Title

Predicting success of energy savings interventions and industry type using smart meter and retrofit data from thousands of non-residential buildings

Author

Miller, Clayton

Publication Date

2017-11-01

Predicting success of energy savings interventions and industry type using smart meter and retrofit data from thousands of non-residential buildings

Clayton Miller
Building and Urban Data Science (BUDS) Group, Department of Building, National University of Singapore (NUS)
Singapore
Institute of Technology in Architecture (ITA), Department of Architecture, ETH Zürich
Zürich, Switzerland
clayton@nus.edu.sg

ABSTRACT

This paper discusses the creation of targeting and segmentation information about non-residential buildings that are equipped with advanced metering infrastructure (AMI) meters, or smart meters. Statistics, model, and pattern-based temporal features are extracted from over 36,000 smart meters. They are then merged with a database of past energy efficiency interventions such as lighting, HVAC, and controls retrofits from 1,600 buildings. The buildings are divided into *Good, Average, and Poor* performing classes according to consumption from before and after the retrofits. Classification models are developed that improve the ability to predict retrofit success and standard industry class by 18.3% and 27.6% respectively over baselines. This study serves as an example of better leveraging smart meter data from non-residential buildings for utility targeted incentive programs. The methodology outlined is preliminary and further models and temporal features are to be tested.

CCS CONCEPTS

• **Information systems** → **Data analytics**; • **General and reference** → *Empirical studies*; • **Applied computing** → *Multi-criterion optimization and decision-making*;

KEYWORDS

Non-residential buildings, Smart meters, Segmentation, Targeting, Retrofit analysis

ACM Reference Format:
Clayton Miller. 2017. Predicting success of energy savings interventions and industry type using smart meter and retrofit data from thousands of non-residential buildings. In *Proceedings of 4th ACM Conference on Systems for Energy-Efficient Built Environments (BuildSys'17).* ACM, New York, NY, USA, 4 pages. https://doi.org/https://doi.org/10.1145/3137133.3137160

BuildSys'17, November 8–9, 2017, Delft, The Netherlands

ACM ISBN 978-1-4503-5544-5/17/11…$15.00
https://doi.org/https://doi.org/10.1145/3137133.3137160

1 INTRODUCTION

As of 2015, there are 64.7 million smart meters in the United States, with 7.3 million of those installed in non-residential buildings[1]. Utilities are searching for ways to convert these data into insight that improve the operations of their distribution systems. They often help building owners improve the efficiency and demand response capabilities of their building through targeted energy audits, retrofit assistance, and improvement incentives. Commercial and industrial buildings are especially a challenge due to their heterogeneity and complex energy-consuming systems as compared to residential. A component of these programs is the ability to target the customers most likely to benefit from incentive programs. A part of this process is known as *segmentation* as it seeks to create groups of similar accounts.

This paper outlines a process performed on a smart meter data set of over 36,000 buildings. These data are aggregated by the Vermont Energy Investment Corporation (VEIC) on behalf of several electrical utilities. The first step in this process is to collect, clean and perform an extensive feature engineering process on the sub-hourly meter data itself. These features are then used to train two different classification models. The first model uses the temporal features from before and after an energy savings intervention to predict the potential success of the implementation. This step is done by combining the meter data from a subset of 1,600 buildings with historical project data from multiple years of energy audits and retrofit applications. The second model predicts the primary standard industry classification (SIC) one-digit code of the building. For commercial facilities, this piece of meta-data is important in understanding the main use of the building. Both of these objectives are valuable for a utility to better target future buildings for retrofit programs. Much of the work in this publication is part of the author's Ph.D. dissertation [8].

1.1 Related work

Previous work in classification of smart meter data focuses primarily the ability to predict demographics, appliance characteristics, and renewable energy integration potential from residential buildings [1, 2, 5]. Residential customer segmentation is a part of this effort [6]. One study used smart meter data for anomaly and behavior detection to uncover sub-optimal consumption [10]. So far,

[1] *How many smart meters are installed in the United States, and who has them?* - https://www.eia.gov/tools/faqs/faq.php?id=108&t=3

no large utility smart meter study has focused on non-residential buildings and their specific challenges.

Another significant body of work is in the area of building retrofit analysis and the prediction of energy intervention success. A large review of energy retrofit analysis toolkits and case studies identifies the status quo of analysis; typically only a single building or a small set of buildings is analyzed [7]. Another study illustrates the use of clustering to group buildings in a community into similar potential retrofit combinations using data-driven approaches [4]. The largest study found on retrofit implementation is the case study using the CityBES tool to predict savings for 540 buildings [3]. These studies do not utilize the combination of a large smart meter or building energy retrofit success data sets.

1.2 Points of departure

This paper makes two advances within the field of utility smart meter data analytics. The first is the focus on a large data set of non-residential buildings. This context is differentiated from the more typical target of residential buildings and customers. Commercial and industrial buildings are inherently more complex and systems-driven than residential, thus they are often neglected in much of the previous research.

The other point of departure is related to the segmentation of commercial buildings according to their potential for energy savings interventions. In the author's knowledge, this the first study of its kind to merge a significant retrofit implementation data set with smart meter data. The number of buildings analyzed from both the smart meter (36,000) and retrofit analysis (1,600) is significantly larger than any previous non-residential building study.

2 TEMPORAL FEATURE EXTRACTION

Temporal features are aggregations of the behavior exhibited in time-series data. They are characteristics that summarize sensor data to inform an analyst through visualization or to use as training data in a classification or regression model. Feature extraction is a step in the process of machine learning as form of dimensionality reduction of data. This process quantifies various qualitative behaviors. A set of 28 temporal features types are extracted in this process from the three categories of statistics, regression model, and pattern-based. Two example features are exhibited in detail along with a heat map illustrating the range of behavior over time. These 28 features are part of a larger temporal feature library for non-residential buildings that is explained further in the literature [8].

The first category of features, basic statistics-based metrics, are created that utilize the time-series data vector for various time ranges to obtain information using concept such as mean, median, maximum, minimum, range, variance, and standard deviation. These metrics are further combined into ratios that describe daily, weekly or monthly behavior. An example of a ratio-based feature is daily maximum versus minimum, or daily load ratio. Distribution descriptors such as skewness and kurtosis are implemented. Finally, metrics that quantify the relationship between the electricity consumption and outdoor weather conditions are created. The use of the Spearman rank order correlation coefficient is an example of such a metric. Monthly coefficients of this type are developed for each month from a time range of 2.5 years of smart meter data. Figure 1 shows a heat map of these data with the x-axis representing the time range and the y-axis representing rows of one dimensional color bands quantifying the coefficient. Each band is a building and has color range from *-1* (highly heating correlated) as red to *+1* (highly cooling correlated) as blue. They are sorted vertically with the most cooling-influenced buildings at the top and most heating-affected at the bottom. This figure shows the general range of behavior related to weather impact and summer cooling and winter heating phases are easily identified.

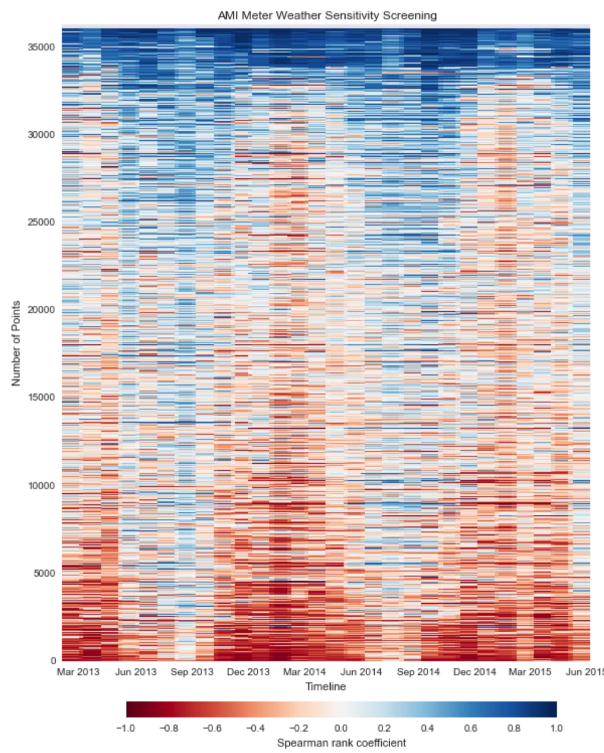

Figure 1: Heat map visual representation of the monthly Spearman rank order correlation coefficient for over 36,000 smart meters. This metric an example of a statistics-based feature that characterizes weather influence.

The secondary category is model-based features that are developed from multivariate, piece-wise regressions using daily consumption and outdoor air dry-bulb temperature. These models approximate the electricity consumption used for heating and cooling, the balance point temperature in which the building transitions from the various phases, and an approximate base-load. Some metrics are extracted from model-fit statistics to characterize appropriateness of the models.

The last category is the pattern-based features extracted from the daily, weekly, monthly and long-term shape and magnitude behavior of the buildings. The first type of pattern-based models is daily load patterns that are created using clustering techniques. Several features identify the percentage of time that a building's daily profiles occur within a particular typical pattern. Another key feature that is extracted quantifies the *volatility* of the consumption of buildings over a medium time range. Figure 2 illustrates an algorithm that detects breakouts, or large shift in steady-state consumption of the building. The breakouts in this figure are shown with a change in color on the heat map bands for each building. The bands are sorted top to bottom from the most to least volatile consumption patterns. The less volatile buildings are most likely from use types such as data centers or laboratories that have a consistent pattern of use over time. More unstable buildings could include schools and dormitories that have regular shifts in occupancy throughout the year.

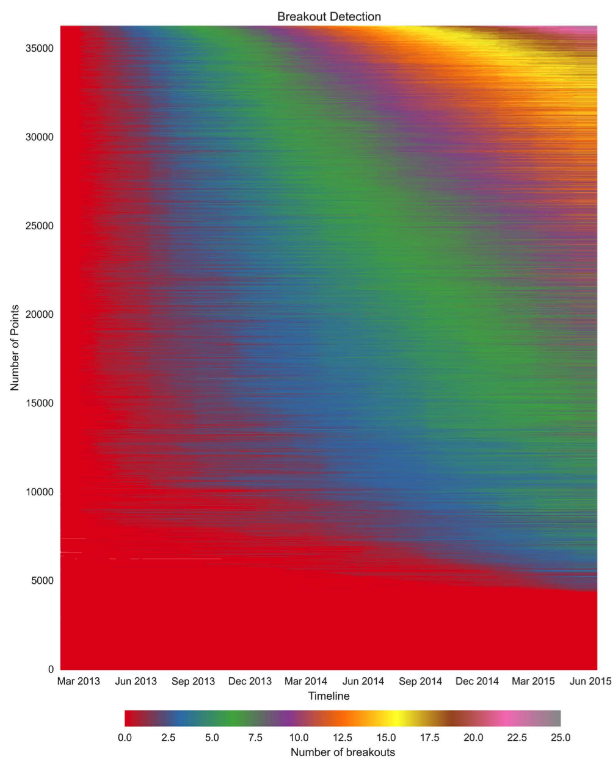

Figure 2: Heat map visual representation of the volatility feature implemented on approximately 36,000 smart meters across 2.5 years of data. This feature uses a breakout detection algorithm to detect pattern-based shifts of consumption.

3 PREDICTION MODELS

3.1 Energy efficiency measure implementation success prediction

The prediction of future energy savings measures implementations using the past data is implemented as a means of targeting potential projects for retrofits. For this proof-of-concept, data from before and after measure implementation are utilized from 1,600 buildings that had one or more actions implemented. The difference in mean daily consumption before and after the measure implementation is calculated to achieve a rough indication of measure success. The measures are divided into three classifications according to where the difference in daily consumption for each account fits in the range of values. In this analysis, the accounts in the lowest 33% are considered "Poor," while the 66% percentile are "Average" and the top 33% are considered "Good." Figure 3 illustrates a breakdown of the measure categories within the tested data set. These retrofits include energy savings interventions on the lighting, heating, cooling, ventilation, and compressed air systems.

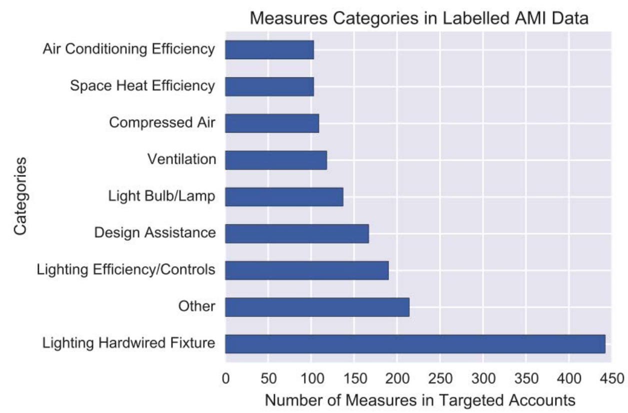

Figure 3: Breakdown of energy savings measure categories implemented on the 1,600 buildings from the retrofit data set

A Random Forest algorithm is implemented to use the temporal features to predict the class of potential measure success (*Good*, *Average* or *Poor*). The labeled data is split into 80% training and 20% test data. Figure 4 illustrates the classification error matrix for this model. The baseline model with these data can predict the success within this set of classification at 32.8% accuracy while the model based on the extracted features achieved 51.1% accuracy, an increase of 18.3%. An important aspect in this analysis is that the misclassification rate between *Good* and *Poor* is less than 20%.

Figure 4: Classification error matrix for prediction of measure implementation success using a random forest model

3.2 Predicting industry class membership

The next task is to characterize the general industry for which a building is being used. The primary usage of a building is a type of meta-data that is not always known in the case of smart meter data. It is used to supplement the energy savings targeting process. Temporal data are used to build a Random Forest classification model to predict the general use type of the building. In this case, the label for use type is the one digit Standard Industrial Class (SIC). This classification type is used because it is the primary tag used by the utilities to segment their customers. These labels include ten classifications from agriculture, finance, two types of manufacturing, mining, public administration, two types of services, transportation, and wholesale uses.

The baseline model correctly predicts the labels with an 18.1% accuracy, while the developed temporal features have an accuracy of 45.7%, an increase of 27.6%. The baseline model represents common practice in which a class is chosen based on the probability distribution of that type occurring in the labeled data set. The temporal feature set more than doubles the likelihood of predicting this piece of meta-data.

4 CONCLUSIONS

This paper illustrates preliminary work in combining and analyzing two key data sets related to performance in buildings: hourly smart meter data and energy savings intervention data. The goal of these data is to supplement a process of targeting buildings for energy conservation implementation measures. Utilization of temporal features is discussed in the context of assisting to label the approximate building use type and predicting measure success implementation. Tests using a classification algorithm showed an 18.3% increase in accuracy of predicting whether it would perform well with a set of performance interventions and a 27.6% increase in accuracy in predicting the industry type of a building.

4.1 Limitations

This work is in the early stage, and there are several fundamental limitations. The first is that only accuracy of classification prediction has been tested and released in this paper. Other classification metrics such as precision, recall, and the $F_{measure}$ should be evaluated and discussed in the targeting context. There is a strong incentive to explore further the ways of analyzing and utilizing classification models in this context to convert the results into real insight for utilities. Also, the author concedes that the overall accuracy of the classification model in this preliminary work is still low relative to what a good targeting process would need. Another limitation is in the way the energy savings implementations success is approximated. The data from before and after the retrofits are analyzed without normalization for weather, occupancy, and other factors that influence consumption beyond the effect of the energy savings technique implemented. This challenge could be overcome by merging weather or building management system data into the process to account for these factors.

4.2 Future work

The biggest opportunity ahead is to characterize missing meta-data and predict measure implementation success for future projects. Much work is also yet to be done to utilize more types of temporal features and input information to bring the overall prediction accuracies higher in absolute terms. Other prediction models from the extensive library found in the machine learning domain should be implemented and various approaches benchmarked against each other. Model prediction can also be improved incrementally as the temporal meter, and measures implementation data are better integrated. The building energy efficiency upgrades could also be analyzed separately to understand not only *if* a building is a good candidate for a retrofit, but also *which* retrofit is most appropriate.

Unfortunately, the exact data from this study is not available publicly due to privacy requirements. A fruitful future direction for this work would be in improving failsafe privacy measures for these type of data to be shared more freely for the benefit of the research community. Another study by the author uses similar techniques on a different, open data set of 507 buildings that are part of the Building Data Genome Project [9]. The open data from this project is available for download and collaboration on Github[2].

ACKNOWLEDGMENTS

The author would like to thank the anonymous reviewers and Semiha Ergan from NYU for their feedback. The author also thanks the collaborators from Vermont Energy Investment Corporation (VEIC), especially Greg Fanslow. This research was funded through an Institute of Technology in Architecture (ITA) Fellowship from the ETH Zürich.

REFERENCES

[1] Adrian Albert and Mehdi Maasoumy. 2016. Predictive segmentation of energy consumers. *Applied Energy* 177 (Sept. 2016), 435–448. https://doi.org/10.1016/j.apenergy.2016.05.128

[2] Adrian Albert and Ram Rajagopal. 2013. Smart meter driven segmentation: What your consumption says about you. *IEEE Transactions on power systems* 28, 4 (2013), 4019–4030. http://ieeexplore.ieee.org/abstract/document/6545387/

[3] Yixing Chen, Tianzhen Hong, and Mary Ann Piette. 2017. City-Scale Building Retrofit Analysis: A Case Study using CityBES. In *Building Simulation 2017*. San Francisco, CA, USA. http://citybes.lbl.gov/BS2017_CityBES_Paper_Final.pdf

[4] Philipp Geyer, Arno Schluter, and Sasha Cisar. 2016. Application of clustering for the development of retrofit strategies for large building stocks. *Advanced Engineering Informatics* (2016). http://www.sciencedirect.com/science/article/pii/S1474034616300167

[5] Srinivasan Iyengar, Stephen Lee, David Irwin, and Prashant Shenoy. 2016. Analyzing Energy Usage on a City-scale Using Utility Smart Meters. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys '16)*. ACM, New York, NY, USA, 51–60. https://doi.org/10.1145/2993422.2993425

[6] J. Kwac, J. Flora, and R. Rajagopal. 2014. Household Energy Consumption Segmentation Using Hourly Data. *IEEE Transactions on Smart Grid* 5, 1 (Jan. 2014), 420–430. https://doi.org/10.1109/TSG.2013.2278477

[7] S H Lee, Tianzhen Hong, Mary Ann Piette, and S C Taylor-Lange. 2015. Energy retrofit analysis toolkits for commercial buildings: A review. *Energy* (Jan. 2015). https://doi.org/10.1016/j.energy.2015.06.112

[8] Clayton Miller. 2017. *Screening Meter Data: Characterization of Temporal Energy Data from Large Groups of Non-Residential Buildings*. Ph.D. Dissertation. ETH Zurich, Zurich, Switzerland.

[9] Clayton Miller and Forrest Meggers. 2017. The Building Data Genome Project: An open, public data set from non-residential building electrical meters. *Energy Procedia* 122 (2017), 439 – 444. https://doi.org/10.1016/j.egypro.2017.07.400 {CISBAT} 2017 International Conference Future Buildings Districts and Energy Efficiency from Nano to Urban Scale.

[10] Kartik Palani, Nabeel Nasir, Vivek Chil Prakash, Amandeep Chugh, Rohit Gupta, and Krithi Ramamritham. 2014. Putting Smart Meters to Work: Beyond the Usual. In *Proceedings of the 5th International Conference on Future Energy Systems (e-Energy '14)*. ACM, New York, NY, USA, 237–238. https://doi.org/10.1145/2602044.2602084

[2] https://github.com/buds-lab/the-building-data-genome