

Thermal profiling of residential energy consumption

Adrian Albert and Ram Rajagopal

Abstract—Demand Response (DR) programs aim to dynamically match consumption on the grid with available supply in real-time. Understanding the patterns in demand of individuals is now being facilitated by granular consumption data collected via *smart meter* sensors that power utility companies have rolled out at scale. In this paper we propose a dynamic model that uses hourly electricity and weather readings to characterize residential users’ thermally-sensitive consumption. From this model we extract useful benchmarks to build profiles of individual users for use with DR programs that focus on temperature-dependent consumption such as air conditioning or heating. We present example profiles generated using our model on real consumers, and show its performance on a large sample of residential users. We then compute metrics that allow us to segment the population dynamically for the purpose of a thermally-motivated DR program. We show that such segmentation and targeting of users may offer savings exceeding 100% of the current performance of such programs.

I. INTRODUCTION AND MOTIVATION

In recent years several states and major cities have passed stringent environmental commitments, such as the pledge of 30% emissions reduction by 2030 in New York, and the 25% emissions reduction plan by 2020 in California. These measures directly affect the operations of energy utility companies, as electricity accounts for $\sim 40\%$ of total energy use and 34% of *GHG* emissions in the U.S. [1]. As such utilities are required to meet efficiency and consumption reduction targets, as well as integrate more renewables in their generation portfolio. Yet these changes are increasing the uncertainty in both supply and demand on the grid, which is becoming a primary operational challenge for utilities.

As a result, there has been a recent effort by utilities to roll out advanced sensing infrastructure to better meter energy consumption and inform demand management practices such as Demand-Response [2]. Millions of “smart meters” have been deployed in CA which are collecting highly granular data (hourly or sub-hourly) about energy consumption. Yet to date little is known at utilities on how to extract information out of this wealth of data, and little precedent exists on how to capitalize on this information to achieve efficiency goals.

This paper is motivated by certain types of Demand-Response designed to reduce loads associated with heating or cooling of a residential premise. Air conditioning and space heating loads are good targets for tailored DR events, since they make up for a sizeable component (27%) of electricity use in the U.S. [3], and are generally deferrable in time, since

the thermal mass of the premise may act as “thermal battery”. Affecting the thermally-sensitive load may be typically achieved through direct load control of the HVAC system (e.g., load curtailment or automatic adjustment of the thermostat setpoint), through adjustable rates (e.g., critical peak pricing), or through incentive schemes [4], [5].

Here we propose a simple model of consumption for a residential premise that is driven by unobserved “occupancy states” that have different responses to ambient weather. These are consumption regimes of a given household that depend on lifestyle (work schedule, familial composition etc.), premise characteristics (heating/cooling mass, square footage etc.), appliance stock, and weather patterns. It is a daunting task to disentangle how much energy each of these components accounts for at a given time (which is the task of a related research direction, *Non-Intrusive Appliance Monitoring* see [6]), especially since obtaining ground-truth information on individual appliance consumption and on premise and user characteristics is intrusive and expensive. But such detailed information might not even be required to design effective DR programs and tailored targeting strategies. However, given information on exogenous covariates (such as outside temperature) we would like to *i)* be able to estimate a high-level thermal response of the given premise, *ii)* assess the likely duration of the current heating or cooling spell, and *iii)* characterize the probability that for a perceived temperature the premise will be either heating or cooling.

The rest of the paper is structured as follows. Section II introduces the profiling problem. Section III discusses the literature on smart meter data analytics. Section IV introduces the consumption model proposed and outlines the algorithms used for estimation. Section VI describes the datasets used throughout the paper. In Section VII model performance on real users is presented. Section VIII discusses benchmarks computed from a large user population, and presents an example application of segmentation and targeting for a DR program. Section IX concludes the paper.

II. PROBLEM STATEMENT

Thermal profiling. A high-level schematic of the thermal profiling methodology proposed here is given in Figure 1. We observe hourly-sampled energy consumption time series (measured in kWh) $\{X_t\}_n$ and outside temperature time series $\{T_t\}_n$ for premises (users) $n = 1, \dots, N$. In the Decoding step we would like to separate the signal $\{X_t\}$ for a given user n into time-consistent portions that are (linearly) driven by temperature, and segments that are not affected by temperature. In effect, we cluster the observations in $\{X_t\}$ according to how they change over time with changes in $\{T_t\}$. To achieve this we allow the premise to experience regimes, or *occupancy states* characterized by different

A. Albert is with the Electrical Engineering Department at Stanford University

R. Rajagopal is with the Civil and Environmental Engineering Department at Stanford University

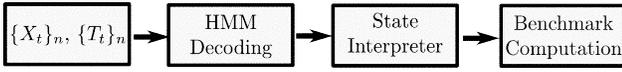


Fig. 1. Thermal profiling methodology: HMM decoding, state interpretation, and benchmark computation.

consumption responses with temperature, and different base levels of activity over time. In addition, we explicitly model a temperature-dependent probability of the premise performing heating or cooling. This formulation is based on the following observations: *i*) heating or cooling spells generally span multiple hours, and given a certain response (say heating) at a given hour, it is likely that the next hour will see the same thermal activity; *ii*) while the premise may have appliances that work on pre-determined schedules (e.g., an automatic thermostat that maintains inside temperature within a certain admissible range), these settings may change dynamically based on user preferences and occupancy, and *iii*) heating or cooling appliances do not consume fixed amounts of energy when they are on, but undergo different operational regimes depending on how much work is needed to respond to the outside temperature experienced.

Interpretation and benchmarks. We develop the following benchmarks to characterize household consumption.

Temperature response levels: We identify time-, activity- and temperature-consistent regimes of consumption and use temperature sensitivities as metrics for a coarse user profile.

Temperature-dependent duration of heating/cooling spells: We compute the characteristic duration of heating, cooling, or non-thermal activity regimes based on temperature.

Likelihood of heating/cooling at a given temperature level: In a relaxation of the fixed setpoints in the breakpoint model [3], we allow different regimes to be triggered with different probabilities based on the outside temperature level. This more flexible model is better suited for the highly-volatile hourly residential consumption patterns observed in the real data, as well as accounts for the unobserved user decisions that may change dynamically with time.

III. LITERATURE REVIEW

In the nascent literature on energy analytics, particular emphasis has been placed on customer segmentation and related demand-management applications. For example, [7] uses 15-minute resolution smart meter data from ~ 200 customers of an utility company in Germany to cluster consumers according to their daily consumption profiles, and to argue for different pricing schemes for each of the clusters. Similarly, [8], [9] develop methods for describing intra-day consumption through a small number of recurring patterns.

A large body of literature exists on modeling weather (and in particular temperature) response of energy use in residential and commercial premises. Most previous studies have been performed on aggregated data (e.g., [10], [11]) because of the lack of intra-day measurements. New sensing capabilities have enabled studies of the thermal response at an hourly level such as [3]. There, the authors break down the temperature profile

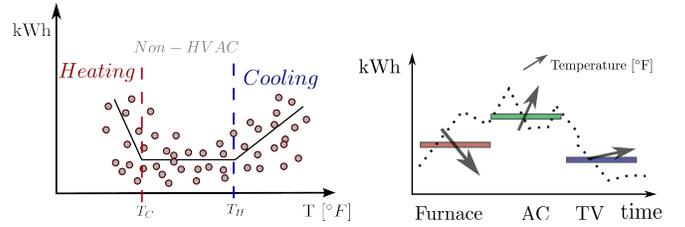


Fig. 2. *Top:* Schematic of the (fixed) breakpoint model benchmark: whole-home reading with thermal and non-HVAC activity [3]; *Bottom:* Occupancy states having different activity levels (height of the horizontal bars), thermal responsiveness (arrow slope), and characteristic duration (width of the bars).

of usage into four categories (base load, activity load, heating season gradient, and cooling season gradient) and model the set point of HVAC equipment such as air conditioning (AC) or heating furnace. However this type of models are rather static - they do not allow for different temperature sensitivities based on temperature, or on different levels of activity. A more complex temperature model is developed in [12], where the existence and operation of the heating appliance is disaggregated from the total consumption of a residential premise. Similar to our model, the authors use a Hidden Markov Model that follows the (hourly) dynamics of a thermal load. However this model does not allow for heating consumption to vary with temperature, which is generally not a valid assumption for real buildings.

In the context the *disaggregation* problem of recovering individual appliance signals from the aggregate load profile [6] much emphasis has been placed on uncovering *all* major end uses through very granular (Hz or kHz-range) readings. Instead, here we propose a simpler, coarser thermal disaggregation in the absence of ground truth data, with the purpose of developing high-level metrics that may serve for DR segmentation and targeting, which is one of the intended uses of disaggregation. For doing so we build on our previous work [13], where we show that dynamic characteristics of consumption computed from individual energy time series using Hidden Markov Models are predictive of both the presence of large appliances and of certain user lifestyles.

IV. OCCUPANCY STATES MODEL

A. The breakpoint model

Currently a popular framework in modelling the temperature response of residential premise energy consumption is the so-called “breakpoint model” [3], [12]. For observed times $t = 1, \dots, \tau$, the model assumes i.i.d energy readings y_t that depend linearly with temperature

$$x_t = \beta_0 + \beta_-(T - T_C)_- + \beta_+(T - T_H)_+ + \epsilon_t, \quad (1)$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$, $(z)_+ \equiv \max(z, 0)$ and $(z)_- \equiv \min(z, 0)$. This model aims to uncover the admissible temperature range of operation of heating and cooling appliances defined by a “cold setpoint” T_C (temperatures below T_C trigger the activation of the heating unit) and a “hot setpoint” T_H (temperatures above T_H trigger the activation of the

cooling unit). Schematically, this model is presented in Figure 2 (top panel) as discussed in [3]. We estimate this model as in [14]. While informative for premises operating on strict HVAC schedules, this model fails to account for the large volatility and heterogeneity observed in the hourly smart meter data. We address some of its shortcomings below.

B. Occupancy states model

We propose a consumption model driven by “occupancy states” as motivated above in Section II where each usage regime has a different activity level and a sensitivity to outside temperature. As discussed in [3] and illustrated in Figure 2, non-HVAC consumption decisions are associated with a negligible temperature sensitivity and high variance; by contrast HVAC-intensive regimes will exhibit a pronounced sensitivity of consumption with outside temperature. Moreover, we allow the occupancy states to be either “sticky” or “transient”, i.e., to have either long or short (bursty) durations, and for the transition between states to depend on temperature. A illustrative scenario is given in Figure 2 (bottom panel). We again stress that in the absence of ground truth about individual appliance consumption it is very difficult to assign meaning to these states; however correlating consumption with weather variables such as temperature or lighting may provide a step towards interpretation.

Model formulation. We view an individual premise as a state machine consuming energy in either of M (unobserved) states $S = \{1, \dots, K\}$. At time t , when recorded outside temperature is T_t , and if the premise is in a given state k , we assume consumption x_t to be described by

$$x_t | S_t = k, T_t \sim \mathcal{N}(\beta_0^k + \beta_1^k T_t, (\sigma^2)^k), \quad (2)$$

where $\mathcal{N}(\cdot, \cdot)$ denotes a standard Gaussian distribution, and $\beta^k = (\beta_0^k, \beta_1^k)$ and $(\sigma^2)^k$ are parameters to be estimated.

We further assume that the sequence of states $\{S_t\}$ follows a Markov process that depends on the outside temperature T_t :

$$P(S_{t+1} = k | S_t = j, T_t) = \frac{\exp(\gamma_0^{k,j} + \gamma_1^{k,j} T_t)}{\sum_i \exp(\gamma_0^{i,j} + \gamma_1^{i,j} T_t)}, \quad (3)$$

i.e., a multinomial logistic regression where the response is the hidden state. Above we made use of the Markov assumption that the state of the system at time $t + 1$ only depends on its state at the current time step t , but not on all past history. As above, we estimate $\gamma^k = (\gamma_0^k, \gamma_1^k)$ from data. We may group all transition probabilities into a temperature-dependent state transition matrix $\{A(T)\}_{ij} \equiv P(S_t = i | S_{t-1} = j, T)$. For simplicity we assume that the initial probability distribution over the states is uniform and does not depend on temperature, i.e., $\pi_k = P(S_1 = k) = \frac{1}{K}$.

Estimation. The formulation above is a simple extension of a Hidden Markov Model [15] for which the means of both the emission and the transition distributions are allowed to depend linearly on an exogenous covariate (temperature). For each user n we wish to estimate $\theta \equiv (\beta^k, (\sigma^2)^k, \gamma^k, k = 1, \dots, K)_n$ using the Maximum-Likelihood framework. This is usually performed via the Baum-Welch algorithm [15], which

is a variant of the EM algorithm [16] for dealing with missing data (here the missing observations are the unobserved states $\{S_t\}$). For this we may write the (complete) log-likelihood for data $\{X_t\}_{t=1}^\tau$ under the model (2, 3) as

$$\begin{aligned} \log \mathcal{L}(\theta) &= \log P(\mathbf{x}, \mathbf{S} | \mathbf{T}, \theta) \\ &= \log \pi(S_0) + \sum_{t=1}^{\tau} \log P(S_t | S_{t-1}, T_{t-1}, \theta_1) + \\ &\quad + \sum_{t=1}^{\tau} \log P(X_t | S_t, T_t, \theta_2), \end{aligned} \quad (4)$$

where $\theta = (\theta_1, \theta_2)$. The terms in the summation above are independent and may be maximized separately. This proceeds using the standard EM algorithm as in [17]:

- 1) Start with a guess θ^0 ;
- 2) *E-step*: compute expected value of complete-data log-likelihood $Q(\theta | \theta^k) = \mathbb{E}[\log \mathcal{L}(\mathbf{x}; \theta^k)]$ as in [17];
- 3) *M-step*: choose $\theta^{k+1} = \operatorname{argmax} Q(\theta | \theta^k)$. For θ_1 in (2) this amounts to a linear regression estimated via ordinary least squares (OLS). For θ_2 the maximization involves estimating a multinomial logistic model as in (3) for each row of the transition matrix.
- 4) Repeat (2) and (3) until convergence.

We compute the the most likely sequence of states \mathbf{S} that fits a given observation sequence \mathbf{x} (the the *decoding* problem) using the standard Viterbi algorithm [15].

Choosing model size. Until now we have assumed the number K of states known; however this is not the case in real applications. Here we adopt a simple selection strategy based on out-of-sample predictive performance of the model. Note that a standard k -fold cross-validation approach [16] is not appropriate here because the random segmentation of the data will violate the serial correlation assumed by the Markov process. To overcome this issue we adopt a deterministic 2-fold cross-validation approach as follows:

- 1) Start with a model of $K = 2$;
- 2) Divide up the time series into an even and an odd sequence, and learn the model (2, 3) of a given K on the even sequence. As discussed in [18], the model parameters learned this way are the same as for the model learned on the full data, with the exception of the transition matrix of the half-chains being A^2 (where A is the transition matrix of the full chain);
- 3) Compute out-of-sample decoding performance (using the Viterbi algorithm) of the even-chain model on the odd-chain model. Here we assess performance using the variance explained (R^2) and Mean Absolute Percentage Error ($MAPE$) metrics defined in the usual way;
- 4) Increment K , and repeat (2) and (3) until the out-of-sample performance reaches a desired threshold α (here we used 0.85 for the R^2 metric and 0.15 for $MAPE$).

V. MODEL BENCHMARKS

We use our model to compute temperature-sensitive (or insensitive) regimes of consumption beyond the three identified in Figure 2. Second, we allow the states to have explicit

duration that is based on temperature. We use a standard result in the analysis of Markov chains [15] to define a mean time spent in a state k for a given temperature T :

$$\tau_k = \frac{1}{1 - A(T)_{kk}}, \quad (5)$$

with $A(T)_{kk}$ the diagonal elements of the temperature-dependent Markov transition matrix as defined above.

Third, we relax the hard breakpoint in (2) to allow for a soft temperature-dependent regime selection. For a given temperature T we may compute the probability that the premise will undergo a regime change according to (3); as an extension we may also compute the long-run probability distribution π of the premise being in either of the K states using a standard result from Markov chain theory [15]:

$$\pi(T) = \pi(T)A(T), \quad (6)$$

from which $\pi(T)$ may be obtained by finding the (normalized) 1-eigenvector of $A(T)$. Furthermore, it is useful for analysis to define an *effective thermal response* $\hat{\beta}_1(T) = \sum_k \beta_1^k \pi_k(T)$ at a given temperature level T . We also define an *effective thermal duration* $\hat{\tau}(T) = \sum_k \tau^k \pi_k(T)$.

Lastly, we recognize that the model proposed here has limitations in discerning between possible confounding effects of temperature and diurnal patterns. That is, because both temperature and human activity follow a typical 24-hour cycle, the regression models in (5) and (2, 3) may over (or under-) estimate the “true” thermal contribution. A partial solution we adopt is to compute average weekly thermal profiles for a household, since averaging will allow errors in the hourly estimations to cancel out to some extent.

VI. DATA DESCRIPTION

In this paper we use three types of data. We first discuss the behavior of our model when estimated on real, high-frequency data for which ground truth HVAC readings were available. We then profile the consumption of a large sample of real users. In each case we collected the appropriate weather time series (at the 5-digit zipcode level) for each of the real premises that we used in our analysis using an online API at www.wunderground.com.

Ground-truth data. For an illustration of our model we used a publicly-available, high-resolution (15 kHz) dataset [19] (the Residential Energy Disaggregation Dataset, or REDD) that contains readings from several individually-monitored appliances as well as whole-home circuits for several houses in Massachusetts and California. We selected a premise (`house_13`) that had separate furnace readings and enough contiguous data aggregated at a hourly level (~ 900 hours between Jan. 9th-Feb. 21st, 2012) for our analysis.

Real premise data. We illustrate user thermal profiling in a real-world context by estimating our model on a large sample of 1,923 premises in a hot climate zone around Bakersfield, CA (zipcodes 93309, 93301, 93304, 93305). We obtained this sample from the Pacific Gas and Electric Company. This is whole-premise data at an hourly level and spans one year from August 30th, 2010 to July 31st, 2011.

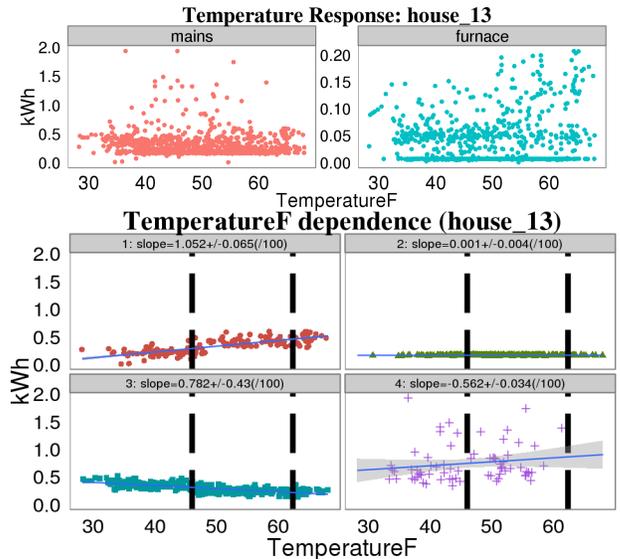


Fig. 3. *Top:* Consumption profile with temperature for whole-home signal (*mains*) and thermal unit (*furnace*) for one example premise in the REDD dataset; *Bottom:* Temperature profiles learned using the occupancy states model separated by state. Dashed lines represent the estimated “hard” breakpoint values estimated from (5).

VII. CASE STUDY: INDIVIDUAL THERMAL PROFILES

A. REDD ground-truth data

In Figure 3 we present temperature profiles of the whole-home signal (*mains*) and thermally-dependent unit (*furnace*) for one premise in the REDD dataset (`house_13`). Note the general similarity of the furnace profile with the sketch in Figure 2 derived from [3], suggesting that both heating and cooling may take place. We learned both the breakpoint (1) and the occupancy state models on the whole-home panel for the premise, and present the results in the bottom panel in the figure. In all figures the two vertical dashed lines represent the estimated “hard” breakpoint values estimated from (5) (at 45°F and 62°F for this example premise). Each panel represents the observations assigned by the model to either of the states uncovered. Estimates of the temperature response slope and corresponding standard error are also indicated.

An example fit for one week of data for `house_13` is given in the top panel in Figure 4 (in-sample $MAPE = 0.11$, day-ahead out-of-sample decoding $MAPE = 0.15$). The two middle panels in the figure present a comparison by hour-of-day of the HVAC (*furnace*) and non-HVAC components, as well as the shares of the decoded states in the HMM fit. Note the larger share of temperature-sensitive states 1 and 3 (red and blue in the figure) for the hours for which furnace activity is recorded. However this simple model will not, in general, offer accurate estimates of thermal energy consumption [3], [12]; yet it may serve to derive useful benchmarks about consumption across different premises for the purpose of comparison and classification. With this in mind, we compare the performance of the two models on detecting thermal activity in excess of a certain threshold (here we used the lower quartile of furnace

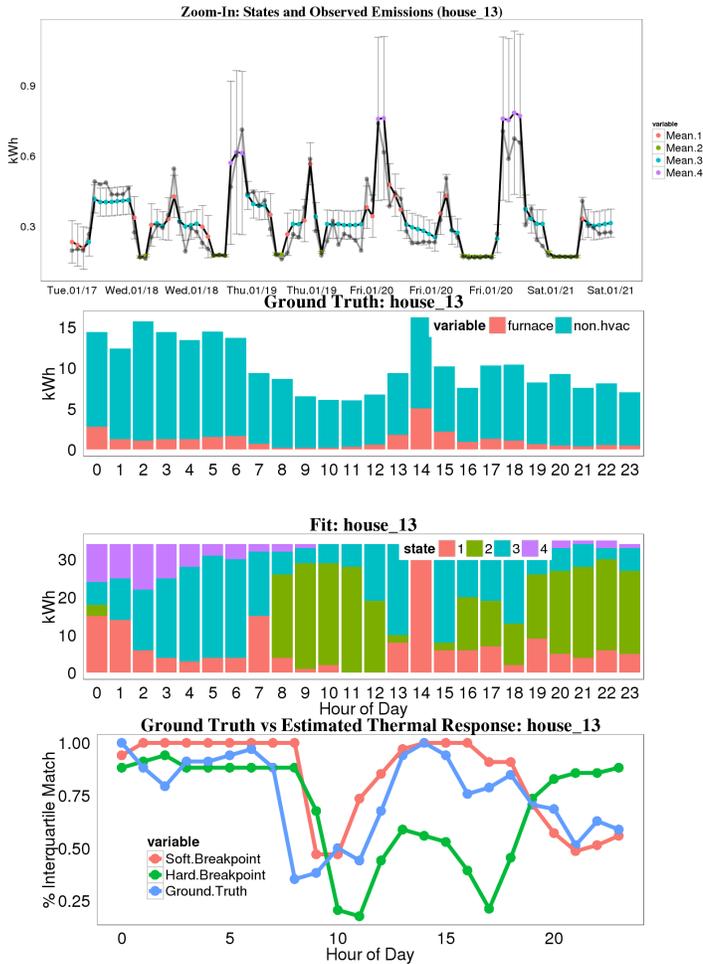


Fig. 4. *Top*: One week of whole-home data for house_13 (gray line) and model estimates (black line) colored by state; *Middle two panels*: hour-of-day comparison of HVAC and non-HVAC consumption with frequency of decoded states; *Bottom*: above lower quartile detection performance.

energy consumption). The results are illustrated in the lower panel in Figure 4 as percentage correct detections by hour-of-day. Used as a detector, our model identifies significant thermal activity more accurately on average than the simple breakpoint model (70% to 54%), and follows the ground truth furnace profile more closely.

B. Two real premises

In Figure 5 we present the yearly consumption profiles of two users (anonimized IDs 3284167 and 3675267) in Bakersfield, CA. For user 1 (left panel), consumption ramps in the summer up very visibly in the afternoon, and follows a double-peak (morning and evening) profile in the winter. A similar, albeit less pronounced profile is followed by user 2 in the right panel of the figure. The differences between the two users become quite clear when looking at their temperature profiles presented in Figure 6. User 1 has a temperature dependence that is similar to the heating-and-cooling profile in Figure 2, while User 2 primarily has cooling activity.

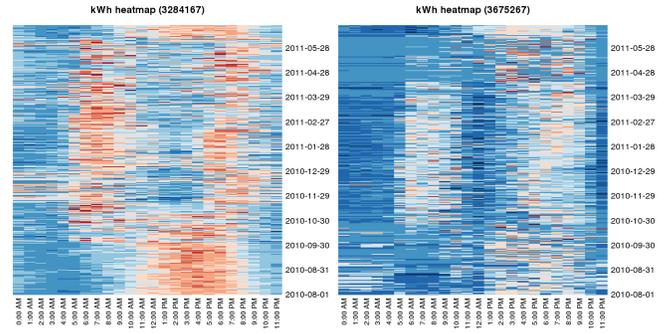


Fig. 5. Yearly consumption profiles for two selected users in Bakersfield, CA. Consumption is color-coded (red is high, blue is low) over 24 hours (horizontal axis) for 365 days (vertical axis).

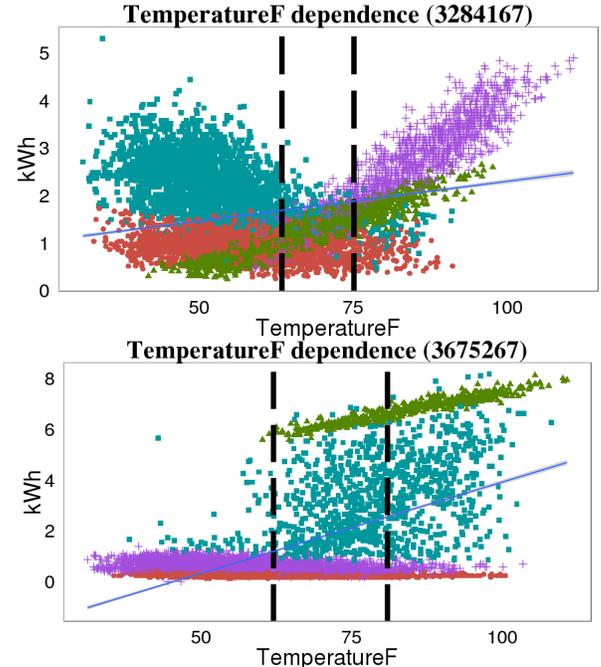


Fig. 6. Temperature profiles and identified thermal regimes for two users.

The temperature profiles in Figure 6 are color-coded by the thermal regimes identified using our model. Both users' consumption may be explained using four thermally-dependent occupancy regimes to a cross-validation $R^2 \geq 85\%$. For User 1, the model identifies two strong cooling states 2 and 4 (green and purple in the figure) for which $\beta_1 > 0$, one temperature-independent state 1 (red in the figure), and one strong cooling state 3 (light blue). Note that in the case of User 2 the model is able to identify a temporally-consistent cooling state 2 (green) that would be otherwise masked by the high-variance cooling state 3 (light blue).

Figure 7 presents the breakdown of thermal occupancy states for the two example users by the summer and winter seasons (as defined by PG&E) and by hour-of-day. For User 1 (top two panels) we notice that the strong cooling state 4 (purple) predominates during summer afternoon hours, while the strong

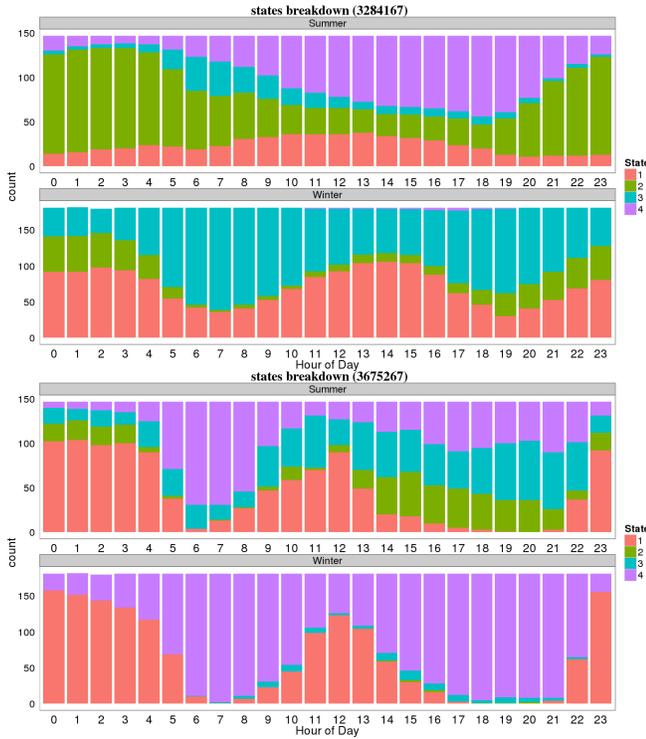


Fig. 7. Seasonal and time-of-day distribution of occupancy states for two real users in Bakersfield, CA.

heating state 3 (light blue) predominates during morning and evening hours in the winter. For User 2 the strong cooling state 2 (green) is identified during summer afternoons; the highly-volatile cooling state 3 (light blue) also predominates during summer afternoons. In contrast, winters for User 2 are spent in occupancy states that are relatively temperature-insensitive; this is understandable since the user resides in a hot area (inland Central California). Note that the time-of-day distribution of the temperature-insensitive state 1 (red) for User 2 is relatively stable over summer and winter; we interpret this as regimes of low occupancy and activity.

We computed the benchmarks introduced above in Section V for the two users, and present the temperature-dependent stationary probability distribution over thermal occupancy regimes in Figure 8. For User 1 the benchmark captures the prevalence of the heating state 3 (light blue) for low temperatures, the medium cooling state 2 (green) for intermediate temperatures, and the high cooling state 4 (purple) for temperatures higher than 75°F. This coincides with the predictions of the breakpoint model (1) indicated by the vertical dashed lines in the figure. Similarly, for User 2 the model identifies the temporally-consistent cooling state 2 (green) as dominant for high temperatures (above 95°F). In the bottom panel in the figure we show the effective thermal response $\hat{\beta}_1$ for the two users. Note that for both users the transition between thermally-insensitive regimes and thermally-sensitive regimes happens gradually as temperature increases. The effective response profile may then be used as means to discriminate

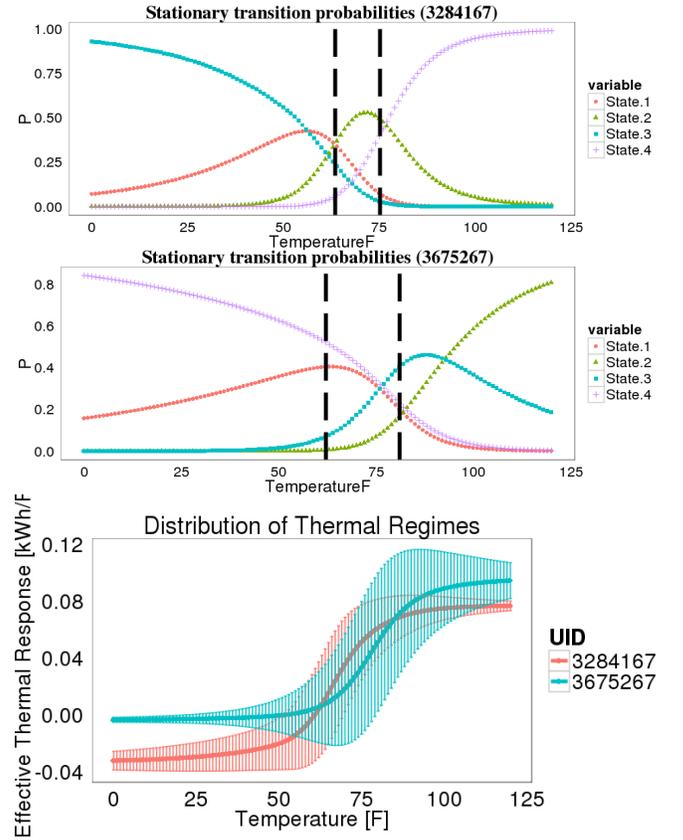


Fig. 8. *Top and Middle*: Temperature-dependent occupancy state probability for two users; *Bottom*: effective thermal response for the two users. Error bars show the effective variance computed as in the Appendix.

among users to decide the best targets for tailored DR events.

VIII. EXAMPLE APPLICATION: THERMAL SEGMENTATION AND TARGETING

A. Profiling a user population

We learned individual occupancy state models for each of 1,923 real households in the PG&E sample. We present model performance results in Figure 9. For the large majority of the users we need just 4 states to achieve at least 85% variance explained (R^2 , see bottom-left panel) out-of-sample in our cross-validation estimation procedure detailed in Section IV. Moreover, the fit performance is very good also on the *MAPE* fit metric (bottom-right panel). Here, out-of-sample cross-validation performance is reported by Viterbi-decoding the observations in the test set using the occupancy state model learned on the train set.

B. A simple segmentation and targeting scenario

In Figure 10 we present a simple segmentation of users in our sample by the effective duration a user spends in a given state $\hat{\tau}(T) = \sum_k \pi_k(T) T_k(T)$ and the effective thermal response at the given temperature T . We classified users according to the corresponding quartiles across the population

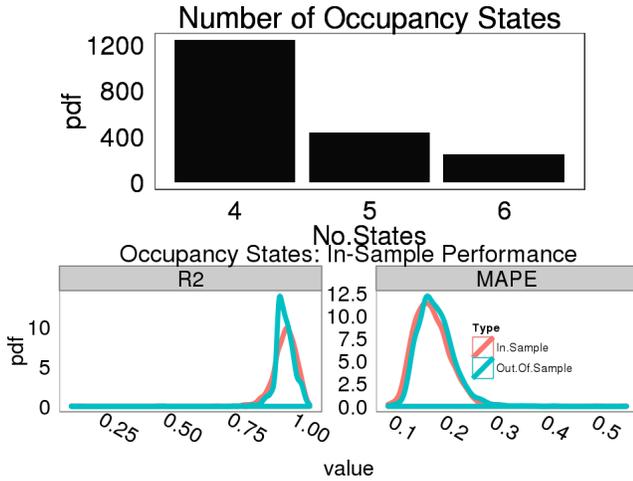


Fig. 9. Model performance on a sample of 1,923 users in four zipcodes near Bakersfield, CA. *Top*: breakdown of model size; *Bottom*: distribution of fit metrics R^2 and $MAPE$ for in-sample and out-of-sample performance.

in which the user’s response falls under these two benchmarks. As such, one user’s effective thermal response may be either zero (very small response), low, or high heating or cooling (5 tiers total). Similarly, regime duration is classified into short, medium, and long.

We next study the following scenario for user selection in a Demand-Response program. Suppose that for a given forecast on temperature levels next hour for which some amount of cooling activity is to be expected in the hot Bakersfield, CA area ($T \in \{45^\circ F, 60^\circ F, 75^\circ F, 95^\circ F\}$) the utility issues DR events asking users to reduce their air conditioning level by $1^\circ F$. This is equivalent to the house experiencing a level of outside temperature T that is cooler by $1^\circ F$; in turn, this action yields an averted energy consumption of $\hat{\beta}_1(T) \times 1^\circ F$. A simple selection problem may read:

$$\max_{\mathbf{x}} \mathbb{E} \left[\sum_i x_i \beta_1(T) \right] \quad (7)$$

$$s.t. \sum_i x_i = N \text{ and } x_i \in \{0, 1\} \quad (8)$$

That is, the system operator wishes to select the subset of users $i \in \{1, \dots, N\}$ (indicated by $x_i = 1$) of a given size N such that the expected savings are maximized. We implemented two possible solutions to this problem for selected groups of users of increasing size N : *i*) random selection (default) and *ii*) a greedy selection strategy that first ranks users according to their effective cooling thermal response and selects the top subset. For this latter strategy we only used users that displayed either a medium or high effective cooling behavior at the given temperature level T . We present the results of this exercise in Figure 11. Expected savings gained by taking into account the effective thermal response far exceed the default (random) numbers as clearly seen in the top panel in the figure. Moreover, the relative performance increases with temperature (since at high temperatures more

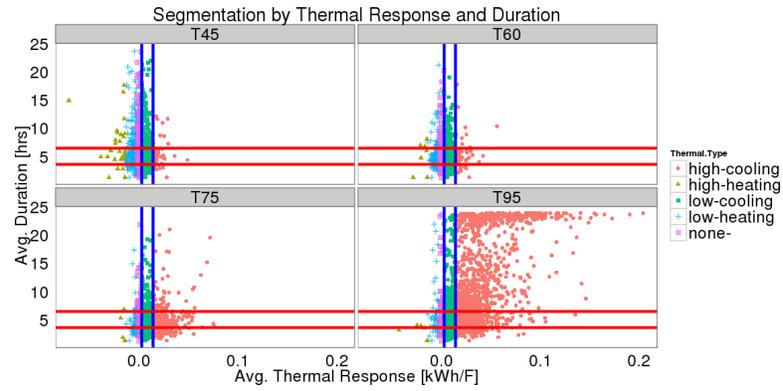


Fig. 10. Segmentation of occupancy states by magnitude/type (heating or cooling) and duration for four temperature levels (shown as contour density plots). Classes are formed according to the quartiles of the respective distributions.

users turn their AC systems on and have stronger effective cooling responses). Savings per $1^\circ F$ of effort for a given user subset size N for thermally-aware targeting exceed those of a random strategy by more than 100% for a hot hour ($T = 95^\circ F$) and a moderate-size target population ($N = 500$).

In the bottom panel of the figure we show the marginal benefit of enrolling additional users in the program at different levels of temperature. Note the “transition point” (which is different at each of the temperature levels in the figure, e.g., $N = 500$ for $T = 95^\circ F$) after which the marginal benefit from the thermal DR scheme is lower than the random selection strategy. This is partly an effect of the sample size and averaging - we had ~ 2000 users in our analysis, so the mean response for subsets larger than a certain size will certainly be larger than the smaller responses. The marginal savings will naturally decrease, but what is surprising is that for groups larger than $N \sim 100$ the marginal average savings decrease more slowly. This can be interpreted as a “phase transition” in program impact - it reinforces the point that most cost-effective savings may be achieved by enrolling a small number of the *right* users.

IX. CONCLUSIONS

We have developed a methodology to construct dynamic energy consumption profiles for individual users that is based on their response to outside temperature. Using this model we computed several benchmarks for characterizing individual premises’ consumption to be used segmentation and targeting for Demand-Response programs. In particular, we envision a situation where the operator may ask customers to affect the setting on their heating or cooling appliances as to avert consumption during certain times (e.g., peak times or particularly hot days). We show that a simple targeting strategy that is aware of the heterogeneity in thermal response may achieve savings (for a given population size) in excess of 100% of the performance of a random selection strategy.

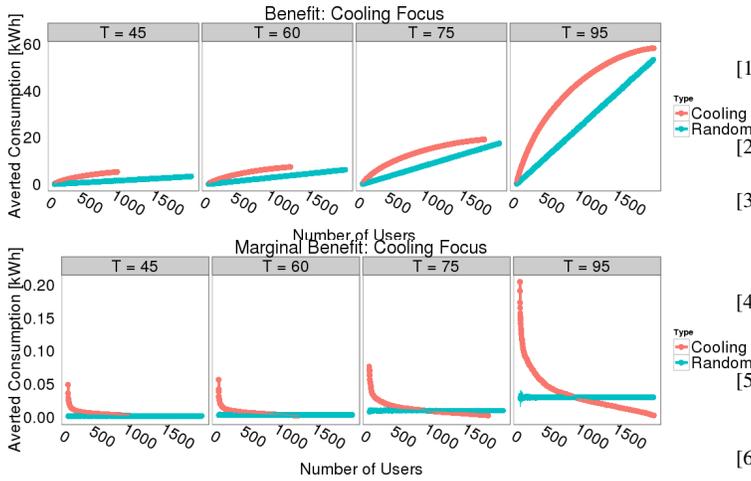


Fig. 11. Maximum targeting effectiveness by effort (number of users enrolled) for four levels of temperature - total (*Top*) and marginal (*Bottom*).

ACKNOWLEDGEMENTS

We thank Sam Borgeson for insightful discussion on this manuscript. We are grateful to Brian Smith and Ann George at PG&E for providing the dataset. A.A. thanks the Precourt Energy Efficiency Center, and its director, James Sweeney, for financial and intellectual support. R.R. acknowledges financial support from ARPA-E and the Powell Foundation.

APPENDIX A EFFECTIVE VARIANCE

For a given user experiencing a temperature level T , the model introduced in this paper allows to define a probability distribution $\pi(T)$ over occupancy states as in Section V. A given state k is drawn from a Gaussian with a temperature-dependent mean as in (2) and constant variance σ_k . This may be expressed by defining a random variable $X_k(T) \sim \mathcal{N}(\beta_0^k + \beta_1^k T, \sigma_k^2)$, and a random variable Z

$$Z(T) = \begin{cases} X_1(T) & \text{w.p. } \pi_1(T) \\ \vdots & \\ X_K(T) & \text{w.p. } \pi_K(T) \end{cases}$$

Then the effective variance of the random variable $Z(T)$ may be computed as (dropping T for convenience)

$$\begin{aligned} \text{Var}(Z) &= \mathbb{E}[(Z - \mathbb{E}[Z])^2] \\ &= \mathbb{E} \left[\left(Z - \sum_k \pi_k \mathbb{E}[X_k] \right)^2 \right] \\ &= \sum_k \pi_k \mathbb{E}[X_k^2] - \left[\sum_k \pi_k (\beta_0^k + \beta_1^k T) \right]^2 \\ &= \sum_k \pi_k \left[(\beta_0^k + \beta_1^k T)^2 + \sigma_k^2 \right] - \left[\sum_k \pi_k (\beta_0^k + \beta_1^k T) \right]^2. \end{aligned} \tag{9}$$

REFERENCES

- [1] Energy Information Administration. Annual energy review. Report DOE/EIA 0384, U.S. Department of Energy Office of Energy Markets and End Use, Washington, DC, 2007.
- [2] A. Faruqui, P. Fox-Penner, and R. Hledik. Smart grid strategy: Quantifying benefits. *Public Utilities Fortnightly*, July 2009.
- [3] Benjamin J. Birt, Guy R. Newsham, Ian Beausoleil-Morrison, Marianne M. Armstrong, Neil Saldanha, and Ian H. Rowlands. Disaggregating categories of electrical energy end-use from whole-house hourly data. *Energy and Buildings*, 50(0):93 – 102, 2012.
- [4] Peter Palensky and Dietmar Dietrich. Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads. *IEEE Transactions on Industrial Informatics*, 7:381–388, 2011.
- [5] Junqiao Han and Mary Ann Piette. Solutions for Summer Electric Power Shortages: Demand Response audits Applications in Air Conditioning and Refrigerating Systems. 2007.
- [6] C. K. Carmel, G. Shrimali, and A. Albert. Disaggregation: the holy grail of energy efficiency? *Energy Policy*, 2013.
- [7] Christoph Flath, David Nicolay, Tobias Conte, Clemens van Dinther, and Lilia Filipova-Neumann. Cluster analysis of smart metering data - an implementation in practice. *Business & Information Systems Engineering*, 4(1), 2012.
- [8] Teemu Räsänen and Mikko Kolehmainen. Feature-based clustering for electricity use time series data. In *Proceedings of the 9th international conference on Adaptive and natural computing algorithms, ICANNGA'09*, pages 401–412. Springer-Verlag, 2009.
- [9] V. Figueiredo, F. Rodrigues, Z. Vale, and J.B. Gouveia. An electric energy consumer characterization framework based on data mining techniques. *Power Systems, IEEE Transactions on*, 20(2), 2005.
- [10] Richard S. J. Tol, Sebastian Petrick, and Katrin Rehdanz. The impact of temperature changes on residential energy use. Working Paper Series 4412, Department of Economics, University of Sussex, November 2012.
- [11] Olivia Guerra Santin, Laure Itard, and Henk Visscher. The effect of occupancy and building characteristics on energy use for space and water heating in dutch residential stock. *Energy and Buildings*, 41(11):1223 – 1232, 2009.
- [12] Dayu Huang, M. Thottan, and F. Feather. Designing customized energy services based on disaggregation of heating usage. In *Innovative Smart Grid Technologies (ISGT), 2013 IEEE PES*, pages 1–6, 2013.
- [13] A. Albert and R. Rajagopal. Smart-meter driven segmentation: what your consumption says about you. *Accepted for publication, IEEE Transactions on Power Systems*, 2013.
- [14] Vito M.R. Muggeo. Estimating regression models with unknown breakpoints. *Statistics in Medicine*, 22:3055–3071, 2003.
- [15] Lawrence R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, pages 257–286, 1989.
- [16] T. Hastie, R. Tibshirani, and J.H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer series in statistics. Springer, 2009.
- [17] Ingmar Visser and Maarten Speekenbrink. depmixS4: An R package for hidden markov models. *Journal of Statistical Software*, 36(7):1–21, 2010.
- [18] Gilles Celeux and Jean-Baptiste Durand. Selecting hidden markov model state number with cross-validated likelihood. *Comput. Stat.*, 23(4):541–564, October 2008.
- [19] J. Zico Kolter and Matthew J. Johnson. REDD: A Public Data Set for Energy Disaggregation Research. In *SustKDD Workshop on Data Mining Applications in Sustainability*, 2011.