

Learning to REDUCE: A reduced Electricity Consumption Prediction Ensemble

S. Aman, C. Chelmis, and V. K. Prasanna

University of Southern California Los Angeles, CA

AAAI Workshop - AI for Smart Grids and Smart Buildings



Demand Response (DR)



- Demand Response is used in smart grids to make the demand adaptive to supply conditions.
- Customers respond to signals from the utility to reduce their consumption during peak periods as per prior agreements.





Prediction of Reduced Consumption



Reduced consumption prediction is useful in following decision-making tasks:

estimating potential reduction during DR (Chelmis et. al., 2015)

intelligently targeting customers for participation in DR (Ziekow et. al., 2015)

performing dynamic DR at a few hours' notice (Aman et. al., 2015) estimating the amount of incentives to be given to the customers (Wijaya et. al., 2014)



Characteristics and Challenges



	Normal Consumption	DR Baseline	Reduced Consumption
Goal	Planning, DR	Curtailment calculation	Planning, DR, dynamic DR
Timing	Outside the DR event	Outside the DR event	During the DR event
Historical data	Readily available	Readily available	Sparse or non-existent
Compute requirements	Offline or real-time	Offline	Real-time for dynamic DR
Profile changes	Gradual	Gradual	Abrupt at DR event boundaries
Prior Work	Several	Several	None

We are the first to address this problem using data from DR experiments done on the USC campus. (Aman et. al., 2016), (Chelmis et. al., 2015)



Key Challenges



- Unavailability of reduced consumption data
- Cancellation of DR event when found violating thermal comfort limits of occupants.
- Reduced consumption is affected by several factors
 - time of day/ day of week
 - reduction strategy
 - human behavior
 - external/environmental factors, e.g., temperature
- Time series models that work well for normal consumption prediction are ineffective for reduced consumption prediction, due to
 - **abrupt changes** in consumption profile at the beginning and end of the DR event
 - **insufficient recent observations** within the DR window for a time series model to be trained reliably



Hypothesis

Historical data from the past DR events can be used as predictors for reduced consumption.



Consumption Sequences





$$\begin{array}{ll} e_{i,j} & - \text{Electricity consumed on day } i \text{ in interval } j \\ \mathcal{E}_{i,s,l} & - \text{Subsequence of daily sequence } \mathcal{E}_i \text{ starting at } s \text{ of length } l \\ L & - \text{Length of the DR interval} \\ d & - \text{The interval when DR begins} \\ J & - \text{Number of intervals in a day} \end{array}$$

$$d > 1$$
$$d + L - 1 \le J$$

Contextual Attributes



- **Time Series attributes:** vary over intervals - temperature, dynamic pricing, occupancy, etc.
- **Static attributes:** same for all intervals day of week, holiday, etc.

Daily Context $C_i = \langle A_i[1], ..., A_i[N_t], B_i[1], ..., B_i[N_s] \rangle$ $A_i[k] = \{a_{i,1}, a_{i,2}, ..., a_{i,J}\}$

Pre-DR Context

$$C_{i,1,d-1} = \langle A_i[1], ..., A_i[N_t], B_i[1], ..., B_i[N_s] \rangle$$

 $A_i[k] = \{a_{i,1}, a_{i,2}, ..., a_{i,d-1}\}$

Correspond to the Daily Sequence and Pre-DR Sequence defined previously.







REDUCE – <u>Redu</u>ced <u>Consumption Ensemble</u>



- $[\hat{\mathcal{E}}_{\epsilon,d,L}]_m$ In-DR sequence predicted by model *m* on day ϵ
- Ensemble Models combine base models that model different behaviors, for e.g., **mean behavior**, **context dependent behavior**, etc.
- Random Forest Models are found to perform better than a single regression tree (Breiman, 2001)

IDS – In-DR Sequence Model



- Models "mean behavior"
- Similar to the averaging approach used by the utilities/ISOs to calculate the DR baseline.
- While utilities average over **similar** non-DR days, IDS averages over **all** DR days.
- Advantages:
 - Low computation cost suitable for real-time predictions
 - Uni-variate model low data collection cost
- Predicted sequence is given by:

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

 ${\cal E}$ is the set of historical DR days



PDS – Pre-DR Sequence Similarity Models



If two DR days have similar pre-DR sequences, their in-DR sequences would be similar.

- Pre-DR sequence Used to select similar DR days
- Pre-DR context
- Similarity is calculated by:

 $SimScore(\epsilon, i) = sim(\langle \mathcal{E}_{\epsilon, 1, d-1}, \mathcal{C}_{\epsilon, 1, d-1} \rangle, \langle \mathcal{E}_{i, 1, d-1}, \mathcal{C}_{i, 1, d-1} \rangle)$

- Selected days are sorted based on decreasing similarity and weighed accordingly.
- Predicted sequence is given by:

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

 \mathcal{E} is the set of historical DR days ω_{ϵ} is the weight on day ϵ

PDS models context dependent behavior





DSS – Daily Sequence Similarity Models

• Daily sequence

Used to discover clusters of daily profiles

- Daily context
- Form daily profiles for each day

$$\mathcal{P}_{\epsilon} = \langle \mathcal{E}_{\epsilon}, \mathcal{C}_{\epsilon} \rangle$$

- Cluster daily profiles and let c_m be the centroid of each cluster
- Probability of a given DR day belonging to a cluster is given by:

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

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

• Predicted sequence is given by:

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



Experiments



- Reduced Consumption Data was collected from 952 DR events (2012-2014) on 32 buildings at USC campus
- Data granularity: 15 minutes (J = 96 intervals per day)
- DR event duration: 1 PM to 5 PM (L = 16 intervals)
- One time series attribute: Temperature
- Seven static attributes: Day of week



Distribution of DR events across buildings



Results – MAPE (1)



🛱 IDS 🛱 Ensemble





Results – MAPE (2)





- 952 DR events (2012 2014)
- 32 USC buildings
- Contextual attributes:
 - temperature (NOAA)
 - day of week

- REDUCE outperforms the baseline IDS for about 70% of the buildings
- It also limits prediction error to <10% for over half the buildings
 considered highly reliable by domain experts (Aman et. al., 2015)
- Overall average error is 13.5%, an improvement of 8.8% over the baseline



Results – Effect of Schedule (1)



We examine two types of buildings:

- Schedule-driven: consisting primarily of classrooms (activities governed by schedules)
- Non-schedule driven: Few or no classrooms

Case Study:

- B21 Non-scheduled:
 - a building with large computer labs, and faculty and graduate student offices
- B28 Non-scheduled:

a campus center building with large meeting spaces, and a grand ballroom with seating for over 1000 people

• B14 – Scheduled:

an academic building with classrooms and few office spaces.

Results:

- For non-scheduled buildings, REDUCE gives superior performance
- For scheduled buildings, IDS performs well due to the presence of repetitive human activity coupled to class schedules

• Corollary: REDUCE would perform better for residential buildings (with non-scheduled activities).



Results – Effect of Schedule (2)







Results – Effect of Schedule (3)







Results – Effect of Schedule (4)





For scheduled buildings, IDS performs well due to the presence of repetitive human activity coupled to class schedules





Results – Effect of Training Data Size



- The performance of REDUCE is not sensitive to the training data size.
- 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.



Results – Effect of Variance in Consumption



🔶 IDS 📥 Ensemble

- Prediction error decreases with increasing average consumption for REDUCE model.
- 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.

Conclusion



We propose a novel ensemble model for reduced consumption prediction

- Achieve an average error of 13.5%, (an improvement of 8.8% over baseline)
- Low computational complexity
- Practical solution for real-time prediction
- Allows domain experts to integrate a variety of contextual attributes

Our proposed model is particularly relevant for:

- buildings for which electricity consumption does not follow a strict schedule (i.e., absence of periodic activities)
- buildings with less historical DR data.

Our results set the foundation for future modeling and practice of DR programs in smart grids.



References



Aman, S.; Chelmis, C.; and Prasanna, V. K. 2016. Learning to REDUCE: A Reduced Consumption Prediction Ensemble. In AAAI Workshop on AI for Smart Grids and Smart Buildings.

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.

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.

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.





Thank You!

