relevance, and Min-Redundancy Features	
Kyoto	**Tulum**
M17 (Kitchen sensor)	M15 (Bathroom sensor)
M38 (Bathroom sensor)	Day of week
M8 (Living room sensor)	M30 (Bedroom2 sensor)
Length of day	M13 (Bathroom sensor)
M32 (Bedroom2 sensor)	Time of day
M18 (Kitchen sensor)	M22 (Bedroom1 sensor)

In Kyoto, one each sensor from Kitchen/Bathroom is selected and two Bathroom sensors also are chosen in Tulum. It makes sense that residents staying in the Kitchen are very likely to do some cooking, while they may consume a plenty of hot water to have a shower in the bathroom. All of these activities are very likely to generate specific energy patterns. The sensors in the Living/Bedroom are also selected by feature selection. The participants movement patterns throughout the space are either directly or indirectly associated with a number of electrical appliances and thus have a unique pattern of power consumption. For example, when the residents perform the activity in the bedroom, they may turn on a light and operate the computer. It should be mentioned that some time-based features (Length of day, Day of week, and Time of day) are selected, indicating that the residents may generate various energy usage patterns under different time periods of a single day or days during a week. Thus, those features can also be acted as important factors for energy prediction.

VI. PREDICTIVE MODELS

In this section, we present machine learning techniques for predicting energy consumption based upon sensor features. It should be noted that we are not planning to predict energy consumption very precisely, since energy usage mainly depends on residential preferences on using electrical appliances in the houses, which is hard to be captured by our current sensors. The purpose of this section focuses on providing users basic information about the distribution of their energy consumption based upon their personalized behaviors. In this section, we focus on using real-valued regression for predicting just the total energy consumption. Generally, the learning models have the form as follows:

$$y = f(x) + \varepsilon \tag{5}$$

Here y denotes the predicted energy consumption, x denotes a vector of known features described above, and ε denotes a zero-mean error term. The goal of our task is to find a function $f(x)$ that has at most deviation from the actually energy usage for all the training data. Two well-known machine learning models were leveraged into this work: a linear regression model and a support vector machine. Additionally, since there may be some concerns of known features that may potentially follow non-linear function, a non-linear kernel function is considered to be applied into support vector machine regression models discussed later.

A. Linear Regression Model

A linear regression model [17] is applied to model the linear relationship between one or more input features $x \in R^n$ and target variable y. The linear regression model

$$y = \beta^T X + \varepsilon \tag{6}$$

for parameters $\beta \in R^n$, and the error term is modeled by a Gaussian distribution. Given a data set $\{x_i, y_i\}_{i=1}^n$, the maximum likelihood estimates of β can be calculated by the least squares method. The relevant parameters of the models are defined as:

$$y = \begin{Bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{Bmatrix}, x = \begin{Bmatrix} x_1^T \\ x_2^T \\ \vdots \\ x_n^T \end{Bmatrix}, \beta = \begin{Bmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_n \end{Bmatrix}, \varepsilon = \begin{Bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{Bmatrix}, \tag{7}$$

The least squares method finds its optimum when the sum of squared residuals S is a minimum value.

$$S = \sum_{i=1}^{n} \varepsilon_i^2 = \sum_{i=1}^{n} (y_i - \beta^T x_i)^2 \tag{8}$$

For linear regression models, the properties of the relevant estimators are easier to determine comparing to nonlinear models. However, there may exist a nonlinear relationship between known features and energy consumption, a support vector machine with a nonlinear kernel will be considered in the next part for exploring such a nonlinear relationship.

B. Support Vector Machine Regression Model

A support vector machine (SVM) model [18] is an optimal algorithm, which maximizes the margin between training examples and class boundary. It can be applied into both classification and regression problems. In the regression case, SVM estimates the model parameters by minimizing the risk, which is measured using Vapniks ε-insensitive loss function. Given a training data set, a SVM regression (SVR) function can be described as follows:

$$y = \omega\phi(X) + b \tag{9}$$

Here ω and b represent the estimator of SVM model, $\phi(X)$ is a map from the original data space of X to a high-dimensional feature space. To obtain the parameters of ω and b, the SVR function can be solved by the following constrained optimization problem by introducing the positive slack variable ξ_i and ξ_i^* as follows:

$$\begin{aligned} &minimize \quad \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{l}(\xi_i + \xi_i^*) \\ &subject\ to \quad \begin{cases} y_i - w\phi(X) - b \leq \varepsilon + \xi_i \\ w\phi(X) + b - y_i \leq \varepsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0 \end{cases} \end{aligned} \tag{10}$$

Finally, by applying Lagrange multipliers and the optimal constraints, the SVR problem can be solved by

$$minimize \quad \begin{cases} -\frac{1}{2}\sum_{i,j=1}^{l}(a_i - a_i^*)(a_j - a_j^*)K(x_i, x_j) \\ -\varepsilon\sum_{i=1}^{l}(a_i + a_i^*) + \sum_{i=1}^{l} y_i(a_i - a_i^*) \end{cases} \tag{11}$$

$$subject\ to \quad \sum_{i=1}^{l}(a_i - a_i^*) = 0 \ and \ a_i, a_i^* \in [0, C]$$

The target function can be computed by

$$f(x) = \sum_{i=1}^{l}(a_i - a_i^*)K(x_i, x_j) + b \tag{12}$$

Where $K(x_i, x_j) = (x_i)*(y_i)$ is defined as the kernel function. The main purpose of kernel functions is to deal with non-linear feature spaces without calculating the map $\phi(x)$ explicitly. For our task, we use a linear kernel $K(x, y) = x*y$ and a non-linear kernel $K(x, y) = (x * y + 1)^2$, which are applied to explore a linear and non-linear relationship between known features and energy consumption respectively. The Sequential Minimal Optimization (SMO) [19] is reported to be an effective method for improving scaling of the training set and computation time of the SVMs. In our study, the SMO is applied to training the SVMs for solving energy prediction problem.

VII. Experimental Results and Analysis

Two series of experiments were performed for energy prediction in smart environments. The first experiment uses the sensor data collected during two months in the Kyoto testbed. In the second experiment, we collected data for three months from the Tulum testbed. We evaluated the performance of the algorithms based on a 10-fold cross validation: the data instances was randomly divided into 10 approximately equal-sized groups, then trained the algorithms over nine of the total groups and tested on the remaining one; we repeated the procedure 10 times for each of the 10 groups, and the average error and correlation coefficient over all the instances were reported. Tables IV and V show the performance of alternative algorithms on two data sets. The algorithms were evaluated by two metrics: (1) correlation coefficient (R), which measures how a regression model fits the data sets; (2) root mean squared error (RMSE) on the energy consumption. Five different time windows were selected for testing the performance on various time scales (1-hour, 4-hour, 6-hour, 8-hour, 1-day).

TABLE IV
Cross Validation Performance of Different Algorithms and Time Windows for the Kyoto Data Set. (Item in Bold Indicates the Best Performing Method)

Time	Linear Regression		SVM(Linear)		SVM(NonLinear)	
	R	RMSE	R	RMSE	R	RMSE
1-hour	0.746	0.410	0.749	0.418	0.644	0.571
4-hour	0.784	1.010	0.792	1.020	0.649	1.540
6-hour	0.779	1.832	**0.830**	1.563	0.732	2.133
8-hour	0.679	1.850	0.753	1.619	0.594	2.227
1-day	0.462	4.620	0.695	4.039	0.719	3.859

TABLE V
Cross Validation Performance of Different Algorithms and Time Windows for the Tulum Data Set. (Item in Bold Indicates the Best Performing Method)

Time	Linear Regression		SVM(Linear)		SVM(NonLinear)	
	R	RMSE	R	RMSE	R	RMSE
1-hour	0.418	0.312	0.386	0.323	0.289	0.193
4-hour	0.634	0.637	0.620	0.645	0.431	1.010
6-hour	0.541	0.850	0.510	0.873	0.253	1.527
8-hour	**0.682**	0.942	0.67	0.944	0.360	1.747
1-day	0.381	2.501	0.159	4.214	0.132	28.401

In the Kyoto data set, as seen in Table IV , the SVM with a linear kernel obtains the best overall performance both on correlation coefficient and RMSE. The simple linear regression model is only marginally worse than the linear SVM, which performs much better than the SVM with non-linear method. For the Tulum training data in Table V, the performance between linear regression model and the linear SVM is very close, which is hard to determine which one is overall better than the other one. Both of these perform better than the non-linear SVM method. We argue that the linear models are preferable for energy prediction in smart environments. It also proves our potential assumptions that the intensity of residents behaviors has a strong linear relationship with energy consumption in houses.

In comparison of the effect on time windows, the Kyoto data instances with 6-hour time window can be best fitted by the predictive models. The regression models can fit best under the 8-hour window for the Tulum data set. It should be noticed that all three regression algorithms are not capable of fitting the Tulum data instances very well under the 1-day scale. The possible reason is that the residents in the Tulum have very similar behavior patterns in one-day scale, which is hard to be captured by the regression models. By comparing the overall performance between the Kyoto and the Tulum, energy consumption in the Kyoto environment can be predicted better by the models. We checked raw power values of the Kyoto and Tulum data sets respectively. We found out that the wave shape of power values in the Kyoto fluctuates dramatically. On the contrary, the Tulum energy usage keeps very smoothly and steadily under different intensity of residential behaviors. That is because there are more and heavier electrical appliances installed in the Kyoto, such as air conditioner, water heater, which may consume much more energy.

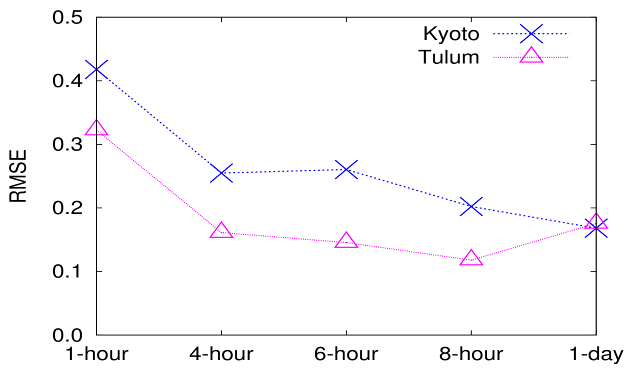

Fig. 4. Comparison of normalized RMSE on different time windows under the SVM with linear kernel.

Due to the different sizes of time windows, it is hard to compare the RMSE of the models under different time windows. Thus, we normalized the RMSE of all time windows to 1-hour scale RMSE as shown in Figure 4. From the figure, for both Kyoto and Tulum, the normalized RMSEs continue to decline dramatically with the size of time windows increases except the Tulum one-day window as discussion above. It is also interesting that the normalized RMSE in the Tulum is overall lower than the Kyoto normalized RMSE although the Kyoto data can be fitted much better by the models based on the standard of correlation coefficient. One possible explanation is that the Kyoto residents generate more energy at unusual times.

VIII. Web-based Energy Prediction Interface

The last component of our system is a behavior-based feedback tool to promote energy sustainability in everyday environments. We focus on a pervasive approach to promote sustainability behavior. Based on regression models described in the previous sections, an end-user application, called CASASviz, is developed as a web-based solution and can therefore run on a computer display or a mobile device. Figure 5(a) shows a line chart for letting residents view and compare predicted energy usage generated by our prediction models and true energy usage in real time. An accompanying activity map and two bar charts, as shown in Figures 5(b), 5(c) and 5(d), also allow users to quickly determine their energy consumption and the corresponding activities in the environment that impacts their

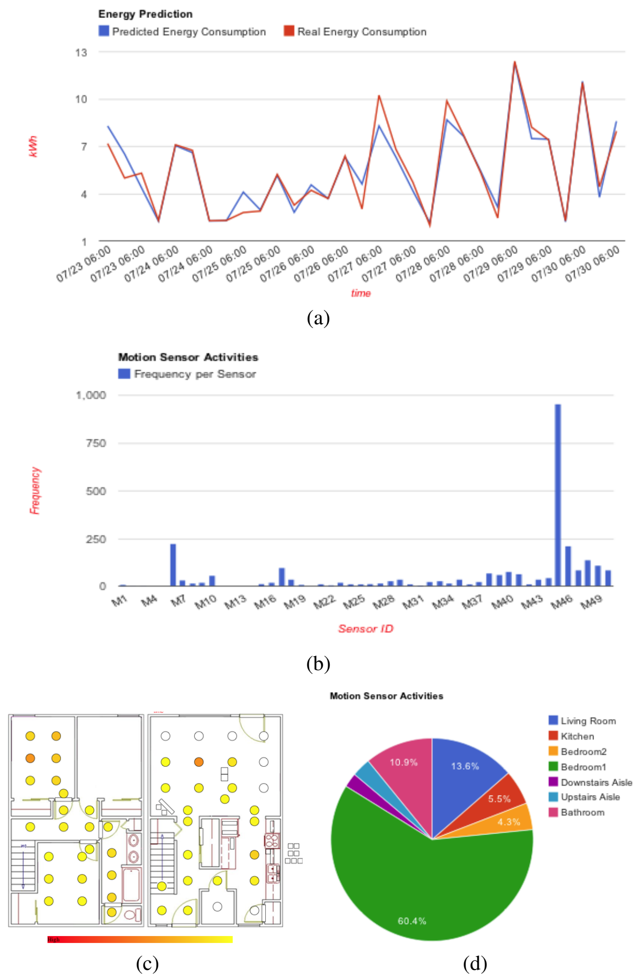

Fig. 5. Screenshots of our CASASviz sysem ((a) Line chart comparing predicted and real energy usage; (b) Frequency chart of motion sensors (c) Activity map of motion sensors; (d) Bar chart of sensor activities in each room).

energy utilization. A preliminary web demo of this tool is available at the website [2].

Although the interface can demonstrate residents' energy usage and their corresponding activities, there are still many open issues regarding this interface. The types of medium need to be more explored to communicate the users effectively, by conducting a survey regarding their requirements. Moreover, the information of the interface can be explained in detail. For example, the users can be noticed about the detailed usage of the appliances while certain activities are associated with higher energy consumption, and further provide the feasible tips to promote behavioral change for energy efficiency. All of these works are planning to be explored in our future work.

Since our approaches based upon user preferences are fully data-driven, our models are applied particularly to the residents living in our smart environments. Without loss of generality, our models can be also set up for other similar smart environments.

[2]http://demo.ailab.wsu.edu/pv/power.html

IX. Conclusion

In this paper, we consider the influence of in-home behaviors on energy usage. In particular, we analyze the distributions of energy consumption from two smart environments. We further identify the role of behaviors for energy usage by using machine learning methods to map activities performed in the environment with their corresponding energy usage. Finally, we proposed a web-based interface to use this information that allows individuals to perform their daily activities in a more energy-efficient manner.

In our ongoing work, we plan to install more sensitive power meters in order to capture more accurate changes in energy consumption. To extend our existing work, we will implement in a greater variety of households, which will allow us to determine whether energy prediction, energy usage trends exist and generalize across multiple settings. Finally, we will evaluate our web-based application to determine effects on sustainable behaviors and energy consumption.

References

[1] Multiple, "Annual energy review 2010," U.S. Energy Information Administration, Tech. Rep., 2011.

[2] T. Ichinose, K. Shimodozono, and K. Hanaki, "Impact of anthropogenic heat on urban climate in tokyo," *Atmospheric Environment*, vol. 33, no. 24-25, pp. 3897–3909, 1999.

[3] Multiple, "Energy efficiency in buildings," Tech. Rep., 2009. [Online]. Available: http://www.wbcsd.org

[4] S. Darby, "the effectiveness of feedback on energy consumption," *Change*, vol. 22nd, no. April, pp. 33–35, 2006.

[5] J. Seryak and K. Kissock, "Occupancy and behavioral affects on residential energy use," in *Proceedings of the Solar Conference*, 2003, pp. 717–722.

[6] S. Darby, "Smart metering: what potential for householder engagement?" *Building Research & Information*, vol. 38, no. 5, pp. 442–457, 2010.

[7] D. J. Cook and S. K. Das, *Smart environments: technologies, protocols, and applications*, ser. Wiley Series on Parallel and Distributed Computing. Wiley-Interscience, 2004.

[8] G. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870 –1891, Dec. 1992.

[9] M. Berges, E. Goldman, H. Matthews, and L. Soibelman, "Learning systems for electric consumption of buildings," in *ASCE International Workshop on Computing in Civil Engineering, Austin, Texas*, 2009.

[10] S. Gupta, M. S. Reynolds, and S. N. Patel, "Electrisense: single-point sensing using emi for electrical event detection and classification in the home," in *Proceedings of the 12th ACM international conference on Ubiquitous computing*, 2010, pp. 139–148.

[11] Google, "Google Power Meter," http://www.google.com/powermeter/about/, 2012.

[12] Microsoft, "Microsoft Hohm," http://www.microsoft-hohm.com/, 2012.

[13] Opower, "Opower," http://opower.com/, 2012.

[14] J. Z. Kolter and J. F. Jr., "A Large-Scale study on predicting and contextualizing building energy usage," in *the Conference on Artificial Intelligence (AAAI), Special Track on Computational Sustainability and AI*, 2011.

[15] E. Limpert, W. Stahel, and M. Abbt, "Log-normal distributions across the sciences: keys and clues," *BioScience*, vol. 51, no. 5, pp. 341–352, 2001.

[16] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 27, no. 8, 2005.

[17] S. Weisberg, *Applied linear regression*. Wiley, 2005, vol. 528.

[18] S. R. Gunn, "Support vector machines for classification and regression," *ISIS technical report*, vol. 14, 1998.

[19] J. Platt, "Sequential minimal optimization: A fast algorithm for training support vector machines," Microsfot Research, Tech. Rep., 1998.