



9-2015

Sometimes, Money Does Grow on Trees: DR-Advisor, A Data Driven Demand Response Recommender System


Madhur Behl

University of Pennsylvania, mbehl@seas.upenn.edu

Rahul Mangharam

University of Pennsylvania, rahulm@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/mlab_papers

 Part of the [Computer Engineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Madhur Behl and Rahul Mangharam, "Sometimes, Money Does Grow on Trees: DR-Advisor, A Data Driven Demand Response Recommender System", *SRC TECHCON 2015*. September 2015.

Publication ID:P084437

Best Paper Award in the Internet of Things Session

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/mlab_papers/81
For more information, please contact repository@pobox.upenn.edu.

Sometimes, Money Does Grow on Trees: DR-Advisor, A Data Driven Demand Response Recommender System

Abstract

Unprecedented amounts of information from millions of smart meters and thermostats installed in recent years has left the door open for better understanding, analyzing and using the insights that data can provide, about the power consumption patterns of a building. The challenge with using data-driven approaches, is to close the loop for near real-time control and decision making in large buildings. Furthermore, providing a technological solution alone is not enough, the solution must also be human centric. We consider the problem of end-user demand response for commercial buildings. Using historical data from the building, we build a family of regression trees based models for predicting the power consumption of the building in real-time. We have built DR-Advisor, a recommender system for the building's facilities manager, which provides optimal control actions to meet the required load curtailment while maintaining building operations and maximizing the economic reward.

Keywords

demand response, machine learning, data, CPS, buildings

Disciplines

Computer Engineering | Electrical and Computer Engineering

Comments

Publication ID:P084437

Best Paper Award in the Internet of Things Session

Sometimes, Money Does Grow on Trees: DR-Advisor, A Data Driven Demand Response Recommender System

Madhur Behl and Rahul Mangharam
Dept. of Electrical and Systems Engineering
University of Pennsylvania
{mbehl, rahulm}@seas.upenn.edu

Abstract—Unprecedented amounts of information from millions of smart meters and thermostats installed in recent years has left the door open for better understanding, analyzing and using the insights that data can provide, about the power consumption patterns of a building. The challenge with using data-driven approaches, is to close the loop for near real-time control and decision making in large buildings. Furthermore, providing a technological solution alone is not enough, the solution must also be human centric. We consider the problem of end-user demand response for commercial buildings. Using historical data from the building, we build a family of regression trees based models for predicting the power consumption of the building in real-time. We have built DR-Advisor, a recommender system for the building’s facilities manager, which provides optimal control actions to meet the required load curtailment while maintaining building operations and maximizing the economic reward.

I. INTRODUCTION

The organized electricity markets in the United States all use some variant of real-time locational marginal price for wholesale electricity. For e.g., PJMs real-time market is a spot market where electricity prices are calculated at five-minute intervals based on the grid operating conditions.

Electricity costs are the single largest component of a large commercial and industrial (C&I) building’s operating budget. For such large consumers, buying and reacting to real-time electricity prices isn’t as simple as paying a flat-rate monthly bill. Their power consumption demands are sensitive to weather conditions and may result in peaks on an extremely hot or an extremely cold day. These peaks are not only operationally inefficient but also extremely expensive. Such customers are increasingly looking to demand response (DR) programs to help manage their electric utility costs. DR programs involve a voluntary response of a building to real-time price signals. In such programs, end-users reduce their electricity load during periods of high prices and receive a financial reward for their load curtailment.

However, to take advantage of real-time pricing and DR programs, the C&I consumers must monitor electricity prices and be flexible in the ways they choose to use electricity. The challenge for large buildings is to be able to predict their aggregate power consumption accurately and at a fast time scale in order to take suitable load curtailment control actions.

There are three big barriers to successfully enabling real-time building electricity prediction and demand response: (a) Each building is designed and used in a different way and therefore, it has to be uniquely modeled. Learning high fidelity models of buildings using non data-driven approaches is very cost and time prohibitive and requires retrofitting the building with several expensive sensors. (b) Secondly, the volatility




Fig. 1: DR-Advisor Architecture

and variance in real-time electricity rates poses a risk for large buildings. They need the capability to respond to the price volatility in a fast and reliable manner. Fig 2 shows an example of the volatility in real-time pricing from the New-England ISO [9]. The nominal price of electricity is \$27.34 but increases to \$672.41 on an extremely hot day in July 2013. (c) Complex models go through a long calculation routine and involve too many factors. It is not easy for a human engineer and a buildings manager to judge if the operation/decision is correct or not or how it was generated in the first place. Therefore, the required solution must be transparent, human centric and highly interpretable.

In 2013, a report by the National Climate Assessment provided evidence that the most recent decade was the nations warmest on record [8] and experts predict that temperatures are only going to rise. Heat waves in summer and polar vortexes in winter are growing longer and pose increasing challenges to an already over-stressed electric grid. With the increasing penetration of renewable generation, the grid is experiencing a shift from predictable and dispatchable electricity generation to variable and non-dispatchable generation. This adds another level of uncertainty and volatility to the electricity grid. Demand response and real-time electricity pricing are considered as an agreed upon means of mitigating the uncertainty of renewable generation and improving the systems efficiency with respect to economic and environmental metrics.

Across the United States, electric utilities and independent system operators (ISOs) are devoting increasing attention and resources to demand response (DR) [6]. Potential peak reduction from demand response markets in U.S. increased by 2,451 MW or 9.3 percent to a total of 28,503 MW from 2012 to 2013 [3]. The revenue to end-users from DR markets




Fig. 2: Volatility in real time prices for New-England ISO [9]. Nominal price: (\$27.34) Peak price: (\$672.41)

with PJM alone is about \$700 million [7]. A recent report [10] estimates that the global C&I DR revenue is expected to reach nearly \$40 billion from 2014 through 2023.

We have built an open-source tool called DR-Advisor (Demand Response-Advisor), which acts as a recommender system for the building's facilities manager and provides building power consumption prediction and control actions for meeting the required load curtailment and maximizing the economic reward. Using historical meter and weather data along with set-point and schedule information, DR-Advisor builds a family of regression trees to learn non-parametric data-driven models for predicting the power consumption of the building (Figure 1). These models can be used for real-time demand response baseline prediction, strategy evaluation and control synthesis, without having to learn complex models of the building. This work has the following contributions:

- 1) We demonstrate the benefit of using regression trees based approaches for estimating the DR baseline power consumption and for evaluating pre-determined DR strategies in real-time. The use of tree based models to address real-time demand response problems is novel.
- 2) The biggest contribution of this work is the fusion of tree based models into DR-Advisor, a simple and highly interpretable open source tool for making demand-response recommendations. It eliminates the cost of time and effort required to build and tune high fidelity models of buildings for DR.
- 3) We evaluate and compare the performance of DR-Advisor on the Department of Energy's (DoE) large commercial reference building.

II. PROBLEM DEFINITION

The time line of a typical DR event consists of three periods (Fig. 5). The main period during which the demand needs to be curtailed is the *sustained response period*. The start of this period, i.e., the time by which the target reduction must be achieved, is the *reduction deadline*. Prior to that deadline, an *event notification* will be issued, at the notification time. The end of the sustained response period – the *release time* – is when the main curtailment is released.

We focus on three challenging problems of demand response

- 1) **DR baseline prediction:** A baseline is an estimate of the electricity that would have been consumed by a customer in the absence of a demand response event. The baseline is the primary tool for measuring curtailment during a DR event which determines the financial payback which




Fig. 3: Example of a demand response timeline.

the customer receives. For this reason the measurement and verification of demand response is the most critical component of any DR program. As shown in Figure 5, actual meter data is compared with the baseline demand to determine the curtailment achieved by the customer.

- 2) **Real time DR strategy evaluation** aims to answer the question of *how can we choose good DR strategies from a pre-determined set of strategies, in real time*. A DR strategy refers to what control actions, and at what times, a system (lighting, HVAC or plug loads) will actuate. In Figure 5, at the event notification time there are N different strategies available to choose from. DR-Advisor predicts the power consumption of the building due to each strategy at every time-step and chooses the best DR strategy.
- 3) **Real time DR strategy synthesis:** The next major challenge with real time demand response for the participant is to figure out what control action to take in the first place i.e. *how to come up with DR strategies ?*, which are suitable for the DR event based on the current state of the building and the weather outside.

In all the the three challenges, a recurring theme is the capability to successfully predict the power consumption of a large building in real-time. Any data-driven method which intends to solve these problems must have the capability to predict the building power consumption of the building under different circumstances and due to different control actions.

III. DR-ADVISOR: DATA-DRIVEN DEMAND RESPONSE

Regression trees are decision trees which predict responses to data. Regression trees belong to the class of recursive partitioning algorithms. At each node of the tree, we check the value of one of the inputs (or features) X_i and depending of the (binary) answer we continue to the left or to the right subbranch. When we reach a leaf we will obtain the prediction of the response Y . The seminal algorithm for learning regression trees from data is the CART algorithm as described in [2]. Contrary to linear or polynomial regression which are global models (the predictive formula is supposed to hold in the entire data space), trees try to partition the data space into small enough parts where we can apply a simple different model on each part. They are conceptually simple yet powerful. Regression trees offer several advantages in addition to being simple, which make them suitable for solving the challenges of demand response and building modeling. We list some of these advantages here:

- 1) Trees require very low computation power, both running time and storage requirements.




Fig. 4: DR-Advisor graphical user interface

- 2) Trees can easily handle the case where the data has lots of features which interact in complicated and nonlinear ways. the predictor variables themselves can be of any combination of continuous, discrete and categorical variables.
- 3) Sometimes, data has missing predictor values in some or all of the predictor variables. This is especially true for buildings, where sensor data streams fail frequently due to faulty sensors or faulty communication links. By design, regression trees can handle missing data better than most algorithms through the use of surrogate variables.
- 4) Tree based models are generally not affected by outliers but regression based models are.
- 5) Trees are highly interpretable algorithms. Complex building models go through a long calculation routine and involve too many factors. It is not easy for a human engineer to judge if the operation/decision is correct or not or how it was generated in the first place. Trees only involve simple *if this then that* rules which are very easy to understand.

A. Ensemble Methods

The problem with trees is their high variance and that they can over fit the data. It is the price to be paid for estimating a simple, tree-based structure from the data. While pruning and cross validation can help reduce over fitting, in DR-Advisor, we use ensemble methods for growing more stable trees. Two families of ensemble methods are usually distinguished: (a) In averaging methods, the driving principle is to build several estimators independently and then to average their predictions. On average, the combined estimator is usually better than any of the single base estimator because its variance is reduced. (b) By contrast, in boosting methods, base estimators are built sequentially and one tries to reduce the bias of the combined estimator. The motivation is to combine several weak models to produce a powerful ensemble. The DR-advisor tool used a combination of cross validated trees, random forest and boosted regression trees as the underlying ensemble methods. For a more comprehensive review we refer the reader to [1].

IV. CASE STUDY

The building under consideration is the DOE Commercial Reference Building simulated in EnergyPlus [5] This virtual test-bed is a large 12 story office building consisting of 73




Fig. 5: Comparison between the actual power consumption and the baseline prediction for July 17th, 2013; a peak load contribution day. The DR event is from 1700-1800 hrs.

zones with a total area of 500,000 sq ft. There are up to 2,397 people in the building during peak occupancy. The building has 2 electric water-cooled chillers, variable air volume (VAV) supply air terminals with reheat and plenum zones and a single gas based boiler. During peak load conditions the building can consume up to 1.6 MW of power. EnergyPlus provides typical meteorological year data files for many sites which are generated as averages of different weather characteristics across the past 15-30 years. However, for the purposes of the simulation we use Actual Meteorological Year (AMY) data from Chicago for the years 2012 and 2013.

The data that we use can be divided into three different categories as described below:

- 1) Weather data, which includes measurements of the dry bulb temperature, wet bulb temperature, relative humidity and wind conditions.
 - 2) Schedule data, which includes fixed temperature set-points schedules of chilled water supply, supply air temperature and zone air temperature on the HVAC side and lighting schedules.
 - 3) Building data, which includes the measurements of zone temperature, lighting, supply air and water temperatures, power consumption etc.
- In addition to these data sets we also train on engineered features like the time of day and the day of week.

A. DR Baseline

On July 17, 2013 a demand response event occurred across the PJM ISO from 1700 hrs to 1800 hrs. We estimate the baseline power consumption of the office building for the DR event for July 17, 2013. The result of this comparison is shown in Figure 5. The lowest root mean square error obtained in this case is only 12 kW on an average consumption of 0.62 MW, which corresponds to a normalised root mean square (NRMSE) of only 2.01%. Using the ensemble methods, the DR-Advisor is able to accurately predict the baseline consumption of the building using just weather and schedule data, which require little to no sensor installations at the building site.

B. DR Strategy Evaluation

As stated earlier, the challenge is DR strategy evaluation is to predict the power consumption profile of the building in




Fig. 6: Left:Rule based demand response temperature set-point reset strategy executed for July 17, 2013. Right: Comparison between the actual power consumption and the predicted power for July 17th, 2013; There is a DR event from 1700-1800 hrs.

real-time due to a fixed policy. Upon receiving the notification of the DR event at 1600 hrs, the zone air temperature set-point for all the zones is increased from a nominal value of 24°C by 2° to 26°C. The chilled water supply temperature set-point is increased from 6.7°C by 1.5° to 7.2°C. At the beginning of the event at 1700 hrs, the zone air temperature set-point is further increased by 2° and the chilled water supply temperature set-point is increased by another 1.5°. This fixed, rule-based strategy is shown in Figure 6(left). The predicted response of the building compared to the actual response due to the fixed strategy is shown in Figure 6(right). We obtain an error of 6.23% for predicting the power consumption of the building in real time during a demand response event.

C. Real-Time DR Strategy Synthesis

Figure 7 shows the power consumption profile of the building using DR-Advisor for the DR event. We can see that using DR-Advisor we are able to achieve a sustained curtailed response of $\sim 300\text{kW}$ over a period of 1 hour as compared to the baseline power consumption estimate.

1) **Revenue from DR:** We use Con Edison utility company's commercial demand response tariff structure [4] to estimate the financial reward obtained due to the curtailment achieved by the DR-Advisor for our Chicago based DoE commercial reference building. The utility provides a \$25/kWh per month as a reservation incentive to participate in the real-time DR program for summer. In addition to that, a payment of \$1 per kWh of energy curtailed is also paid. For our virtual test-bed, the peak load curtailed is 331kW and a total of




Fig. 7: Real-Time DR synthesis using the mbCRT algorithm for July 17, 2013. A curtailment of 300kW is sustained during the DR event period.

327.4kWh of energy was saved. Considering ~ 5 such events per month for 4 months, this amounts to a revenue of \$39,700 for participating in real-time DR only for summer. This is a significant amount, especially since using DR-Advisor does not require an investment in complex modeling or sensor retrofits for a building.

V. CONCLUSION

DR-Advisor, a data-driven open source tool has been presented. We show how regression tree based methods provide an excellent way to predict the power consumption response of a large commercial building while being simple and interpretable. The use of regression trees based methods to address problem of real-time demand response for large scale buildings is novel. DR-Advisor achieves a prediction accuracy of **94-97%** for DR baseline and DR strategy evaluation. DR-Advisor can achieve a sustained curtailment of **300kW** during a DR event. Using a DR pricing structure from Con Edison utility, we estimate a revenue of \sim \$40,000 for the DoE reference building over one summer. The biggest advantage of DR-Advisor is that it completely bypasses cost and time prohibitive process of building high fidelity models of buildings.

ACKNOWLEDGEMENTS

This work was supported by STARnet, a Semiconductor Research Corporation program, sponsored by MARCO and DARPA.

REFERENCES

- [1] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [2] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. *Classification and regression trees*. CRC press, 1984.
- [3] Federal Energy Regulatory Commission et al. Assessment of demand response and advanced metering. 2008.
- [4] Con Edison. Demand response programs details.
- [5] M. Deru, K. Field, D. Studer, et al. U.s. department of energy commercial reference building models of the national building stock. 2010.
- [6] Charles Goldman. Coordination of energy efficiency and demand response. *Lawrence Berkeley National Laboratory*, 2010.
- [7] PJM Interconnection. 2014 demand response operations markets activity report. 2014.
- [8] Jerry M Melillo, TC Richmond, and Gary W Yohe. Climate change impacts in the united states: the third national climate assessment. *US Global change research program*, 841, 2014.
- [9] New-England ISO. Real-time maps and charts, archives., 2013.
- [10] Navigant Research. Demand response for commercial & industrial markets market players and dynamics, key technologies, competitive overview, and global market forecasts. 2015.