Personalized energy efficiency program targeting with association rule mining

Pawel Zawadzki [1]*, Ye Lin, Frances Dahlquist, Tingwen Bao, Anne-Lise Laurain, Kevin Johnson*
Pacific Gas & Electric Company

ABSTRACT

Finding customers predisposed to energy efficiency (EE) offerings is one of the most challenging elements of EE marketing campaigns. As part of PG&E's efforts to improve EE targeting, the Data Analytics and Governance (DAG) team is employing data mining methodologies to develop personalized EE product and program recommendations for non-residential customers. In many industries, similar recommendation systems have proven to be a highly cost effective way of matching customers with product or service offerings and are now widely used by retailers such as Amazon and Walmart, entertainment providers such as Netflix or online advertising companies such as Google. To build such a recommendation system at PG&E, we used association rule mining to discover meaningful relationships between historical adoption of EE products and various customer characteristics such as machine-learnt cooling loads, industry codes, and participation in Peak Day Pricing (PDP). We have discovered about 10,000 statistically significant rules that form the basis for personalized EE recommendations for commercial customers and provide insights into EE adoption patterns.

Introduction

Customer segmentation is one of the most popular approaches to targeting (Linoff and Berry 2011). According to the 2015 Utility Marketing survey, out of 20 surveyed utilities, 15 use segmentation for residential customers and 11 for small and medium businesses (Hutson 2016). Customer segments are typically designed to separate the universe of applicable customers into smaller and increasingly similar subsets; a unique combination of product recommendations, messaging and outreach channel is used to communicate with each subset. While this approach is very appealing, dividing customers into unique segments is difficult as different products may require different segmentation schemas. For instance, an EE measure focused on efficient HVAC systems may be driven most heavily by the weather sensitivity of a customer's load and climate zone, while an EE measure focused on lighting may be driven by a different combination of factors such as NAICS code, operating hours and rate structure. Furthermore, customers often do not belong to mutually exclusive segments, but rather exhibit behavior along a continuum, making segmentations frequently arbitrary (Linoff and Berry 2011).

An alternative approach to customer targeting is personalization, where each customer receives unique recommendations based on specific characteristics. PG&E and other utility companies use this approach to target large commercial customers, who are often assigned dedicated sales representatives. These reps are often industry experts who are able to provide highly accurate and personalized energy efficiency recommendations. In order to be cost

[1] Corresponding Author: ppz1@pge.com

effective, the process of developing curated recommendations, however, needs to be limited to a relatively small number of large customers. As a consequence, smaller businesses are less likely to receive high quality personalized recommendations. This may contribute to the fact that small business customers only account for 5% of non-residential EE savings despite being responsible for 21% of non-residential load (PG&E 2015).

Algorithmic, data driven targeting has the potential to provide personalized recommendations and improve the effectiveness of targeting campaigns. Such customer to product matching systems are now widely used in many industries such as retail (Amazon and Walmart), entertainment providers such (Netflix and U-Verse) or online advertising (Google), and have proven to be a highly cost effective way of reaching out to a large number of diverse customers (Finger 2014). So effective in fact that several companies, including Netflix, have held open competitions aimed at improving recommendation algorithms effectiveness with prizes as high as 1 Million USD (Netflix 2009). In the energy industry however, recommender systems remain relatively unexplored. Ironically, electric and gas utilities are well positioned to build such systems, having access to large amounts of customer data including AMI data, billing data, EE adoption data, etc.

In this paper, the Data Analytics and Governance (DAG) team describes the development of the "EE Recommender" an algorithmic and data driven match making system that provides personalized EE recommendations for commercial customers. This recommender system uses association rule learning (Han, Kamber and Pei 2010) to discover EE adoption patterns, i.e. relationships between various customer characteristics and EE products. Adoption patterns are then applied to current customers. EE Recommender system provides highly tunable recommendations that can be used by program managers, marketing teams, and sales representatives to better target commercial customers, especially small and medium business (SMB) customers.

Requirements

An energy efficiency customer to product match making system needs to meet a number of requirements that are unique to the utility industry. As opposed to retailers or online advertising companies who seek to maximize number of items sold or profit, utilities in California and other states need to maximize energy efficiency savings; EE is the highest priority procurement resource in California (CPUC 2013). For this reason, energy efficiency recommendation system should be able to make recommendations not only to past EE adopters but also customers who have not yet adopted an EE measure. In addition, it should also drive adoption from different end uses in order to capture more available savings.

For the above reasons, two of the most popular approaches for recommender systems, collaborative filtering and content based recommendations (Jannach et al. 2010), may not be optimal for an energy efficiency recommendation system. In collaborative filtering, recommendations are made based on the similarities between users (or items). The underlying idea for this type of recommender system is that users with similar preferences (hence the word collaborative) will buy similar products. To capture customer preferences, collaborative filtering requires a large amount of customer interaction data whereas utility customers typically adopt at most a few EE measures. For customers who have not yet rated or purchased any products, recommendations cannot be made (the so-called cold start problem of collaborative filtering). A utility industry EE recommendation system, on the other hand, should be able to make recommendations to customers, despite the fact that they have not yet adopted any EE measures. Content based recommenders need much less data, but they only provide recommendations similar to products that a customer has already adopted. For instance, if a customer has adopted an exterior lighting product, a content based recommender will suggest a different exterior lighting product. This behavior is clearly undesirable for an EE recommendation system since suggesting only similar products neglects potential savings from different end uses.

Algorithm

To build a recommender system that meets the requirements outlined above, the DAG team developed a knowledge-based recommendation engine that uses association rule learning (Han, Kamber and Pei 2010) to discover EE program adoption patterns. Adoption patterns link various customer characteristics with EE products and take the following form: "IF customer characteristics THEN energy efficiency product". Recommendations are then made to customers who exhibit the IF characteristics, but have not yet adopted the product.

Association rule learning is a method of discovering interesting relationships between variables in large relational databases and has been widely used for shopping basket analysis (Han, Kamber and Pei 2010). For instance, Wal-Mart discovered that customers who purchase Barbie dolls have a 60% likelihood of also purchasing one of three types of candy bars (Palmeri 1997).

The strength of an association between X ⇒ Y can be characterized in multiple ways, the most popular being (Han, Kamber and Pei 2010):

- Rule support (*Supp*): the proportion of transactions that contain both X and Y
- Confidence (*Conf*): the proportion of the X containing transactions that also contain Y
  - For instance, if the proportion of transitions that contain bread, butter and both is 0.1, 0.2 and 0.05, respectively, then *Supp*(bread) = 0.1, *Supp*(butter) = 0.2, *Conf*(bread ⇒ butter) = *Supp*(bread, butter)/*Supp*(bread) = 0.5.
- Conviction (*Conv*): quantifies the frequency that the rule makes an incorrect prediction
  - For instance, the conviction for the above "bread ⇒ butter" example is Conv(bread ⇒ butter) = [1 – *Supp*(butter)] / [1 – *Conf*(bread ⇒ butter)] = 1.6 indicating that if there is no association between bread and butter, bread without butter should appear 60% more often.

Figure 1 shows a flow chart of the algorithm used to develop PG&E's EE Recommender. Two sets of customers are built: Set A is composed of 43,000 non-residential customers (unique service account ids) who have adopted one or more EE products between 2010 and 2015. Set B is composed of all non-residential customers within the PG&E territory (630,000 unique service account ids). Customer profiles are composed of a variety of characteristics including North American Industry Classification System (NAICS[2]) codes, meter location, climate zone, rate

[2] https://www.naics.com

schedule, behavioral characteristics, and electricity usage patterns. Behavioral characteristics include participation in various programs such as past EE adoption and Peak Day Pricing[3]. Electricity usage patterns include, among others, seasonality, day of week, peak hours/opening hours, baseline usage, and load factor.

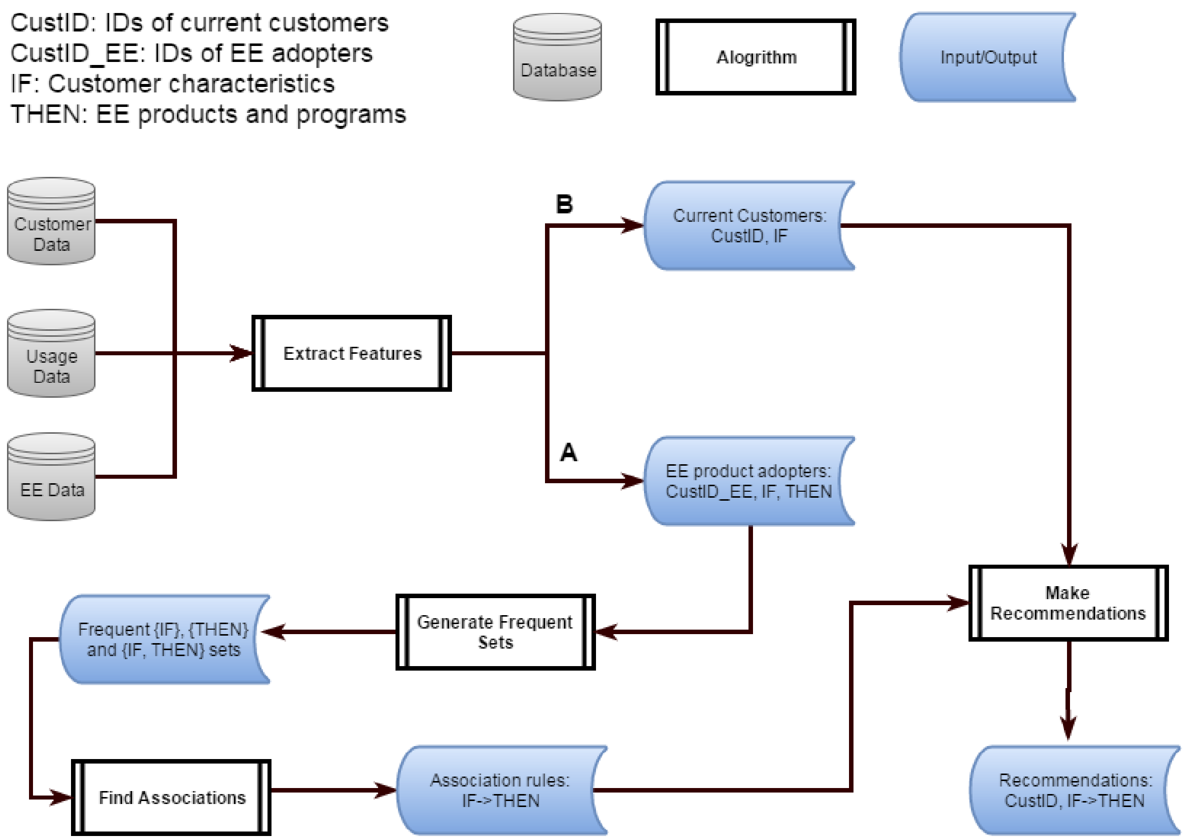

Figure 1- Schematic representation of the EE Recommender algorithm.

Profiles of EE adopters (set A) are then joined with the EE products that they have adopted (120,000 EE products in total). Because many historical products are no longer part of the EE portfolio we group them into 92 end use categories. For instance LED and CFL exterior lighting products are grouped into the exterior lighting category. In this way, EE recommendations will remain valid even when the product list changes and marketing teams or sale representative will be able to promote the most cost effective products.

Customer profiles (set A) matched with EE products groupings are then used to calculate frequent {IF}, {THEN} and {IF, THEN} sets: the {IF} set includes customer characteristics, the {THEN} set is a single element set that contains an EE product grouping, and {IF, THEN} is the combined set of customer characteristics and a product grouping. Thus constructed frequent sets are fed into the FP-growth frequent pattern mining algorithm which generates IF ⇒ THEN

[3] PG&E program that gives businesses a discount on regular summer electricity rates in exchange for higher prices on 9-15 Event Days per year.

association rules (Han et al. 2010). The resulting rules are then pruned to remove redundancy: if a rule 1 is a superset of another rule 2 having the same or higher confidence as rule 1, then rule 1does not provide additional knowledge and can be removed.

Figure 2 shows example association rules for walk-in cooler incentives. The most strongly associated rule (rank 1) links customers who are on secondary voltage, located in the climate zone Z04, operate 8 am 10 pm (day_late flag), active Mon-Sun (all_week flag), on the A10 rate schedule and belong to "Limited Service-Restaurants" industry code (NAICS code 722211). This adoption pattern has conviction of 3.27; meaning that if the association is purely by chance we should see 227% more incorrect predictions compared to what is contained in the EE adoption database.

In addition to product suggestion, EE Recommender also estimates potential total kWh and kW savings. Potential savings are calculated as median savings for customers with the IF characteristics who have adopted the product. Note that such estimated savings need to be adjusted for adoption rate which may vary for different communication channels, messaging, timing etc.

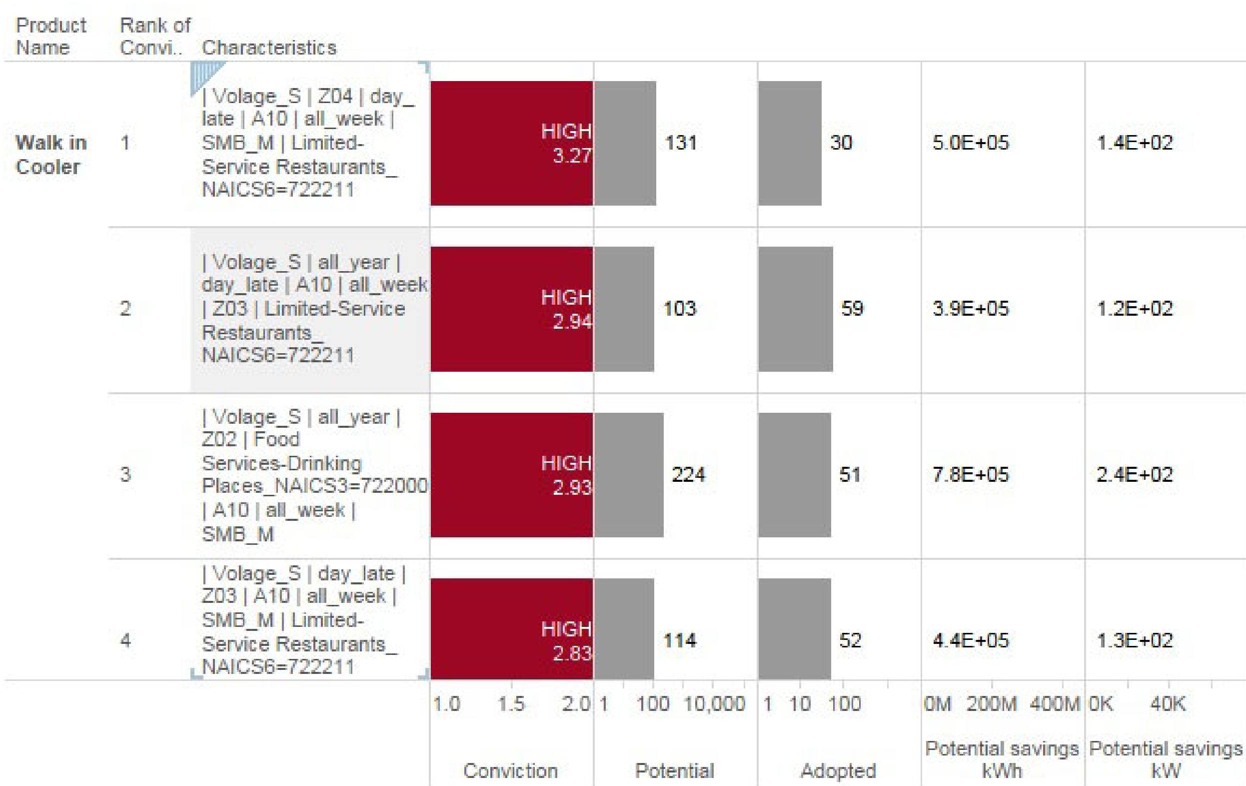

Figure 2- Example association rules for Walk-in Cooler Incentive (Each association rule also comes with estimated saving potential)

Note: Conviction: measures the association strength;

Potential: the number of customers who have not yet adopted the program;

Adopted: the number of customers who have done so

Potential savings: expected savings for all customers who have not yet adopted the program.

Validation

To evaluate the performance of PG&E's EE Recommender a train-test procedure was followed; during this procedure association rules are calculated using a train set and recommendations are generated for a test set[4]. The set recommendations is then compared with what the test customers actually did in the past. There are three possible outcomes of this comparison:

- Recommendations match what customers have adopted (True positives, TP);
- Customers have not adopted recommended products (False positives, FP);
- Products adopted by customers have not been recommended (False negatives, FN).

Using the above statistics we then calculate:

- Precision, P = TP / (TP + FP);
- Recall, R = TP / (TP + FN);
- F1 score: a combined metric equal to 2 P R / (P + R).

The precision essentially quantifies how many of the recommended items are relevant. Precision of 0.5 implies that half of the recommendations agree with what customers did. The recall, on the other hand, quantifies how many of the relevant items are recommended. Recall of 0.5 implies half of the products customers have adopted are also recommended by EE Recommender. To measure the overall performance of our match making system we calculate a combined metric: the F1 score and monitor the maximum value. For the current algorithm, the maximum F1 score is 0.36 at conviction of 1.4.

Figure 3 shows precision, recall and F1 score as a function of conviction cutoff where rules with conviction greater than the cutoff are retained for recommendations. Not surprisingly, higher conviction leads to higher precision as highly statistically significant recommendations are more likely to be accurate ones. Recall, on the other hand decreases with the conviction as more recommendations are being excluded. In essence, recommendations with higher conviction are more accurate but also fewer.

Note that the test set used for the evaluation is not the ground truth. Many false positives may exist; for example, because customers were not aware of the product or because they are already using an energy efficient product. For this reason the false positives are likely to be overestimated and the precision underestimated. Despite this and other deficiencies of the test set, we used it for development purposes; it is important to be able to quickly estimate the performance of the algorithm in order to decide what improvements should or should not be accepted.

[4] Train and test sets are generated by splitting data into two random samples, 80% and 20% of data, respectively.

In addition to train-test set validation, we are testing the performance of EE Recommender in two targeting campaigns. In the first use case we will provide LED products recommendations to commercial customers. In the second use case we will target a wide variety of EE product to customers on an A10 rate schedule.

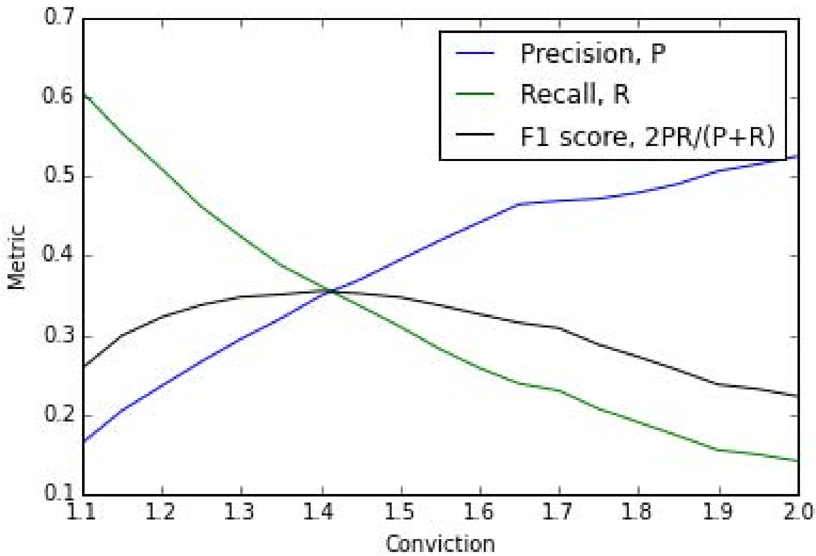

Figure 3- Precision, recall and F1 score calculated on the test set.

EE Recommender

We have discovered about 10,000 association rules that link various customer characteristics with EE offerings. These adoption patterns form the basis of the multilevel targeting capabilities of our EE Recommender system. Below we discuss a few selected capabilities:

- **Identifying top customers for product targeting**: A program manager or marketing team can retrieve a list of customers for a specific product. The list can be ordered by the strength of association between customer characteristics and the product, making it possible to prioritize a marketing campaign or to test the effectiveness of specific language on a small subset of customers.

- **Targeting customers at the account level**: sales representatives or call center consultants can provide personalized EE product recommendations. Figure 4 shows suggestions for a customer that belongs to the “Child Day Care Services” NAICS code.

For this specific customer PG&E can recommend HVAC quality maintenance, HVAC tune-up, power management software and lighting replacement. Lighting replacement is most strongly associated with the customer characteristics and similar customers have saved on average 1,520 kWh and 0.1 kW.

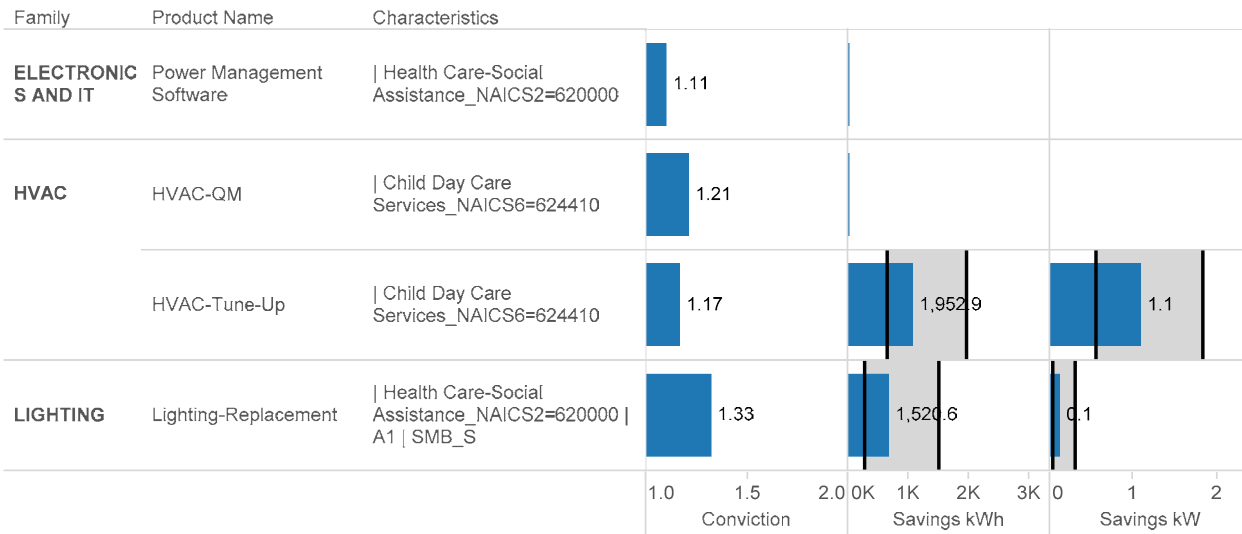

Figure 4- Personalized recommendations for a child day-care center.

Note: Conviction: measures the association strength;

Savings: median the kWh and kW savings for similar customers who have adopted the product, the shaded areas indicate 25th to 75th percentile range.

- **High level overview of the relationship between various customer characteristics and EE products**: This feature can aid program managers or marketing teams to understand possible drivers behind EE product adoption. Figure 5 shows the associations between industry NAICS codes and families of EE products. Not surprisingly, lighting products can be recommended to most industries, but pumps and fans are only strongly associated with agriculture, forest and fishing (NAICS = 720000), mining (NAICS = 210000), and utilities (NAICS = 220000). As opposed to segmentation, this high level overview clearly links customer characteristics with EE products.

	Family							
NAICS	APPLIANCES	BOILERS AND STEAM SYSTEMS	ELECTRONICS AND IT	FOOD SERVICE TECHNOLOGY	HVAC	LIGHTING	PUMPS AND FANS	REFRIGERATION
NAICS2 independent	HIGH			MED	HIGH	HIGH	HIGH	HIGH
Accommodation & Food Services_NAICS2=720000	HIGH			HIGH	LOW	HIGH		HIGH
Administrative-supt-Waste Mgmn_NAICS2=560000						MED		
Agriculture, Forestry, Fishing_NAICS2=110000	MED					HIGH	HIGH	LOW
Arts, Entertainment Recreation_NAICS2=710000					HIGH	HIGH		
Construction_NAICS2=230000						MED		
Educational Services_NAICS2=610000			MED		HIGH	HIGH		
Finance and Insurance_NAICS2=520000					HIGH	MED		
Health Care-Social Assistance_NAICS2=620000			HIGH		HIGH	HIGH		
Information_NAICS2=510000					HIGH	MED		
Manufacturing - 2_NAICS2=310000	MED			MED		LOW		MED
Manufacturing - 3_NAICS2=320000						MED		
Manufacturing - 4_NAICS2=330000					HIGH	HIGH		
Mgmnt of Companies-Enterprises_NAICS2=550000					HIGH	MED		
Mining_NAICS2=210000							HIGH	
Othr Svcs (excpt Public Admin)_NAICS2=810000	MED	HIGH			MED	HIGH		
Prof'l, Sci-Tech Svc_NAICS2=540000					MED	HIGH		
Public Administration_NAICS2=920000					MED	HIGH		
Real Estate & Rental & Leasing_NAICS2=530000					HIGH	HIGH		
Retail Trade - 1_NAICS2=440000	MED			HIGH		HIGH		HIGH
Retail Trade - 2_NAICS2=450000				HIGH		HIGH		MED
Transportation & Warehousing_NAICS2=490000						HIGH		
Transportation_NAICS2=480000						HIGH		
Utilities_NAICS2=220000							HIGH	
Wholesale Trade_NAICS2=420000						HIGH		

Figure 5 - EE opportunities for industries

- **Locational EE savings potential**: Recommendations can be aggregated geospatially providing estimates for locational load reduction. This capability can be particularly useful as a component of an integrated demand side management (IDSM) system. Figure 6, for instance, shows a map of PG&E's service territory and estimated kWh savings at ZIP code level for all EE Recommender suggestions with conviction higher than 1.2. Higher savings in California's Central Valley (the greener areas in Fig. 6) are primarily due to high savings from agriculture pump recommendations.

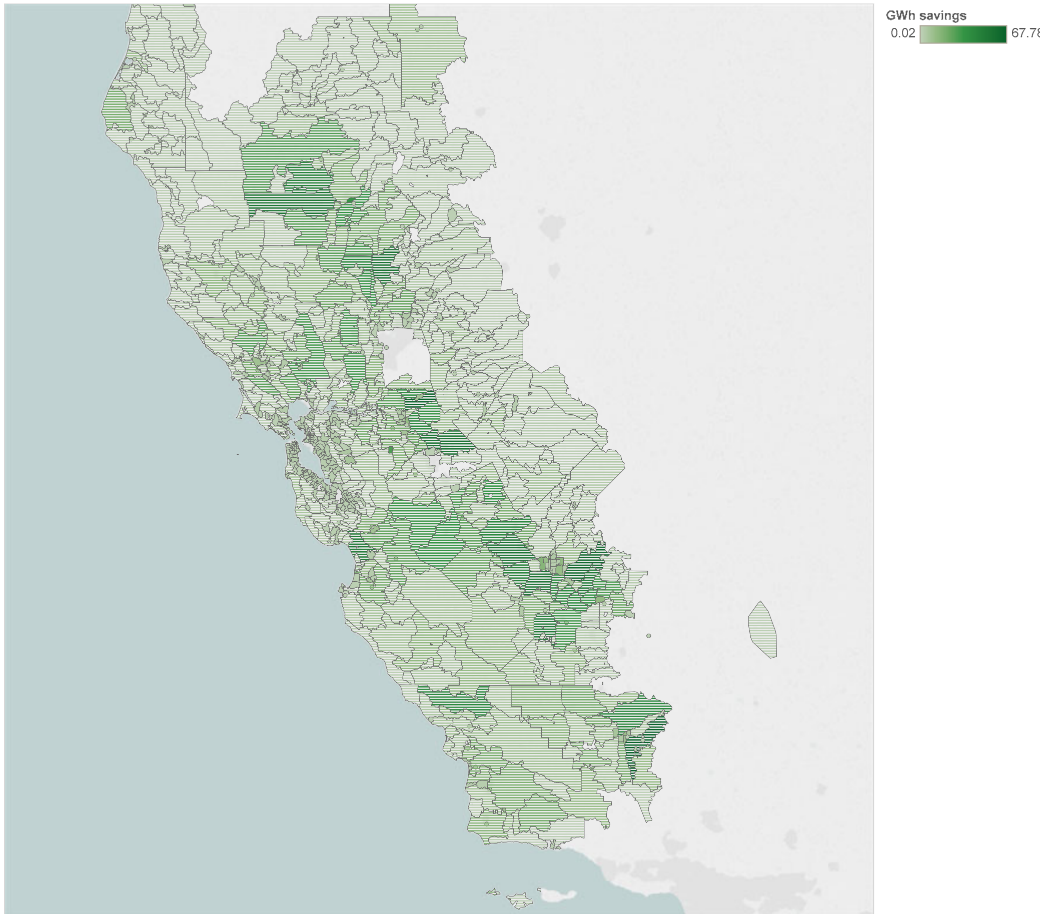

Figure 6 - EE savings potential for ZIP code areas within PG&E's territory summed over all EE recommendations with conviction greater than 1.2. Note, however, that savings estimates represent total savings not adjusted for adoption rate.

Conclusions

Pacific Gas and Electric's Data Analytics and Governance team has developed a knowledge-based recommendation system for commercial customers. The developed algorithm is able to make recommendations to customers who have not previously adopted any EE measure and therefore avoids the so-called cold start problem of collaborative filtering. Recommendations are made by matching customer characteristics with EE product adoption patterns discovered with association rule learning. EE Recommender is able to provide both personalized and aggregated suggestions and can serve as a targeting tool for marketing teams, sales representatives and program managers.

Acknowledgments

We would like to acknowledge the members of the marketing and business energy solution teams for numerus insightful discussions. We are grateful to Timotej Gavrilovic and Grace He; this work would not be possible without their support. We would like to thank JP Dolphin and Laura Winger for their careful reading of the manuscript.

References

CPUC (California Public Utilities Commission). 2013. *Energy Efficiency Policy Manual* http://www.cpuc.ca.gov/uploadedFiles/CPUC_Public_Website/Content/Utilities_and_Industries/Energy_-_Electricity_and_Natural_Gas/EEPolicyManualV5forPDF.pdf

Finger, Lutz. 2014. *Recommendation Engines: The Reason Why We Love Big Data* *http://www.forbes.com/sites/lutzfinger/2014/09/02/recommendation-engines-the-reason-why-we-love-big-data/#74676d2e218e*

Han, J., M. Kamber, J. Pei. 2012. *Data Mining: Concepts and Techniques*, *Third Edition*. Morgan Kaufmann Publishers

Hutson, S., 2016. *Utility Marketing Survey 2015*. https://www.esource.com/marketing-survey

Jannach, D., M. Zanker M., A. Felfernig, G. Friedrich. 2010. *Recommender Systems: An Introduction*. Cambridge University Press

Linoff, G. S., and M. J.A. Berry. 2011. *Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management, Third Edition*. Wiley

Netflix, 2009, "Nextflix Prize" Accessed March 01, 2016 http://www.netflixprize.com/

Palmeri C. 1997. *Believe in yourself, believe in the merchandise* www.forbes.com/forbes/1997/0908/6005118a.html

PG&E. 2015. *PG&E SMB Foundational Customer Insights* (Internal report)