

Predictability of Energy Use in Homes

David Lachut

Computer Science and Electrical Engineering
University of Maryland Baltimore County
Baltimore, Maryland
dlachut1@umbc.edu

Nilanjan Banerjee

Computer Science and Electrical Engineering
University of Maryland Baltimore County
Baltimore, Maryland
nilanb@umbc.edu

Sami Rollins

Computer Science
University of San Francisco
San Francisco, California
scrollins@cs.usfca.edu

Abstract—Predictability of home energy usage forms the basis of many home energy management and demand-response systems. While existing studies focus on designing more accurate prediction algorithms, a comprehensive energy management solution requires a broad understanding of prediction accuracy at different granularities, for example appliance and home, as well as different time horizons, for example an hour, day, or week into the future. In this paper, we undertake an analysis of predictability of power draw of appliances and whole-home energy consumption at four different time horizons: an hour, a quarter-day, a day, and a week in the future. Our analysis presents two research contributions. Our first contribution is a diverse dataset, GreenHomes, that includes appliance power draw and whole-home energy consumption data from seven homes across three states in the United States over a two-year period. Our second and primary contribution is a set of insights into the predictability of home energy usage. We show that simple statistic-based algorithms perform as well as sophisticated machine learning algorithms and time-series based predictors. These simple algorithms can considerably reduce the computational need for large-scale predictive analysis of home energy data. We also show that appliance-level power draw is more predictable than whole-home energy consumption at shorter time horizons while home-level energy consumption is more predictable at longer time horizons. Finally, we show that there is large variation in predictability across homes. This variation may be attributed to home type and points to the need for personalized energy management systems.

I. INTRODUCTION

The ability to predict future energy requirements is a critical component of a variety of applications that seek to conserve or improve management of energy resources. Utilities, for example, use forecasting of future demand to determine how to manage energy generation. With the recent widespread deployment of the smart grid, forecasting at the scale of individual homes and even appliances becomes necessary for enabling effective demand response systems and user-side energy management. A system that is able to accurately estimate energy requirements at the scale of individual appliances in the home, for example, can enable feedback and recommendations to support load shifting to reduce peak demand or help home owners better understand how they can modify usage to minimize electricity bills.

Building predictive models for time series data is a well studied area [1], however, most prior efforts in the domain of home energy have focused on designing better models for a small subset of scenarios [2], [3], [4], [5]. There is research, for example, on better prediction of appliance energy

usage at short time horizons like an hour into the future [6]. Similarly, some research efforts have focused on whole home energy consumption predictions at time frames of a day in the future [7], [8]. A comprehensive energy management solution, however, requires a broad understanding of prediction accuracy at different granularities, for example appliance and home, as well as different time horizons, for example an hour, day, or week into the future. A broad understanding of predictability may be used to inform the design of energy management solutions by allowing a designer to tune the system to tolerate prediction error, or to make smart decisions about integrating or excluding appliances known to have unpredictable usage patterns.

In this paper, we undertake a comprehensive study to answer the following question: *how predictable is home energy usage?* Our goal is not to design a new prediction algorithm, a topic that has been comprehensively studied in the literature [9], [8], but to instead derive insights into the predictability of home energy usage with application to demand-response systems in particular and home energy management in general. To answer the above question, we study the predictability of energy consumption at the scale of individual appliances, and single homes. We also compare the prediction accuracy at the time horizons of an hour into the future, a quarter day into the future, a day into the future, and a week into the future. For our analysis, we leverage three well-known machine learning predictors: k -nearest neighbor, a bayesian predictor, and support vector machine and a time-series based predictor (ARMA(1,1)). We compare these predictors with a simple algorithm that uses the energy consumed in the last time frame as the predicted energy consumed in the next time frame. Finally, we compare the performance of power feature average power consumed with time-of-day features such as hour of the day and day of the week. The quantitative analysis of the predictive models presents two novel research contributions.

Our first contribution is a diverse, longitudinal dataset on home energy consumption—the GreenHomes dataset. The dataset consists of power consumption data from appliances and whole home energy consumption data collected over a period of three years (on a rolling basis) from seven homes. The homes in the dataset are diverse: rentals and user owned homes; an off-grid home and a home partially powered by solar panels; and homes that are powered completely by the grid. The residents of the homes vary from graduate students (low income group) to professionals (middle and high income group), single-user and family homes, and include buildings in Arkansas, Maryland, and California. The dataset is the basis of our analysis and overcomes the shortcomings of publicly

available datasets that are either collected over a short period of time [10], sometimes do not include data for both appliances and the whole home [11], and are often collected in one geographic region [11], [10].

Our second and primary contribution is a set of insights that provide a comprehensive understanding of energy usage and can inform the design of demand-response and energy management systems. We show, for example, that simple algorithms that use basic power features like power consumed in the previous time frame have comparable performance to sophisticated machine learning algorithms and time series predictors. These simple algorithms can considerably reduce the computational need for large scale predictive analysis of home energy data. We also show that appliance-level energy usage is more predictable than whole home energy at shorter time horizons while home level energy consumption is more predictable at longer time horizons like a week. Finally, we show that there is a large variation in predictability across homes that can be attributed to home type.

II. GREENHOMES DATASET

The goal of our work is to gain a comprehensive understanding of the predictability of home energy usage. To achieve this goal, we impose the following three requirements for the dataset that serves as the basis for our analysis. First, it must provide data at different granularities including appliance power draw data as well as whole-home energy consumption data. Second, it must be longitudinal to enable derivation of predictability results that are significant, especially for longer time horizons like a week. Third, the dataset must be diverse, for example with respect to number of home occupants, in order to derive general insights into home energy usage. There exist several publicly available datasets on home energy consumption, however all suffer from drawbacks that make them unfit for our study. REDD is a popular dataset used in energy disaggregation research [10]. The dataset consists of highly granular appliance-level power draw data, however, the publicly available dataset released has only 30 days of usage data. While the Smart* dataset [11] is richer and provides circuit, home, and micro-grid data at a fine granularity, it consists of data from only three homes. Moreover, it does not include appliance power consumption data.

The GreenHomes dataset is a comprehensive home energy consumption dataset that our group has been collecting from seven deployments on a rolling basis over a period of 2 years. The characteristics of the collected data are illustrated in Table 1. Our deployments are diverse and include an off-grid home powered entirely by solar panels, a grid-tied house partially powered by solar, and five houses that are powered solely by the grid. Our deployments span houses in Arkansas, Maryland, and California, and hence include a large spectrum of weather zones. Four of our subjects own their homes while three houses are rental apartments. Our dataset includes data from more than 80 appliances and includes refrigerators, cooking appliances (microwave and ovens), entertainment appliances (televisions and gaming consoles), lighting, and washers/dryers. Two of our subjects (site 1 and site 4) moved during the span of our data collection and the 21 and 14 appliances noted in Table 1 include two distinct sets of appliances for

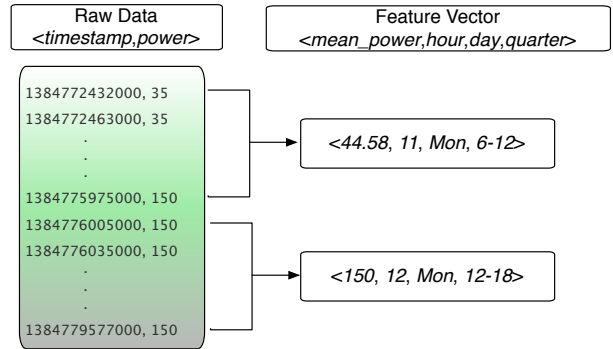


Fig. 2. Feature extraction for one hour scale.

the two locations. As shown in the table, our data set includes 5–20 months of data per site.

A. Data Collection System

We use off-the-shelf energy meters [12] to collect power draw information from several devices in each home. Our seven deployments measure a variety of appliances including televisions, lamps, microwaves, and computers. In some homes, we use a meter [13] to directly collect whole-home energy usage data.

In each home, we have deployed a low-power FitPC [14] that runs a client component built on the HomeOS [15] platform. The client periodically (once every 30 seconds) polls a dual-radio gateway [16] that retrieves instantaneous power draw readings from each individual meter. This information is then pushed to a centralized server where it is stored in a standard mysql database. Our current deployments focus on data collection, however our server does implement a secure, RESTful API that provides access to the raw data. We also provide users with an iPhone application that displays graphs of usage data as well as the ability to control individual appliances remotely [17]. We have prototypes of both a web application and an Android application. We have also tested a prototype that collects contextual information such as when the user is home and what the user is doing (activities like cooking, entertainment, chores, and work) when energy usage is high [18]. As future work, we plan to use this additional contextual information to inform our prediction algorithms.

B. Data Cleaning and Feature Extraction

The raw data collected by our system consist of power-draw values in Watts and corresponding timestamps. These data points must be processed into usable form for our prediction analysis. The first step is simply to remove impossible values: raw data points with wattages less than zero for instance are removed. Moreover, we have found that close to 22% of the time our meters malfunctioned or reported stale data. This is likely due to wireless transmission failures or power-outages. We do not use those time periods for our predictions.

For each appliance and home we monitor, our system collects instantaneous power draw every 30 seconds, however to make predictions about energy use we aggregate these

Site#	# of Appliances	Home data	# of months	# of occupants	Type of home
Site 0	8	Yes	12	2	rental, on-grid
Site 1*	21	Yes	5	1	rental, on-grid
Site 2	11		12	2	owned, on-grid
Site 3	9	Yes	18	1	owned, on-grid
Site 4*	14		20	2	rental, on-grid
Site 5	6		14	3	owned, grid-tied
Site 6	10	Yes	5	2	owned, off-grid

Fig. 1. Characteristics of the GreenHomes dataset. Our dataset includes data from a wide range of homes. Site 1 and Site 4 subjects moved during the course of the data collection.

(timestamp, power draw) values into a set of labeled data. For each home or appliance, we first compress the data into hour-, six-hour-, day-, or week-long units, depending on the scale at which we wish to make predictions, by calculating the mean power draw of the device in that time period. Figure 2 demonstrates this process on the hour scale. Note that the average power draw times the length of this time period is the energy use for that time period. Using the starting time for each of these time periods, we calculate time-based features: hour of the day, day of the week, and quarter of the day. Using the average power draws we calculated, we select power features such as the average power use in the immediately prior time period. The features we select and calculate are assembled into ordered sequences to be used as feature vectors in our prediction algorithms.

III. PREDICTION ALGORITHMS

For our analysis, we use three machine learning algorithms and one time series-based predictor: Naive Bayes Predictor [19], K-Nearest Neighbor predictor [20], Support Vector Machine predictor [21], and autoregressive and moving average predictor of order (1,1) [22]. The machine learning algorithms output a predicted class of power draw while the ARMA predictor outputs a continuous power draw value. To divide the power consumption values into classes, we cluster the average power draw for different time horizons into clusters using an unsupervised clustering algorithm similar to DBSCAN [23]. The predicted power consumption for the machine learning algorithms, therefore, is the power consumption value corresponding to the centroid of the predicted class. While the above set of algorithms used in our analysis is not exhaustive, they represent samples for a wide spectrum of learning and time-series based algorithms used popularly in demand-response systems. For instance, the Naive Bayes Predictor is a probabilistic predictor that assumes independence of features and SVM is a non-probabilistic algorithm that models dependencies between features. Auto-regressive and Moving Average prediction (ARMA) is a time-series-based predictor that builds a polynomial model for the system using a mix of regression and moving average models.

IV. INSIGHTS AND RESULTS

In this section, we discuss the insights inferred from our prediction analysis and their implications to home energy management and demand-response systems.

A. Experimental Setup

Algorithms: Our analysis considers four prediction algorithms: 5 Nearest-Neighbor (5NN), Naive Bayes predictor (Bayesian), Support Vector Machine (SVM), Auto-regressive and moving average model (ARMA(1,1)) outlined in §III, and a simple statistic-based algorithm (SIMPLE). SIMPLE gives us a baseline for comparison and uses the average power draw in the previous time cycle as the predicted average power draw in the next time cycle. It does not perform any training or feature aggregation, for instance, the average power draw predicted from 9AM-10AM is simply the average power draw between 8AM-9AM. The machine learning algorithms and the time series predictor train on two weeks of prior data and, unless otherwise noted, use a single feature: the power draw in the previous time cycle. Our analysis also considers the features time of the day, day of the week, and quarter (6-hour block) of the day.

Time Horizons: To understand how far in advance energy needs may be forecasted, we consider four time horizons: next hour, next quarter of the day (6 hours), next day (24 hours), and next week.

Prediction Accuracy: The primary metric we consider is Prediction Accuracy. In our analysis, a prediction is considered accurate if the difference between the predicted and actual power draw value is within 5% of the maximum power draw ever observed for the appliance. Consider a TV that draws 0 Watts when off, 30 Watts when in a low-power, suspended state, and 150 Watts when on. If the TV is in the suspended state drawing 30 Watts and the algorithm predicts 35 Watts the prediction will be considered accurate—it is within 7.5 Watts ($150 \times .05$) of the maximum power draw of the device.

B. Insights Derived

Insight 1: Simple statistics yield prediction accuracies comparable to sophisticated machine learning algorithms and time series predictors

Figure 3(a) shows the mean prediction accuracy for the power draw of 80 appliances in our seven deployment sites. The error bars are the standard deviations in accuracy for the appliances. Similarly, Figure 3(b) illustrates the prediction accuracy for the whole-home energy consumption averaged over four homes where whole-home energy consumption data was available (see Table 1). The figures show that the machine learning algorithms have comparable performance. Though SVM and Bayesian perform best for time horizons of an hour,

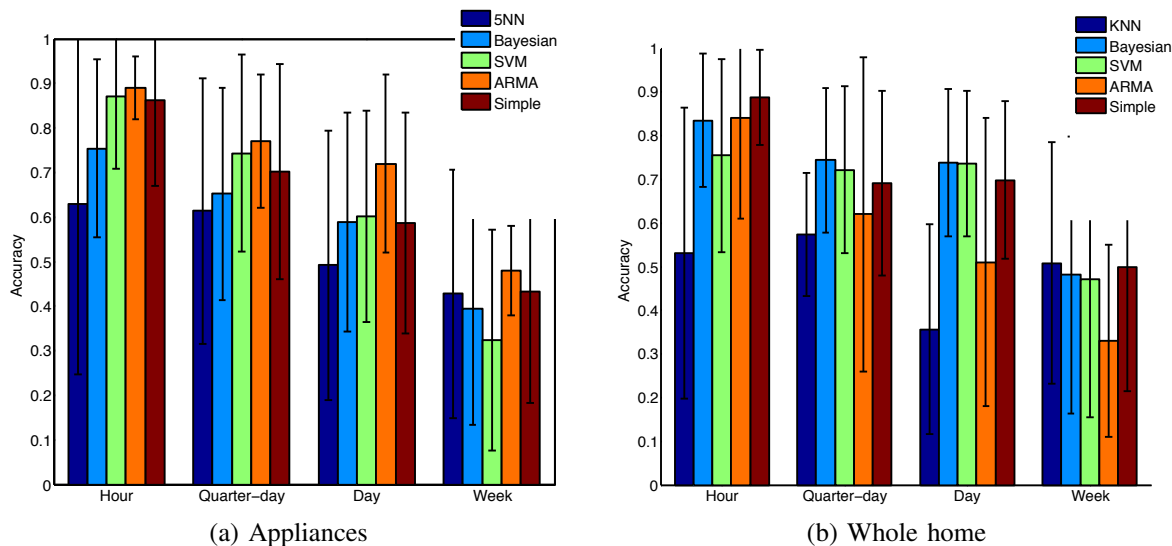


Fig. 3. (a) Prediction Accuracy for three machine learning algorithms, a time series predictor, and a simple statistic-based predictor for appliances for four time horizons. The graph shows the average prediction accuracy for 80 appliances in seven homes. The error bars depict the standard deviation across appliances. The simple statistic-based algorithm performs comparably to the machine learning algorithms and the time-series predictor. (b) Prediction Accuracy for three machine learning algorithms, a time series predictor, and a simple statistic-based predictor for whole-home energy consumption for the four time horizons. Again, the simple statistic-based algorithm performs comparably to the more sophisticated algorithms.

quarter-day, and a day, 5NN performs best at a longer prediction horizon of a week. The time series model, ARMA (1,1), performs best across all algorithms. The salient result from the figures, however, is that the performance of SIMPLE is similar to the three sophisticated machine learning algorithms and close to the performance of the time series predictor. For a prediction horizon of an hour, for instance, the SIMPLE algorithm has a average accuracy of 82%, compared to the 85% accuracy that the SVM algorithm provides. The performance is comparable for time horizons of a quarter-day, day, and a week, as well as for predicting the whole-home energy consumption.

To understand why a simple statistic-based algorithm performs as well as the machine learning-and time-series based predictors, we perform another set of experiments. These experiments evaluate the following two advantages that learning and time-series based predictors provide over simple statistics-based schemes like SIMPLE. First, they allow building predictive models based on a *set* of training data, for example, if the feature for training the predictive model is the average power draw in the last hour, the training set for the machine learning algorithm or the ARMA predictor uses a collection of samples of the form $\langle P_{-1}, P_0 \rangle$ where P_{-1} is the power draw in the last hour and P_0 is the power draw in the present hour. The model parameters are estimated as an aggregate fit for the set of training samples while minimizing an error metric like mean-squared error. Secondly, machine learning algorithms allow aggregation of disjoint features, for example, a learning algorithm can consider, in addition to power features, time-of-day features such as hour of the day and day of the week. A machine learning algorithm may then aggregate features using basis functions like radial basis [24]. We observe, however, that predictive models built for energy consumption data in our dataset do not benefit from these advantages.

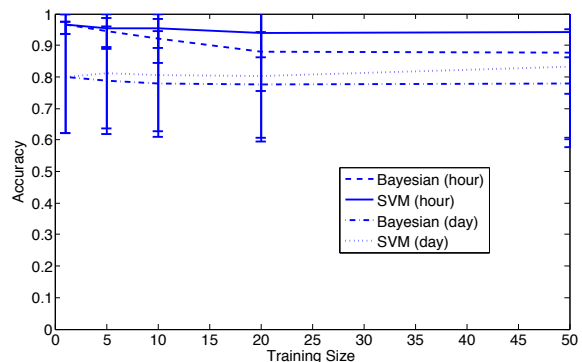


Fig. 4. The figure compares the average prediction accuracy for all appliances in Site 3 for different training set sizes. We used the SVM and Bayesian algorithms for this experiment. As the training size is increased, change in accuracy is minimal.

To understand the benefit of using several prior data points for training, in Figure 4 we vary the training size from 1 sample to 50 samples and test the accuracy of predicting appliance usage for SVM and Bayesian (the highest-performing machine learning algorithms) using time horizons of an hour and a day across the appliances in Site 3. We used Site 3 for this experiment since our dataset has both appliance power draw and whole-home energy consumption data from this home over 18 months. We find that an increase in the training size does not change accuracy much. In fact, the net change is less than 5%, confirming that training a model using prior data points are unlikely to be advantageous.

Next, Figure 6 considers whether combining the power draw feature with the additional features such as time of the day, day of the week, and quarter of the day improves prediction performance for appliances usage and the home energy

Appliance	Feature	dist-covariance	ts
Lamp (hour)	Power	0.426	1738.68
Lamp (hour)	Time	0.142	555.97
Lamp (week)	Power	0.016	0.659
Television (day)	Power	0.548	53.15
Television (day)	Time	0.024	3.13
Television (week)	Power	0.365	6.97
Washer (week)	Power	0.059	2.53
Washer (week)	Time	0.085	2.84
Home (hour)	Power	0.801	2139.08
Home (hour)	Time	0.005	16.95
Home (week)	Power	0.497	14.34
Two Homes (hour)	Power	0.800	353.19
Two Homes (week)	Power	0.824	5.49

Fig. 5. The table compares the dependence between the power and time features, and the power draw for appliances and whole-home energy consumption for different time horizons. To calculate this dependence, we use distance correlation and test statistic values.

usage. In this experiment, we use SVM and consider the results for all sites. Again, the results show that using additional features does not improve performance. Additionally, Figure 6 shows that our models achieved greatest accuracy when based only on information about prior power measurements; adding the time-based features did not improve accuracy.

To understand this artifact further, we used a statistical metric, distance-correlation [25] to quantify the dependence between features and home energy consumption and appliance power draw. Distance-correlation models the dependence between two random vectors. If the feature vector is denoted by X and the power draw of an appliance is denoted by Y , and A and B denote the standard matrix of euclidean distances with rows or column means subtracted and grand mean added, the distance covariance $dCov^2 = (1/n^2) \cdot \sum_{i,j} A_{ij} \cdot B_{ij}$, and the distance correlation $R^2 = dCov(X, Y)^2 / (dCov(X) \cdot dCov(Y))$. The distance correlation is a normalized value between 0 and 1, where a result of 0 means the two random vectors being compared are fully independent, and a higher value implies a higher degree of statistical dependence. We also present the value of the test statistic (ts) [25] that also models the dependence between the two random vectors. The test statistic is similar to distance correlation, and though not bounded, a higher test statistic implies a greater statistical dependence between two sequences. Table 5 presents the distance correlation and ts values for a set of appliances for a single site (site 3), the whole-home energy consumption of the site, and the aggregate energy consumption of two homes. While the test statistic shows that the power and time features have *some* dependence on the power draw of appliances, the distance correlation values are higher for power features. This provides evidence of a stronger tie between power features and home and appliance energy consumption. A more qualitative argument for this observation is that the power features encapsulate more information than the time of the day features. For instance, the power draw by an appliance in the last quarter-day encapsulates a time parameter, the 6-hour quarter-of-day block, and a power parameter, the power draw during that time. Unfortunately, a time feature like quarter of the day is single dimensional and encodes sparser information.

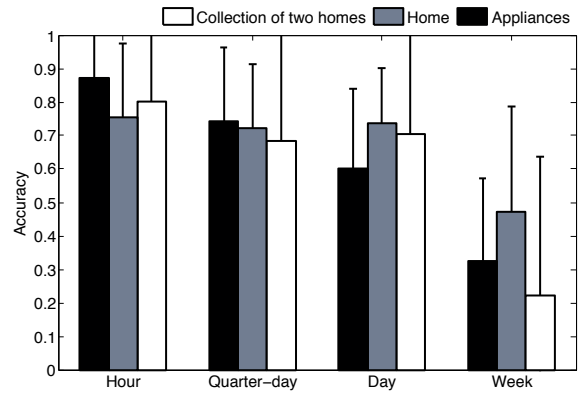


Fig. 7. Comparison of prediction accuracy for SVM for whole-home energy consumption, aggregate energy consumption of two homes, and per-appliance power draw for four different time horizons.

Overall, these results suggest that there is an inherent *unpredictability* and *randomness* associated with energy usage and provide evidence that using machine learning or time-series predictors for appliance- or home-level forecasting is unlikely to be successful, overall. We note, however, that these results provide an aggregate view across all appliances and later in this section we consider whether specific appliance categories are more predictable.

Implication to Energy Management: Demand-response systems rely heavily on predictive models for appliance usage as well as home and neighborhood energy consumption. Most demand-response systems must, in realtime, determine future energy usage and make adaptive decisions for balancing demand with supply or load balancing demand. One of the implications of Insight 1 is the feasibility of replacing sophisticated learning algorithms with simple statistics for prediction. Since a simple statistic-based algorithm that keeps track of the power draw in the last time cycle has low computational and storage needs, it can substantially reduce the cost and improve the performance of demand-response systems.

Insight 2: Predictions of appliance usage are more accurate than prediction of home energy usage at short time horizons while predictions of home energy usage are more accurate at longer prediction horizons

In our next set of experiments, we compare the prediction accuracy of appliance power draw and whole-home energy consumption for the four time scales. For this experiment, we use the SVM algorithm and use only the power draw feature. The results, shown in Figure 7, demonstrate that at shorter time horizons—hour, quarter-day, and day—the prediction accuracy for appliance usage is higher than for whole-home energy consumption. However, at the longer horizon of a week the algorithm has a higher prediction accuracy for the whole-home energy consumption. The prediction accuracy for the energy consumed by a collection of two homes does not show any trends.

To quantitatively understand this behavior, we again consider the distance-correlation between the power features and average power draw. Table 5 presents the results for a subset of appliances and the whole home for short and long time

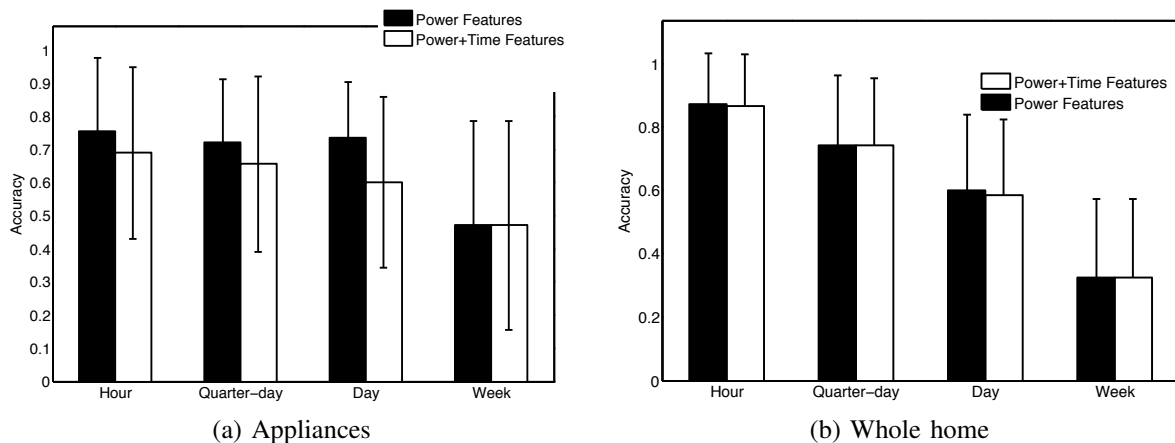


Fig. 6. The graphs show experiments performed using SVM. (a) Comparison of the prediction accuracy for appliances using power features and a combination of power and time-of-day features. (b) Comparison of the prediction accuracy for home energy consumption using power features and a combination of power and time-of-day features.

horizons for site 3. The table shows that for the power features the distance correlation values are higher for appliances at shorter horizon than whole-home energy consumption, but is lower for longer horizons. The test statistic values also show the same trend. A more qualitative argument for the above behavior can be made based on the factors that affect appliance and whole-home energy consumption. To predict the power draw by a device an hour in advance, our model uses a single feature: the power draw in the previous hour. This approach works well for appliances whose use is highly context dependent. For instance, the power use of a lamp likely depends on room occupancy and incoming sunlight. If a room is occupied and there is little sunlight in one hour, it is likely to be occupied and have little sunlight in the next hour resulting in similar use of the lamp in both hours. A whole home, on the other hand, has multiple concurrent users, a large number of potential values for its instantaneous power draw, and a high volume of relevant contexts to cause a specific amount of energy use in a given moment. These differences suggest that our predictor, when predicting a whole home's energy consumption, will show different accuracies for the same timeframes. This is borne out by our experiments. We only achieve 70% accuracy when predicting whole home energy use for one hour, versus 85% for predicting a single appliance. At longer timeframes, a week for instance, accuracy of whole home is better than appliances—42% compared to 30%. This is because the users have the tendency of performing the same set of activities over time frames of a week. For instance, washing and drying will be performed once a week, although not necessarily on a particular day of the week. Therefore, for homes the cumulative usage of all appliances show more predictability over longer time horizon of a week.

Implications to Energy Management: The above result has strong implications to the way demand-response systems must be designed. Since the prediction accuracy for longer time horizons, both for appliances usage and whole-home energy consumption, is close to 50%, a demand-response system must account for this error in prediction. An approach to account for this error is to use stochastic optimization frameworks [26] where the predicted power draw is probabilistic. Moreover,

proposed recommendation systems that utilize optimization frameworks to balance demand with supply or balance demand with a user-defined goal [27] must leverage the high prediction accuracy of homes at longer time horizons and high prediction accuracy of appliances at shorter time horizons. A method of leveraging this dichotomy would be to determine the energy savings required over a longer time horizon like a week, but make energy saving recommendations for changing appliance usage over shorter time horizons. For example, if the goal of a homeowner is to reduce his electricity bill by 20% in a week, the system can predict the home energy usage over a week to determine the energy usage goal, and then make recommendations for reducing usage of appliances like televisions or lighting at short horizons such as one hour.

Insight 3: There is large variance in predictability across appliance types

In our next set of experiments, we perform an exhaustive analysis on the predictability of different appliance and home types. For our analysis, we first divide all the home appliances into seven categories: (1) Computers (PCs, laptops, desktops); (2) Cooking appliances (microwaves, ovens); (3) Entertainment appliances (televisions, gaming consoles); (4) Lighting (lamps); (5) Fridges (including mini-fridges); (6) WashingDrying (clothes washers and dryers); and (7) Miscellaneous (Keizers, Sonos, treadmills, and other appliances).

Figure 8 compares the prediction accuracy for the seven appliance categories for all five algorithms. The predictions are made for a time frame of one hour in the future. Figure 8(a) shows the prediction accuracy for all predictions and Figure 8(b) illustrates the prediction accuracy for instances when the power draw of the device was positive (i.e., the device was on). We can draw several conclusions from the figure. First, unlike Insight 1, we find that for certain appliance types, using machine learning algorithms and time-series predictors does yield better prediction results. The SIMPLE algorithm, for instance, has a poor accuracy of 15% and 60% for refrigerators and lighting while the SVM algorithm has better accuracy of 60% and 91%. For other appliance types, such as cooking appliances, computers, and washers/dryers, the

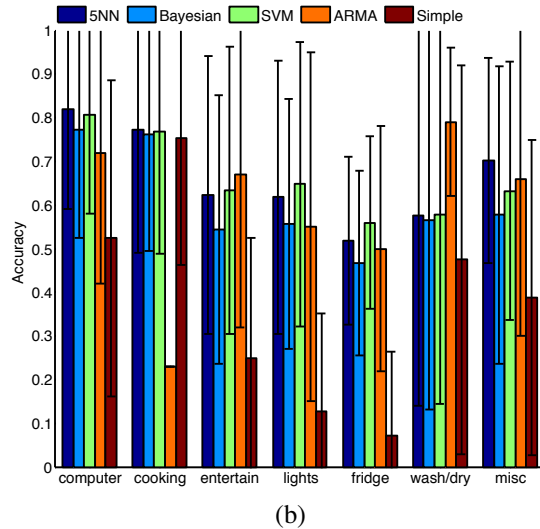



Fig. 8. (a) The accuracy of predicting the power draw of seven appliance types across seven sites in the GreenHomes dataset. (b) The accuracy of predicting the power draw of seven appliance types when the devices had a positive power draw. For both graphs the predictions were performed for a time horizon of one hour in the future.




Fig. 9. The prediction accuracy for the seven appliance types for different time horizons using the SVM predictor.

SIMPLE algorithm performs as well as the other algorithms. This implies that there is some variation in the correlation between the power draw by appliances like lighting between hours t and $t - 1$. A training set, therefore, helps in better prediction of these appliance types. Our second observation is that certain appliance types are more predictable than others. Computers and cooking appliances have fairly high accuracy of prediction while entertainment and lighting have medium or low prediction accuracies. There is high chance that a computer or cooking appliance will be used in time t if it is used in the previous hour. Washers and dryers shows high accuracy for prediction in Figure 8 (a), however, that is because they are switched off most of the time. As Figure 8 (b) shows, the accuracy of prediction for washers/dryers when they are on is less than 50%.

At the time horizon of one hour computers and cooking appliances are more predictable than refrigerators, washers, and lighting. Figure 9 compares the predictability of different




Fig. 10. A comparison of the prediction accuracies of whole-home energy consumption for user-owned and rental apartments for four time horizons. We use SVM for this experiment.

appliance types for all four time horizons using SVM. As described in Insight 2, appliances have poor predictability at longer time horizons, and the predictability monotonically degrades from hour to week. The drop in accuracy is most visible for cooking appliances. The accuracy drops from 95% to 15%, as the prediction horizon changes from an hour to a week, demonstrating that power draw of cooking appliances in the last week is not an accurate predictor and usage changes from week to week.

Insight 4: There is a large variation in predictability across homes and predictability is impacted by number of occupants, home type (off-grid, grid-tied, or on-grid), and home ownership type (rental or personally owned)

We next compare the predictability of whole-home energy consumption for four sites in our study. The homes in our study include off-grid, grid-tied, on-grid, rentals, and user-owned

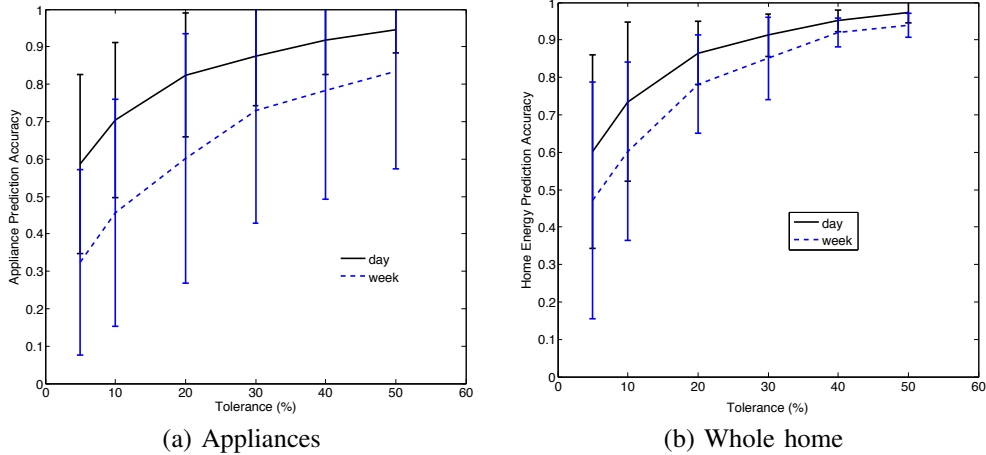


Fig. 12. (a) Accuracy as a function of tolerance averaged over all appliances in the dataset. (b) Accuracy as a function of tolerance averaged over all homes where the dataset has whole-home energy consumption data.




Fig. 11. A comparison of appliance power draw prediction accuracies for the seven sites in our deployment.

homes. Figure 10 compares the predictability of the whole-home energy consumption for the four different time horizons for rentals and user-owned apartments. The figure shows a general trend where rental apartments are less predictable than user-owned houses. We conjecture that this is due to the difference in incentives towards saving energy in user-owned homes and rental apartments. In rental apartments, the electricity bill is often included in the monthly bill (that is true for the two sites in the figure), and hence residents do not have the incentive to follow regular schedules for appliance usage to save electricity. On the other hand, for user-owned homes, the homeowner is responsible for his electricity bill, and hence the households follow more regular schedules of appliance usage. As future work, we plan to further validate this conjecture through a usability study.

Figure 11 compares the predictability of all appliances for all seven sites using SVM. There is large variance in the predictability of appliances across homes. Site 3 shows the highest predictability for appliance usage for the four time horizons. Site 3 is a user-owned house that is occupied by

a single user, and therefore, is more predictable. Site 5, a house that is grid-tied and partially powered by solar panels, also shows high predictability of appliances. Site 0 and Site 4 are rental apartments that show lower predictability. The off-grid and grid-tied homes have more regular appliance usage, and are more predictable, than the on-grid home appliances. The figure, overall, shows that there is a large variation in predictability across homes, and there are several factors like house type and number of occupants that affect predictability of energy usage.

Implications to Energy Management: Insights 3 and 4 demonstrate the need for *personalized* energy management systems. The predictability of usage varies considerably across homes and appliance types. It is imperative, therefore, that any energy saving recommendation system or demand management system should be fine-tuned to the household. Insight 3 also demonstrates that certain appliance types have more predictable usage, which should be considered when designing a demand-response system. We hypothesize, for example, that the user may be most likely to reduce usage of appliances that show lower predictability at shorter time horizons, for example Entertainment appliances. The reasoning behind this conjecture is that low predictability shows use patterns that are already irregular and, thus, may be changed with little impact. A notable exception, however, would be Refrigerators, which show low predictability at all time horizons. In this case, the salient insight is that for this class of appliance alternative to predictive modeling such as asking the user is necessary to determine energy needs.

C. Discussion

Our prior insights assume that precisely accurate prediction is necessary for demand response and energy management. There are, however, applications where a more coarse-grained definition of accuracy might suffice, for instance, certain applications might be interested in predicting whether an appliance is on or off [7], and other applications for demand response and load balancing might be more tolerant to the actual error between the predicted and actual energy consumption. We conclude with an experiment that explores the tradeoff between

this tolerance and accuracy of prediction for a broader class of applications. Recall that our previous definition of accuracy considered a prediction accurate if the predicted power draw was within a 5% of the maximum observed power draw of the appliance. In this case, the *tolerance* is 5%. For this experiment, we vary the *tolerance* from 5% to 50%, and show the accuracy of prediction of appliances and whole-home energy consumption for two prediction horizons—hour and week (Figure 12 (a) and (b)). The figure shows that if the definition of tolerance is relaxed and the application can tolerate higher errors, SVM can have an accuracy of close to 90% for longer horizons. Hence, such applications can reliably use appliance and whole-home energy consumption predictions for longer horizons.

V. RELATED WORK

Our work builds on previous work on forecasting algorithms applied to various application domains, predictive models for home energy usage, and general demand-response systems. Here we compare and contrast our work with the most relevant literature.

Forecasting Algorithms: Forecasting and developing predictive models is a well studied area with applications to weather monitoring [28], modeling predicted energy harvesting [29], and traffic forecasting [30]. Forecasting algorithms include time-series based algorithms like autoregressive and moving average models, kalman filters [31], [2] and machine learning algorithms such as Bayesian predictors [32] and support vector machines [33]. Our work is complimentary to the algorithms developed for forecasting. Our goal is to use existing algorithms to understand how predictable home energy usage is. For our analysis we use the *k*-nearest neighbor, support vector machine, a bayesian predictor, and a ARMA predictor. Our study can inform the design of demand response systems for homes.

Predictive models for home energy usage: This paper builds on previous work on developing predictive models for home energy usage. Most of the focus in this area has been on developing prediction algorithms that produce the accuracy required for specific applications like matching demand and supply in homes. The algorithms developed include autoregressive models [6], and machine learning algorithms [4], [34], and consider contextual data like weather in conjunction with home energy consumption data. Since the papers written on predictive models has focussed on specific applications, there is a lack of studies that perform a comparison analysis of predictability of energy usage of appliances, homes, and a collection of homes at different timeframes such as an hour, quarter-day, day, and week into the future. The goal of this paper is to draw insights on predictability of energy usage at different granularities and time horizons.

Demand-response and Energy Management: Insights into predictability of energy usage in the home can inform demand-response systems in particular and home energy management in general. Hence, our predictive analysis is complementary to several demand-response systems in homes [27], [35] and can be used to improve their performance. These include systems that flatten peak energy consumption [36], methods to predict energy generation and consumption [37], balance

energy demand with supply [38], and minimize whole building energy consumption. The GreenHomes dataset is an useful addition to rich datasets like REDD [10] and Smart* [11] that provide publicly available energy consumption data for homes, appliances, and micro-grids.

VI. CONCLUSION

In this paper, we study the predictability of home energy consumption through a analysis of appliance power draw and whole-home energy consumption in seven homes in the United States. Our dataset, GreenHomes, consists of data on appliance power draw and home energy consumption over a period of 2 years. Our predictive analysis, using three machine learning algorithms, a time-series predictor, and a simple statistic-based algorithm, presents several interesting insights that have implications to the design of home energy management and demand-response systems. We show that simple algorithms that use basic power features like power consumed in the previous timeframe have comparable performance to sophisticated machine learning algorithms and time series predictors. These simple algorithms can considerably reduce the computational need for large-scale predictive analysis. We also show that appliance-level energy usage is more predictable than whole-home energy at shorter time horizons while home-level energy consumption is more predictable at longer time horizons. Moreover, there is a large variation across appliance types and homes, and depends on the type of home. This insight calls for the need for personalized home energy management solutions.

REFERENCES

- [1] V.G. Vovk. Universal forecasting algorithms. *Information and Computation*, 96(2):245 – 277, 1992.
- [2] Hae Young Noh and Ram Rajagopal. Data-driven forecasting algorithms for building energy consumption, 2013.
- [3] A. Barbato, A. Capone, M. Rodolfi, and D. Tagliaferri. Forecasting the usage of household appliances through power meter sensors for demand management in the smart grid. In *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, pages 404–409, 2011.
- [4] Holger Ziekow, Christoph Goebel, Jens Strucker, and Hans-Arno Jacobsen. The potential of smart home sensors in forecasting household electricity demand. In *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, pages 229–234, 2013.
- [5] Richard E. Edwards, Joshua New, and Lynne E. Parker. Predicting future hourly residential electrical consumption: A machine learning case study. *Energy and Buildings*, 49(0):591 – 603, 2012.
- [6] Rayman Preet Singh, Peter Xiang Gao, and Daniel J. Lizotte. On hourly home peak load prediction. In *SmartGridComm*, pages 163–168. IEEE, 2012.
- [7] Ngoc Cuong Truong, Long Tran-Thanh, Enrico Costanza, and Sarvapali D. Ramchurn. Towards appliance usage prediction for home energy management. In *Proceedings of the Fourth International Conference on Future Energy Systems, e-Energy '13*, pages 287–288, New York, NY, USA, 2013. ACM.
- [8] Ngoc Cuong Truong, James McInerney, Long Tran-Thanh, Enrico Costanza, and Sarvapali D. Ramchurn. Forecasting multi-appliance usage for smart home energy management. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI'13*, pages 2908–2914. AAAI Press, 2013.
- [9] M.S. Kandil, S.M. El-Debeiky, and N. E. Hasaniien. Long-term load forecasting for fast developing utility using a knowledge-based expert system. *Power Systems, IEEE Transactions on*, 17(2):491–496, 2002.
- [10] J Zico Kolter and Matthew J Johnson. Redd: A public data set for energy disaggregation research. In *proceedings of the SustKDD workshop on Data Mining Applications in Sustainability*, pages 1–6, 2011.

- [11] Sean Barker, Aditya Mishra, David Irwin, Emmanuel Cecchet, Prashant Shenoy, and Jeannie Albrecht. Smart*: An open data set and tools for enabling research in sustainable homes. *SustKDD*, August, 2012.
- [12] Zwave energy meters. <http://aeotec.com/z-wave-home-energy-measure>.
- [13] Z-wave current transformers. <http://www.vemarsas.it/it/e-commerce/adaptive-automation-system/z-wave-clamp-power-meter---two-clamps-detail>.
- [14] Fitpc computer. <http://www.fit-pc.com/web/>.
- [15] Colin Dixon, Ratul Mahajan, Sharad Agarwal, AJ Brush, Bongshin Lee, Stefan Saroiu, and Paramvir Bahl. An operating system for the home. In *NSDI 2012*.
<http://micasaverde.com/vera.php>. Micasaverde vera2.
- [17] 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.
- [18] Sami Rollins and Nilanjan Banerjee. Using rule mining to understand appliance energy consumption patterns. *IEEE PerCom*, 2014.
- [19] Pat Langley and Stephanie Sage. Induction of selective bayesian classifiers. In *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, UAI'94, pages 399–406, San Francisco, CA, USA, 1994. Morgan Kaufmann Publishers Inc.
- [20] Sahibsingh A Dudani. The distance-weighted k-nearest-neighbor rule. *Systems, Man and Cybernetics, IEEE Transactions on*, (4):325–327, 1976.
- [21] Terrence S Furey, Nello Cristianini, Nigel Duffy, David W Bednarski, Michel Schummer, and David Haussler. Support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics*, 16(10):906–914, 2000.
- [22] Hirotugu Akaike. Fitting autoregressive models for prediction. *Annals of the institute of statistical mathematics*, 21(1):243–247, 1969.
- [23] <http://en.wikipedia.org/wiki/DBSCAN>. Dbscan clustering algorithm.
- [24] Abdulhamit Subasi. A decision support system for diagnosis of neuromuscular disorders using dwt and evolutionary support vector machines. *Signal, Image and Video Processing*, pages 1–10, 2013.
- [25] Gábor J Székely, Maria L Rizzo, and Nail K Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769–2794, 2007.
- [26] M.V.F. Pereira and L.M.V.G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1-3):359–375, 1991.
- [27] Tanuja Bapat, Neha Sengupta, Sunil Kumar Ghai, Vijay Arya, Yedendra Babu Shrinivasan, and Deva Seetharam. User-sensitive scheduling of home appliances. In *Proceedings of the 2Nd ACM SIGCOMM Workshop on Green Networking*, GreenNets '11, pages 43–48, New York, NY, USA, 2011. ACM.
- [28] Brad Colman, Kirby Cook, and Bradley J Snyder. Numerical weather prediction and weather forecasting in complex terrain. In *Mountain Weather Research and Forecasting*, pages 655–692. Springer, 2013.
- [29] Ekaterina Vladislavleva, Tobias Friedrich, Frank Neumann, and Markus Wagner. Predicting the energy output of wind farms based on weather data: Important variables and their correlation. *Renewable Energy*, 50:236–243, 2013.
- [30] M. Papadopouli, E. Raftopoulos, and H. Shen. Evaluation of short-term traffic forecasting algorithms in wireless networks. In *Next Generation Internet Design and Engineering, 2006. NGI '06. 2006 2nd Conference on*, pages 8 pp.–109, 2006.
- [31] R. H. Shumway and D. S. Stoffer. An approach to time series smoothing and forecasting using the em algorithm. *Journal of Time Series Analysis*, 3(4):253–264, 1982.
- [32] Jaime S Ide, Pradeep Shenoy, J Yu Angela, and R Li Chiang-shan. Bayesian prediction and evaluation in the anterior cingulate cortex. *The Journal of Neuroscience*, 33(5):2039–2047, 2013.
- [33] A Mellit, A Massi Pavan, and M Benganem. Least squares support vector machine for short-term prediction of meteorological time series. *Theoretical and Applied Climatology*, 111(1-2):297–307, 2013.
- [34] A. Barbato, A. Capone, G. Carello, M. Delfanti, M. Merlo, and A. Zaminga. House energy demand optimization in single and multi-user scenarios. In *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, pages 345–350, 2011.
- [35] 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.
- [36] 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.
- [37] 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.
- [38] Nilanjan Banerjee, Sami Rollins, and Kevin Moran. Automating Energy Management in Green Homes. In *ACM Sigcomm Workshop on Home Networks*, August 2011.