Learning to REDUCE: A Reduced Electricity Consumption Prediction Ensemble

Saima Aman
Computer Science Department
University of Southern California
Los Angeles, California
Email: saman@usc.edu

Charalampos Chelmis and Viktor K. Prasanna
Electrical Engineering Department
University of Southern California
Los Angeles, California
Email: chelmis@usc.edu, prasanna@usc.edu

Acknowledgement

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000192.

Disclaimer

Learning to REDUCE: A Reduced Electricity Consumption Prediction Ensemble

Saima Aman
Computer Science Department
University of Southern California
Los Angeles, California
Email: saman@usc.edu

Charalampos Chelmis and **Viktor K. Prasanna**
Electrical Engineering Department
University of Southern California
Los Angeles, California
Email: chelmis@usc.edu, prasanna@usc.edu

Abstract

Utilities use Demand Response (DR) to balance supply and demand in the electric grid by involving customers in efforts to reduce electricity consumption during peak periods. To implement and adapt DR under dynamically changing conditions of the grid, reliable prediction of *reduced consumption* is critical. However, despite the wealth of research on electricity consumption prediction and DR being long in practice, the problem of reduced consumption prediction remains largely un-addressed. In this paper, we identify unique computational challenges associated with the prediction of reduced consumption and contrast this to that of normal consumption and DR baseline prediction. We propose a novel ensemble model that leverages different sequences of daily electricity consumption on DR event days as well as contextual attributes for reduced consumption prediction. We demonstrate the success of our model on a large, real-world, high resolution dataset from a university microgrid comprising of over 950 DR events across a diverse set of 32 buildings. Our model achieves an average error of 13.5%, an 8.8% improvement over the baseline. Our work is particularly relevant for buildings where electricity consumption is not tied to strict schedules. Our results and insights should prove useful to the researchers and practitioners working in the sustainable energy domain.

Introduction

One of the critical challenges confronting modern societies is the need to attain energy sustainability. Buildings account for about 40% of the energy consumption worldwide (UNDP 2010) and novel energy optimization measures adopted in buildings can significantly contribute to energy sustainability. With the advent of Smart Grids, buildings are now being fitted with smart meters that record electricity usage every 15 minutes or less (Aman et al. 2015). Mining large amounts of electricity consumption data collected by smart meters provides insights into peak demand periods for buildings. Electric utilities can use these insights to ask building occupants and facility managers to reduce consumption during anticipated peak demand periods, a practice popularly know as **Demand Response** (DR). DR is defined as: "change in electric usage by end-use customers from their *normal* consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized" (FERC 2010). We address *reduced consumption* prediction during DR, which can help utilities in:

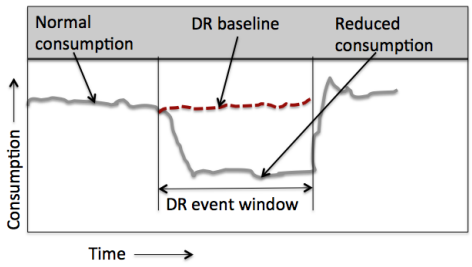

Figure 1: **Normal consumption, Reduced consumption, and DR baseline vis-a-vis a DR event.**

- estimating the extent of potential reduction during DR *before* the DR event occurs (Chelmis et al. 2015);
- performing dynamic DR at a few hours' advance notice whenever necessitated by the dynamically changing conditions of the grid, such as due to the integration of intermittent renewable generation sources (Aman et al. 2015);
- intelligently targeting customers for participation in DR based on a prediction of their reduced consumption and modifying such selection in real-time as needed (Ziekow et al. 2013); and
- estimating the amount of incentives to be given to the customers (Wijaya, Vasirani, and Aberer 2014).

Techniques that work well for normal consumption prediction, such as time series models, are ineffective for reduced consumption prediction due to 1) *abrupt changes in the consumption profile* at the beginning and end of the DR event (Figure 1); and 2) *insufficient recent observations* within the DR window for a time series model to be trained reliably. Instead, historical data from past DR events can be used as predictors for reduced consumption. Another challenge is that *reduced consumption is affected by several factors* such as the time of day and day of week, and DR factors such as curtailment strategy, human behavior, as well as environmental factors, such as temperature.

Table 1: **Normal consumption, Reduced consumption, and DR baseline: Key characteristics and challenges**

Prediction Task	Goal	Prior work	Timing	Historical Data	Compute Requirements	Profile changes
***Normal* Consumption**	Planning, DR	Several	Outside the DR event	Readily available	Off-line or real-time	Gradual
Counterfactual **DR Baseline**	Curtailment calculation	Several	During the DR event	Readily available	Off-line	Gradual
Reduced Consumption	Planning, DR, dynamic DR	None	During the DR event	Sparse or non-existent	Real-time for dynamic DR	Abrupt at the DR event boundaries

Our contributions in this paper are:

- We identify key characteristics and challenges of reduced consumption prediction problem.
- We use diverse predictors in a novel ensemble that use different sequences of daily electricity consumption on DR event days, as well as contextual attributes, for reduced consumption prediction. The *low computational complexity* of our model makes it ideal for real-time applications such as *dynamic demand response* (Aman et al. 2015).
- We evaluate our model on a large real-world dataset from a university microgrid. Our model achieves an average error of 13.5%, an 8.8% improvement of over the baseline.

Related Work

Electricity consumption prediction is studied in three contexts: 1) ***normal* consumption**, 2) **reduced consumption**, and 3) **DR baselines**, which differ greatly in terms of their scope and characteristics (Table 1). While electricity consumption is a widely studied problem (Martinez Alvarez et al. 2011), (Mathieu, Callaway, and Kiliccote 2011), (Alzate and Sinn 2013), the problem of reduced energy consumption prediction is a new and open problem with little existing research (Chelmis et al. 2015). This can be attributed to factors such as the unavailability of reduced consumption data; the human factors causing variance in response to DR; and cancellation of DR when found violating occupants' thermal comfort limits. The utilities have so far focused more on predicting normal consumption or DR baselines. DR baselines are calculated during a DR event (Park et al. 2014), (Coughlin et al. 2009) and estimate *the amount of electricity that would have been consumed in absence of a DR event* (Figure 1). They are *counterfactual* in that they give a theoretical measure of what the customer did not do, but would have done in absence of a DR event. Utilities generally use simple averaging models for DR baseline predictions due to their simplicity and reduced computational requirements (Coughlin et al. 2009). DR baselines are used to measure the extent of curtailment achieved during a DR event.

Preliminaries

Definition 1 *A **DR event** for a building is the period during which the building's electricity consumption is reduced (for e.g. by turning devices off or by turning them down to a lower consumption setting than normal operation).*

A day in which a DR event occurs is called **DR day** (Wijaya, Vasirani, and Aberer 2014).

Definition 2 *A **daily sequence** of electricity consumption observations for a building on the i-th DR day, $\mathcal{E}_i = \{e_{i,1}, e_{i,2}, ..., e_{i,J}\}$, where $e_{i,j}$ is the observation at time interval j, and J is the number of intervals in a day.*

For simplicity, we assume all data to be sampled at the same frequency; hence all daily sequences are of the same length J. The set of all daily sequences from DR days for a building is an $I \times J$ matrix, $\mathcal{E} = (\mathcal{E}_1, \mathcal{E}_2, ..., \mathcal{E}_I)^T$ where I is the number of DR days observed for the given building.

Definition 3 *The **pre-DR sequence** $\mathcal{E}_{i,1,d-1} = \{e_{i,1}, e_{i,2}, ..., e_{i,d-1}\}$ with $d > 1$, is a subsequence of $\mathcal{E}_i$ beginning at interval 1 and ending just before d, the interval at which a DR event begins.*

Definition 4 *The **in-DR sequence** $\mathcal{E}_{i,d,L} = \{e_{i,d}, e_{i,d+1}, ..., e_{i,d+L-1}\}$, with $d > 1$, and $d + L - 1 \leq J$, is a subsequence of $\mathcal{E}_i$ that begins at time interval d, the interval at which a DR event begins and is L intervals long.*

Contextual Attributes

We consider two types of contextual attributes: **time-series attributes** that are defined for each time interval of the day, and **static attributes** that remain same for all intervals of a day. Examples of the former include temperature, humidity, dynamic pricing, occupancy, etc. while for the latter include day of week, holidays, etc. For N_t distinct time-series attributes, and N_s static attributes, we define the following:

Definition 5 *The **daily context** for a building on the i-th DR day is a tuple, $\mathcal{C}_i = \langle A_i[1], \ldots, A_i[N_t], B_i[1], \ldots, B_i[N_s] \rangle$, where $A_i[k] = \{a_{i,1}, a_{i,2}, ..., a_{i,J}\}$ is the k-th time series attribute and $B_i[k]$ is the k-th static attribute.*

Definition 6 *The **pre-DR context** for a building on the i-th DR day is a tuple, $\mathcal{C}_{i,1,d-1} = \langle A_i[1], \ldots, A_i[N_t], B_i[1], \ldots, B_i[N_s] \rangle$ where $A_i[k] = \{a_{i,1}, a_{i,2}, ..., a_{i,d-1}\}$ is a subsequence of the k-th time series attribute from interval 1 to just before interval d when the DR begins, and $B_i[k]$ is the k-th static attribute.*

Table 2: Notations

Symbol	Description
I	Number of DR days observed
J	Number of observations in a day
L	Length of the DR event window
$e_{i,j}$	Electricity consumed on day i in interval j
$\mathcal{E}_i$	Daily DR sequence for day i
$\mathcal{E}_{i,s,l}$	Subsequence of $\mathcal{E}_i$ starting at s of length l
$\mathcal{C}_i$	Daily context for day i
$\mathcal{C}_{i,s,l}$	Subsequence of $\mathcal{C}_i$ starting at s of length l
$A_i[k]$	Vector of k-th time series attribute for day i
$B_i[k]$	k-th static attributes for day i

Problem Definition

We formulate the problem of predicting reduced electricity consumption for a building during a DR event as the problem of calculating the values of in-DR sequence $\mathcal{E}_{i,d,L}$ for DR day i, given the pre-DR sequence $\mathcal{E}_{i,1,d-1}$ and pre-DR context $\mathcal{C}_{i,1,d-1}$ for the day i, and the set of daily sequences $\mathcal{E}$ and daily contexts $\mathcal{C}$ from the historical data.

REDUCE Ensemble

We propose REDUCE (Reduced Electricity Consumption Ensemble) that learns to combine outputs from **three base models** using a Random Forest[1] approach (Breiman 2001) to do reduced consumption prediction. To train the ensemble, we form a set of predictor tuples $([\hat{\mathcal{E}}_{\epsilon,d,L}]_{IDS}, [\hat{\mathcal{E}}_{\epsilon,d,L}]_{PDS}, [\hat{\mathcal{E}}_{\epsilon,d,L}]_{DSS})$, where $[\hat{\mathcal{E}}_{\epsilon,d,L}]_m$ is the value predicted by base model m; and corresponding set of responses $\mathcal{E}_{\epsilon,d,L}$ from the observed values for each day ϵ.

IDS: In-DR Sequence Model

IDS is similar to the approach used by the utility for predicting DR baselines. While utilities average over a set of past similar (non-DR) days, IDS averages all in-DR sequences from past DR days. Thus, the in-DR sequence for each building during the i-th DR day is given by:

$$[\hat{\mathcal{E}}_{i,d,L}]_{IDS} = \frac{1}{|\mathcal{E}|}\sum_{\epsilon=1}^{|\mathcal{E}|}\mathcal{E}_{\epsilon,d,L} \tag{1}$$

This model offers two key advantages: 1) low computational complexity, as computation time is independent of the length of the DR event and the size of the historical data, making it ideal for real-time predictions for dynamic DR, and 2) it is a univariate model that only depends on electricity consumption values and does not require additional variables, which would increase data collection costs (Aman, Simmhan, and Prasanna 2015).

[1]Implemented by the randomForest R package (Liaw and Wiener 2002).

PDS: Pre-DR Sequence Similarity Model

This model considers the contextual attributes and the electricity consumption values on the DR day *before* the beginning of DR, for selecting "similar" DR days from the past data. Our *hypothesis is that if two DR days have similar pre-DR sequences, their in-DR sequences would be similar*. Thus, for each DR day i in the historical data we first form a tuple of pre-DR sequence and pre-DR context $\langle \mathcal{E}_{\epsilon,1,d-1}, \mathcal{C}_{\epsilon,1,d-1}\rangle$. Similarly, we form a tuple for DR day i. The similarity score between each DR day ϵ in the historical data and given DR day i is given by

$$\begin{aligned} SimScore(\epsilon, i) = sim(&\langle \mathcal{E}_{\epsilon,1,d-1}, \mathcal{C}_{\epsilon,1,d-1}\rangle, \\ &\langle \mathcal{E}_{i,1,d-1}, \mathcal{C}_{i,1,d-1}\rangle). \end{aligned} \tag{2}$$

where sim can be any similarity measure.

Next, we sort historical days based on their similarity score to DR day i in descending order. We then predict the in-DR sequence on a given day as a weighted average of historical in-DR conditions, such that higher weights are assigned to days with a higher similarity score. The weights are chosen to exponentially decrease with decreasing similarity score. The predicted in-DR sequence is given by

$$[\hat{\mathcal{E}}_{i,d,L}]_{PDS} = \frac{1}{|\mathcal{E}|}\sum_{\epsilon=1}^{|\mathcal{E}|}\omega_\epsilon \times \mathcal{E}_{\epsilon,d,L} \tag{3}$$

where weights $\omega_\epsilon = exp(-\lambda)$, and $0 < \lambda \leq 1$ is the decay rate that determines the rate of decrease of weights with decreasing similarity score.

DSS: Daily Sequence Similarity Model

This model considers the entire daily sequences and contexts in the historical data to first discover clusters of *daily profiles* for each building. We define daily profiles $\mathcal{P}_\epsilon = \langle \mathcal{E}_\epsilon, \mathcal{C}_\epsilon\rangle$ for each building to consist of tuples of daily sequences $\mathcal{E}_\epsilon$ and daily contexts $\mathcal{C}_\epsilon$ for each DR day ϵ in the historical data. We cluster the daily profiles using k-means clustering (Dudani 1976) into N_k clusters, $C = \{C_1, C_2, ..., C_{N_k}\}$. The number of clusters N_k is estimated by minimizing the within cluster sum of squares. The centroid c_m of each cluster C_m can be interpreted as the characteristic profile of the cluster:

$$c_m = \frac{1}{N_k}\sum_{\epsilon=1}^{N_k}\mathcal{P}_\epsilon \tag{4}$$

For a given day i, we calculate the probability of i belonging to cluster C_m using the pre-DR part of the daily profile for the i-th day and finding their similarity to the pre-DR part of the centroid vector's profile:

$$P(i \in C_m) = \frac{1}{\alpha\|\mathcal{P}_{i,1,d-1} - \mathcal{P}_{c_m,1,d-1}\|_2} \tag{5}$$

where α is a constant used to normalize the probability values between 0 and 1:

$$\alpha = \sum_{1}^{N_k}\frac{1}{\|\mathcal{P}_{i,1,d-1} - \mathcal{P}_{c_m,1,d-1}\|_2} \tag{6}$$

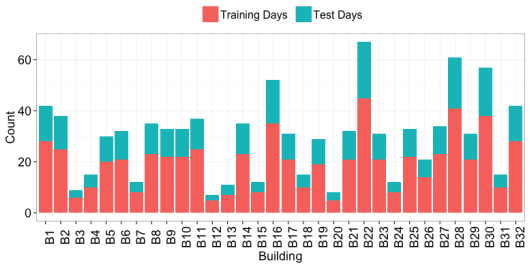

Figure 2: Distribution of DR events across buildings

The in-DR sequence of DR day i can then be calculated as the weighted sum of characteristic in-DR sequences with weights equal to the probability of DR day i belonging to the respective clusters, as follows:

$$[\hat{\mathcal{E}}_{i,d,L}]_{DSS} = \frac{1}{N_k} \sum_{m=1}^{N_k} P(i \in C_m) \times \mathcal{E}_{c_m,d,L} \qquad (7)$$

The clustering step is performed only once for each building on historical data whose size is small. However, the clustering step can be repeated periodically when more historical data is accumulated.

Time Complexity

IDS has $\mathcal{O}(1)$ run-time complexity. PDS involves a one-time step of sorting historical days based on similarity, with $\mathcal{O}(nlogn)$ complexity, where n is the number of days in the historical data. Thereafter, prediction with PDS is of $\mathcal{O}(1)$. DSS involves k-nn clustering as a one time step, while prediction is of $\mathcal{O}(1)$. REDUCE uses the random forest method with time complexity of $\mathcal{O}(nlogn)$ for the training step. Prediction is of $\mathcal{O}(1)$ complexity. Given this low time complexity, our proposed model is ideally suited for making *real-time predictions*.

Experiments

Dataset. Reduced consumption data was collected from 952 DR events (2012-2014) on 32 buildings in a University microgrid [Anonymized] (Figure 2). Consumption reduction was achieved via DR strategies (Piette, Kiliccote, and Dudley 2012) that directly reduce the loads or alter temperature settings. Temperature data was collected from the NOAA weather station located on the university campus.

Evaluation. We use MAPE (Mean Absolute Percentage Error) for evaluation. As a relative measure, MAPE is independent of the scale of consumption of a building (Aman, Simmhan, and Prasanna 2015). It is defined as $MAPE = \frac{1}{n}\sum_{i=1}^{n} \frac{|e_i - \hat{e}_i|}{e_i}$, where e_i and $\hat{e}_i$ represent observed and predicted electricity consumption respectively. (Our code and sample data is publicly available [Anonymized].) IDS, being popular with the utilities, is used as the baseline for comparing the performance of the ensemble.

Parameters. We use 15-minute granularity data resulting in $J = 96$ intervals per day. The DR events occurred on weekdays during 1 PM ($d = 54$) to 5 PM, the peak load period designated by the local utility [Anonymized]. Thus, the length of DR events is set at $L = 16$. We use one time-series attribute (i.e., $N_t = 1$), temperature, and seven ($N_s = 7$) static attributes to represent day of week based on a simple 1-of-7 encoding scheme.

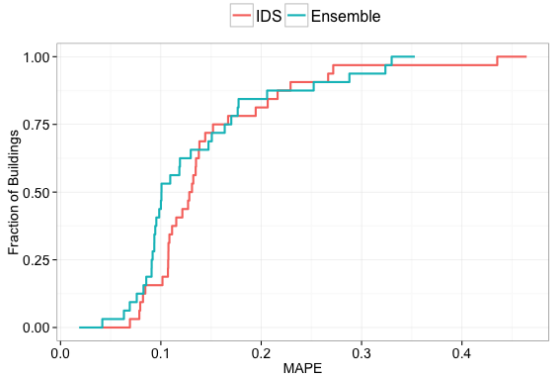

Figure 3: Cumulative density function (CDF) of MAPE

Results. Figure 4 shows the MAPE values for individual buildings, while Figure 3 shows the cumulative distribution function of average MAPE of all buildings. We observe that our ensemble REDUCE outperforms the baseline IDS for about 70% of the buildings. It also limits prediction error to $\leq$ 10% for over half of the buildings, which is considered highly *reliable* by domain experts (Aman, Simmhan, and Prasanna 2015). Overall, it achieves an average error of 13.5% and standard deviation of 7.3%, which is an improvement of 8.8% over the baseline. The baseline IDS performs reasonably well with 14.8% average error and 7.4% standard deviation (Figure 3) indicating that historical time of day averaging is a strong predictor for reduced electricity consumption. Although simple, it derives its predictive power from the errors being averaged out over the entire dataset, as well as from electricity consumption being strongly related to repetitive patterns of human activities. This repetition is more pronounced for campus buildings where activities are tightly coupled to class schedules.

Effect of Schedule. We examine two types of buildings: 1) *schedule-driven*, consisting primarily of classrooms, and 2) *non schedule-driven*, with few or no classrooms. For *B15*, *B21*, *B28*, and *B29*, which are non-schedule driven, REDUCE gives superior performance (Figure 4). IDS does not perform well here as there are no significant repetitive patterns in human activity such as those found in classroom-only buildings where activities are tightly coupled to class schedules. To further assess this difference, we analyze three buildings: B21, a building with large computer labs, and faculty and graduate student offices, B28, a campus center building with large meeting spaces, and a grand ballroom with seating for over 1000 people, and B14, an academic building with classrooms and few office spaces.

Figures 5a and 6a show that REDUCE gives low error for both B21 and B28. The error for B21 is low for all days of

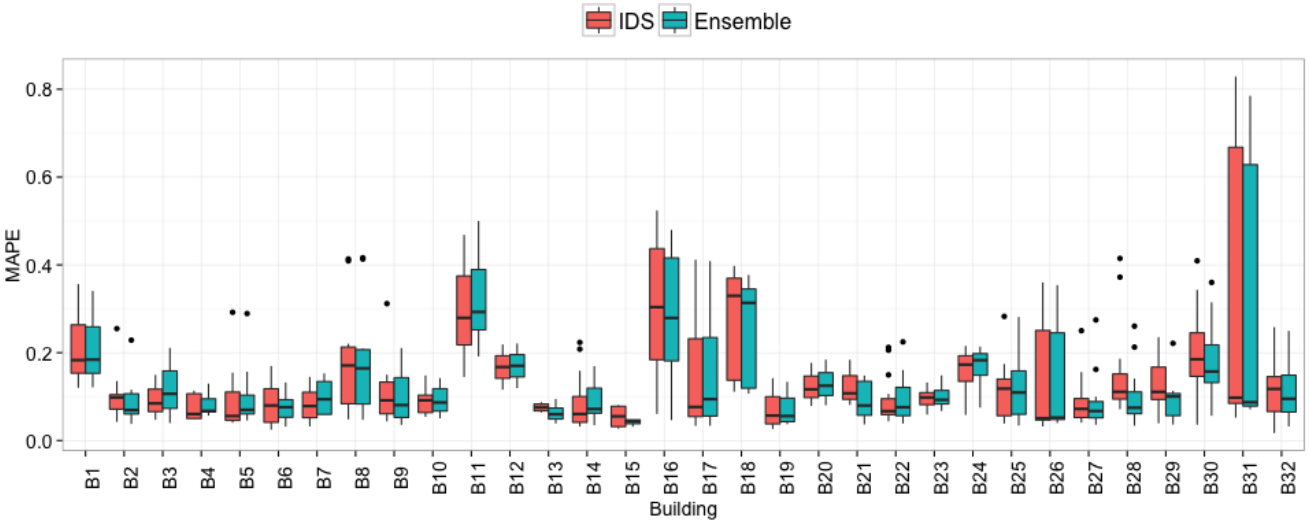

Figure 4: MAPE across buildings

week but one (Figure 5b), while for B28 it is low for all days of week (Figure 6b). We observe similar behavior across seasons (Figures 5c and 6c). On the contrary, IDS outperforms REDUCE for B14 (Figure 7a). REDUCE performs more consistently on Tuesdays, Thursdays, and Fridays, as compared to IDS (Figure 7b). It is notable that in Fall, when classes are scheduled, IDS performs well; however, in Summer when few classes are offered, and a variety of events may be occurring, REDUCE outperforms IDS (Figure 7c).

Insight 1: REDUCE gives superior performance when applied to buildings which do not follow a tight schedule. As a corollary, *we expect REDUCE to achieve similar performance for residential buildings, where human activities do not follow strict schedules*, and hence, the performance of averaging models such as IDS will deteriorate.

Effect of Training Data Size. Contrary to our ensemble REDUCE, the performance of IDS deteriorates with increasing size of the training data, which can be attributed to noise being introduced into the training dataset (Figure 8).

Insight 2: The performance of REDUCE is not sensitive to the training data size. As a corollary, *REDUCE would allow accurate predictions to be made with fewer historical data* which is useful for new buildings as well as for reducing computational and storage requirements.

Effect of Variance in Consumption. We observe that prediction error decreases with increasing average consumption for our ensemble REDUCE model, while it does not change for IDS (Figure 9). This could be attributed to more stable and predictable behavior for larger buildings, though it needs further investigation to understand this behavior. Also, for smaller buildings, the electricity consumption values are small; so even when the predicted value is offset by a small amount, it translates to a large percentage error.

Insight 3: The performance of REDUCE slightly improves for larger buildings.

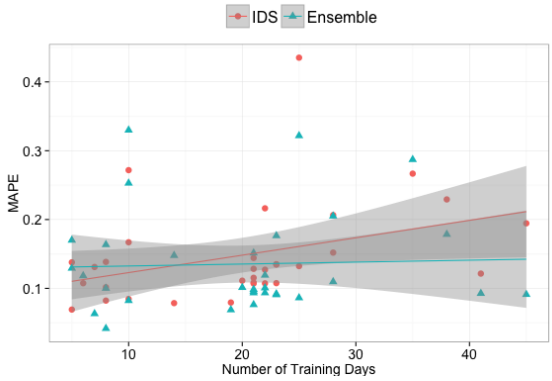

Figure 8: MAPE as a function of training data size

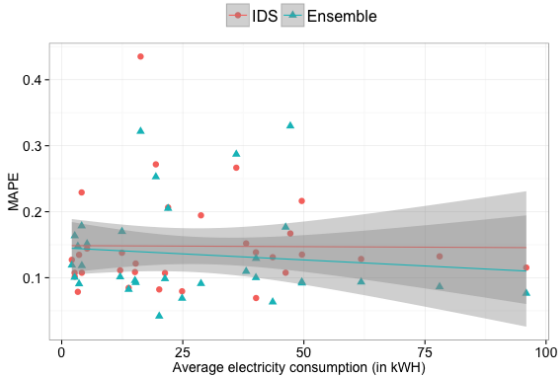

Figure 9: MAPE versus mean electricity consumption

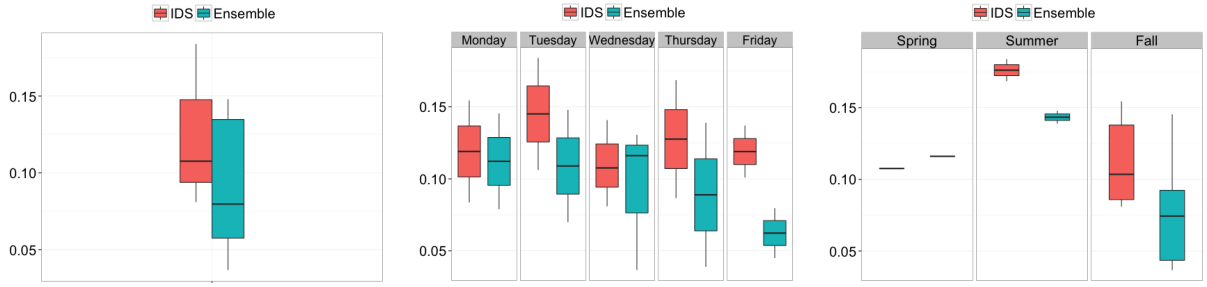

(a) For all test days (b) By day of week (c) By seasons

Figure 5: MAPE values for B21: a building with large computer labs, and faculty and graduate student offices

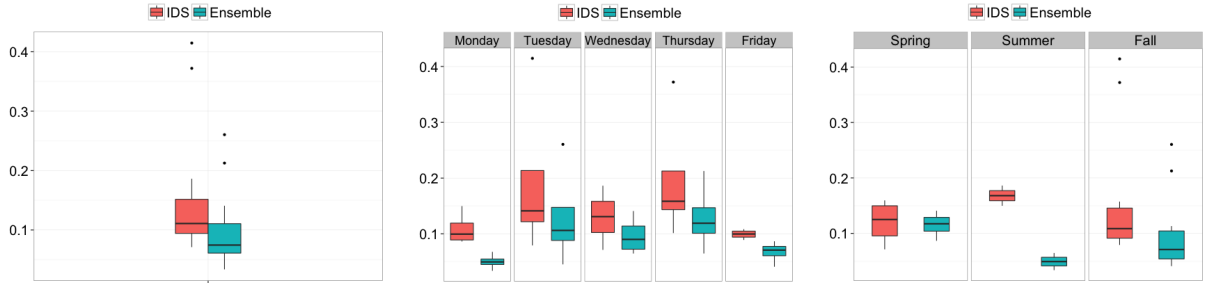

(a) For all test days (b) By day of week (c) By seasons

Figure 6: MAPE values for B28: campus center with meeting and event spaces and a grand ballroom

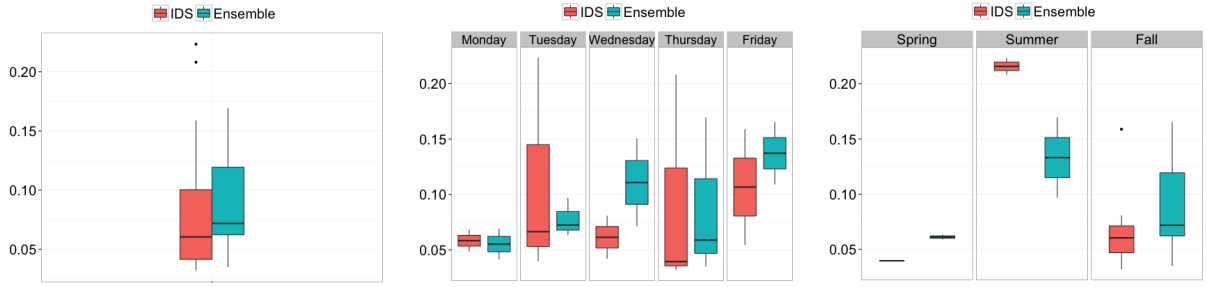

(a) For all test days (b) By day of week (c) By seasons

Figure 7: MAPE values for B14: an academic building with large proportion of classrooms and some faculty offices

Conclusion

We address the *reduced consumption prediction* problem in the energy sustainability domain that is relevant for successful implementation of dynamic demand response (DR) by the electric utilities. Standard models for electricity consumption prediction, such as the time series models, are unsuitable for this problem due to abrupt changes in consumption profile at the beginning and end of DR events. We propose a novel ensemble model to make predictions using pre-DR , in-DR, and all-day consumption sequences, to provide *superior performance*, achieving an average error of 13.5%, which is an improvement of 8.8% over averaging based baseline approach. With *low computational complexity*, our approach provides a practical solution that can be applied for real-time prediction. Also, our model provides a *simple and generalizable* approach allowing domain experts to integrate a variety of contextual attributes that could affect reduced electricity consumption. Our results indicate that the strength of our model is particularly relevant for: 1) *buildings for which electricity consumption does not follow a strict schedule* (i.e., absence of periodic activities), and 2) *buildings with less historical DR data*. We believe that our results and insights set the foundation for future modeling and practice of DR programs in the smart grid domain.

Acknowledgments

This material is based upon work supported by the United States Department of Energy under Award Number DEOE0000192, and the Los Angeles Department of Water and Power (LA DWP). The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, the LA DWP, nor any of their employees.

References

Alzate, C., and Sinn, M. 2013. Improved electricity load forecasting via kernel spectral clustering of smart meters. In *IEEE International Conference on Data Mining*.

Aman, S.; Frincu, M.; Chelmis, C.; Noor, M.; Simmhan, Y.; and Prasanna, V. 2015. Prediction models for dynamic demand response: Requirements, challenges, and insights. In *IEEE International Conference on Smart Grid Communications*.

Aman, S.; Simmhan, Y.; and Prasanna, V. 2015. Holistic measures for evaluating prediction models in smart grids. *IEEE Transactions in Knowledge and Data Engineering* 27(2).

Breiman, L. 2001. Random forests. *Machine learning* 45(1):5–32.

Chelmis, C.; Aman, S.; Saeed, M. R.; Frincu, M.; and Prasanna, V. K. 2015. Predicting reduced electricity consumption during dynamic demand response. In *AAAI Workshop on Computational Sustainability*.

Coughlin, K.; Piette, M. A.; Goldman, C.; and Kiliccote, S. 2009. Statistical analysis of baseline load models for non-residential buildings. *Energy and Buildings* 41(4).

Dudani, S. A. 1976. The distance-weighted k-nearest-neighbor rule. *IEEE Transactions on Systems, Man and Cybernetics* 4.

FERC. 2010. Assessment of demand response and advanced metering. *Federal Energy Regulatory Commission*.

Liaw, A., and Wiener, M. 2002. Classification and regression by randomforest. *R News* 2(3):18–22.

Martinez Alvarez, F.; Troncoso, A.; Riquelme, J. C.; and Aguilar Ruiz, J. S. 2011. Energy time series forecasting based on pattern sequence similarity. *Knowledge and Data Engineering, IEEE Transactions on* 23(8):1230–1243.

Mathieu, J. L.; Callaway, D. S.; and Kiliccote, S. 2011. Variability in Automated Responses of Commercial Buildings and Industrial Facilities to Dynamic Electricity Prices. *Energy and Buildings*.

Park, S.; Ryu, S.; Choi, Y.; and Kim, H. 2014. A framework for baseline load estimation in demand response: Data mining approach. In *IEEE International Conference on Smart Grid Communications*.

Piette, M. A.; Kiliccote, S.; and Dudley, J. H. 2012. Field Demonstration of Automated Demand Response for Both Winter and Summer Events in Large Buildings in the Pacific Northwest. *Lawrence Berkeley National Lab Technical Report, LBNL-6216E*.

UNDP. 2010. Promoting energy efficiency in buildings: Lessons learned from international experience. Release Date: 2010.

Wijaya, T. K.; Vasirani, M.; and Aberer, K. 2014. When bias matters: An economic assessment of demand response baselines for residential customers. *IEEE Transactions on Smart Grid* 5(4).

Ziekow, H.; Goebel, C.; Struker, J.; and Jacobsen, H.-A. 2013. The potential of smart home sensors in forecasting household electricity demand. In *IEEE International Conference on Smart Grid Communications*.