2020 Building Performance Analysis Conference and
SimBuild co-organized by ASHRAE and IBPSA-USA
Chicago, IL
August 12-14, 2020

IMPACT ANALYSIS OF PERSONALIZED THERMOSTAT DEMAND RESPONSE

Kunind Sharma[1], Michael Kane[2]
[1]Northeastern University, Boston, MA, USA

ABSTRACT

Demand response (DR) aims to curtail peak electric load to avoid running inefficient power plants, reduce costly transmission capacity increases, and improve grid reliability. One method of DR for residential settings involves the utility or third part remotely adjusting thermostat temperatures. Cooling setpoints may be increased 2°F – 6°F for hours, irrespective of occupant comfort. To maintain their comfort, occupants override these setpoint changes impacting the reliability of DR services provided to the grid.

This paper presents a preliminary impact analysis of replacing the current DR approach of uniform thermostat setbacks for an entire population with personalized and context aware DR based on a dynamic model of occupant thermal comfort behavior. Over 151k interactions with ecobee thermostats are analyzed to predict the impact of such personalized DR on overrides, energy curtailment magnitude, and reliability.

INTRODUCTION

Increasing load from electric vehicles and electrification of home heating with heat pumps; loss of reliable, yet dirty, fossil generation sources; and growth in the supply of power from variable renewable energy sources is driving a paradigm shift in the control of the grid. Legacy methods of controlling generation to match stochastic uncertainty of load no longer suffice. Load will need to be more flexible, controllable, and efficient. Buildings, which consume ~40% of the energy on the electric grid (U.S. Energy Information Administration, 2015), will be a prime target of flexible load programs for grid operators (Neukomm et al., 2019), a \$15B/yr. market opportunity (Hledik et al., 2019).

Residential HVAC systems can provide load flexibility through remote control of their thermostats. (Marie-Andrée et al., 2011; Singla and Keshav, 2012). In such a case, three independent variables to control load can be defined as (1) the amount the thermostat is changed to curtail energy (i.e., setback), (2) the duration that setback is applied before reverting to the previously set temperature, and (3) when that setback begins. This is the framework current demand response (DR) programs use. These programs can reduce aggregate peak energy consumption by 30% (Seiden et al., 2017) for peaks that may last a few hours. If many homes are available to aggregate and less curtailment is needed (either duration or power), then the duration and/or setback can be reduced and homes can be cycled to provide a small amount of curtailment that provides the desired net effect (Almassalkhi et al., 2018; Jin et al., 2017).

Thermostat setpoint changes naturally effect indoor temperature, which affects occupant comfort. However, DR events with large setbacks and durations may cause occupants to be so uncomfortable with the indoor environment that they override the grid's control of their thermostat, increasing the home's energy use and negating the benefits to the grid. For DR events up to 8 hrs., up to 30% of occupants may opt-out and override (Seiden et al., 2017). This can lead to the DR program failing to provide the promised curtailment to the grid, resulting in tens of millions of dollars in performance penalties each year (PJM Interconnection LLC, 2018).

To maximize curtailment and achieve acceptable reliability from overrides, DR programs must understand what causes opt-outs and overrides. For example, a $2°F$ setback may not seem like much, but if a thrifty customer already set their thermostat to the edge of their comfort zone to save on their electric bill, such a change could push them over the edge of discomfort, leading to an override (Cappers et al., 2016). Comfort is based on the indoor environment, but also on interactions of contextual and personal factors (Schweiker et al., 2020) Current DR programs neglect these confounding factors, and simply deploy a uniform setback and duration to all occupants.

Thermal comfort is a well-established area of study (Fanger, 1973). Yet, nearly all research has focused on the steady-state temperature at which an occupant, or a group of occupants are comfortable (Kim et al., 2018; Mirakhorli and Dong, 2016). These steady-state temperatures are not appropriate for use during DR events when the indoor temperature changes due to the thermostat changes and building thermal dynamics. Only a few truly dynamic comfort models exist (Chen et al., 2015; Kane, 2018; Zhang et al., 2010), which suffer from too little experimental data to fit them to contextual and personal factors that affect thermal perception and behavior.

This paper has three objectives: (1) present a framework for modeling the dynamic behavior of overrides, (2) develop a neural network-based approach to personalized and context aware override prediction, and (3) quantify the potential for personalized and context aware DR to increase curtailment and reliability compared to uniform setbacks and durations applied across the population.

METHODS

Overall this paper aims to answer the question "How can personalized and context-aware DR durations improve DR reliability and energy curtailment?" This is investigated through four questions. **(Q1)** Assuming no overrides, how do the DR setback and duration affect the percent of energy curtailed? **(Q2)** How do overrides affect the answer to Q1? **(Q3)** If a DR program uses the answer to Q1 to predict curtailment, but only achieves the answer to Q2, what is the resulting error in curtailment? **(Q4)** How much will the percent of energy curtailed and estimation error improve if a neural network is used to predict when an override may occur, and set the duration just short of the amount of time, using environmental, contextual, and personal features?

Dataset Curation and Preprocessing

The answer to these questions lies in the ecobee Donate your Data (DyD dataset (Ecobee, Inc., 2018). Specifically, valuable information lies in the 151k thermostat interactions in 1,181 homes where only one occupant is home. This dataset was collected from users of ecobee thermostats who voluntarily selected the ecobee app to "donate your data to science." Self-reported single occupant homes were selected to remove the effect of social factors on thermal comfort behavior. Thermostat's and remote occupancy sensors were used to classify if the user was home, thus experiencing the indoor temperature in a way that may result in an override.

The DyD dataset provides the following data relevant to this study at 5-min. intervals, over multiple years for some users: indoor temperature(s), outdoor temperature, cooling setpoint, the duty-cycle of each piece of cooling equipment (i.e., the number of seconds over each 5-min. interval the relay to the equipment was turned on), PIR motion sensor(s), the event that caused the setpoint change (e.g. manual, scheduled, or a smart event), and the timestamp of each data record.

Two filtering methods were applied to the data. (1) When the setpoint is changed multiple times, each within 5-min. of the previous, the setpoint is averaged over the interval, conflating the actual setpoints. To estimate the actual setpoints, such setpoint measurements were de-averaged assuming the change occurred in the middle of the 5-min interval. (2) PIR occupancy sensors have a 3% false negative rate, yet only a 33% true positive rate. It is important in this study to only include data when the occupant is home and exposed to the indoor environment and available to override the DR event, while still maximizing the amount of data. As such, the method from (Pedersen et al., 2017) was employed to fill in gaps of negative occupancy data less than 30-min. long, assuming the occupant was just temporarily occluded from the PIR sensor's view.

Assumptions

The analysis methods used to answer the research questions are based on the following assumptions that simplify and define the scope of this study. **(A1)** Occupant behavior for any setpoint change of a given setback is equivalent to a DR event with the same setback, contextual, and personal factors. **(A2)** Pre-cooling and recovery of indoor temperature are not considered. **(A3)** Equipment runtime is proportional to energy consumption. **(A4)** A two-stage cooling system is assumed to consume twice the peak energy as a single stage. **(A5)** Thermal loads from solar, occupants, and internal loads are negligible. **(A6)** The indoor temperature is at equilibrium before each setpoint change. **(A7)** The setpoint remains constant during a DR event, and returns to the previous setpoint at the end, unless overridden, in which case the setpoint is set to the override setpoint. **(A8)** Per A6, the system is at equilibrium when a setpoint increases, causing the cooling equipment to remain off for a transient period until the steady state of the new setpoint is reached, and vice versa, a setpoint decreases causes the equipment to remain on at full power until equilibrium is reached. **(A9)** The number of customers available to a DR program is fixed, and the program wishes to maximize the amount of curtailment they can provide to the grid, thus they wish to maximize the curtailment provided by each user.

Feature Extraction

Assumption (A1) enables a study of users' responses to DR events by analyzing their responses to any setpoint changes. The first step toward such an analysis is identifying all SCs in the dataset where an occupant is present and detected by a motion sensor for at least one sample within the length of a maximum DR event (i.e., 4 hrs. in this study). In consideration of summer DR events, the system should be in cooling mode during the SC and the SC must increase the setpoint. By (A6) if the equipment is not running in the period before the SC, then a positive SC will not result in energy savings. Therefore, only SCs with equipment running in the period before the SC (1 hr. in this study) should be considered. Finally, SCs with bad data should not be considered, that is ones with missing values or gaps in the time record.

For each of these SCs, the following features were extracted to enable analysis of curtailment and overrides: the outdoor $\left(T_{O_{SC}}, T_{O_{prev}}, T_{O_{MSC}}\right)$, indoor $\left(T_{I_{SC}}, T_{I_{prev}}, T_{I_{MSC}}\right)$, and setpoint temperature $\left(T_{SP_{SC}}, T_{SP_{prev}}, T_{SP_{MSC}}\right)$ respectively at the time of the SC, the time of the previous SC, and the time of the following manual setpoint change (MSC), if one occurs within 4 hrs. Furthermore, the time from the SC to the MSC is recorded as ΔT_{MSC}, if it occurs.

Additionally, vectors of the following data, sampled every 5-min. from 1 hr. before the SC to 4 hrs. after the SC, were extracted to fit thermal models of each home: the outdoor, indoor, and setpoint temperatures along with the duty cycle of all cooling equipment.

Two conditions were set to qualify a SC valid for analysis. First being, the HVAC equipment should run for at least 24 minutes to know that the system is trying to reach its steady state instead of already being in a steady state. Second, missing data SC's should not be included in the filtered data. Instances where the data was lost or not sampled due to power-outage, network issues, etc. would result in incorrect computation of the systems energy use and hence, SC's with missing data should be ignored.

To train the neural network, the following additional environmental, contextual, and personal variables were extracted and calculated. Environmental variables included the indoor and outdoor relative humidity at the current SC and the previous. Contextual variables included the time between the current SC and the previous, the minute of the day (from midnight), the day of the week, and the month. Personal variables included the percent of time occupied in the hour before the SC, as well as the following non-time-varying personal variables: the percent of times the user increased the thermostat vs. lowered, the harmonic mean of the user's cooling setpoint increases, the percent of time the user's time to override from the SC was less than 10 min. (overall in the data these delays are bimodal, with many happening nearly instantaneously and another peak around 30 min. that depends on the amount of setback), the user's mean time to override when greater than 10 min., the user's mean number of MSCs per year, the user's mode time of day for MSCs during weekdays, and mean time of day during weekends.

From the 5 hrs. of 5-min. sampled indoor and outdoor temperatures and duty cycles for each SC_i, the thermal dynamics of each home can be modeled. The primary factors that govern indoor temperature T_{in} are approximately related by the first order differential equation (1): the thermal mass, mc_p, the equivalent insulation of the envelope $(RA)_{eq}$, the outdoor temperature T_{out}, and the heat from the HVAC, $\dot{Q}_{HVAC}$.

$$mc_p\dot{T}_{in} = (RA)_{eq}(T_{out} - T_{in}) + \dot{Q}_{HVAC} \qquad (1)$$

Following (1), a two parameter model (2) can be fit to the data: $\beta_R^{(i)}$ represents the building insulation divided by thermal mass, $\beta_Q^{(i)}$ represents the maximum heat remove rate of all equipment, and U is the duty cycle of all equipment, per assumption (A4).

$$\dot{T}_I(t) = -\beta_R^{(i)}T_I(t) + \beta_Q^{(i)}U(t) + \beta_R^{(i)}T_O(t) \qquad (2)$$

The β parameters can be fit to the data for each timestep, from 1 through K=300 (5 min. sampling over 5 hrs.), for each SC by reformulating this into a pseudo inverse problem in discrete time (3). A separate β model is created for each SC as opposed to a single model for each home since these properties may change throughout a year, e.g., windows open or increased infiltration due to wind.

$$\begin{bmatrix} T_O^{(2)} - T_I^{(2)} & U^{(2)} \\ \vdots & \vdots \\ T_O^{(K)} - T_I^{(K)} & U^{(K)} \end{bmatrix} \begin{bmatrix} \beta_R^{(i)} \\ \beta_Q^{(i)} \end{bmatrix} = \begin{bmatrix} T_{in}^{(2)} - T_{in}^{(1)} \\ \vdots \\ T_{in}^{(K)} - T_{in}^{(K-1)} \end{bmatrix} \qquad (3)$$

From this model, the energy use, i.e., duty cycle per (A3), and dynamics under different scenarios can be estimated. The steady state duty cycle before each SC_i, $\overline{U}_{prev}^{(i)}$ is estimated by (4) per (A6). Similarly, the estimate for the duty cycle after the SC once equilibrium is reached $\overline{U}_{SC}^{(i)}$ follows (5), and the equilibrium duty cycle after the MSC $\overline{U}_{MSC}^{(i)}$ follows (6). The duration of the SC transient $\Delta t_{tr_{SC}}^{(i)}$ until steady state is reached follows (7) and the transient time from the MSC to steady state is $\Delta t_{tr_{MSC}}^{(i)}$, following (8).

$$\overline{U}_{prev}^{(\mathrm{i})} = \left(\frac{\beta_R^{(\mathrm{i})}}{\beta_Q^{(\mathrm{i})}}\right)\left(T_{sp_{prev}}^{(\mathrm{i})} - T_{out_{prev}}^{(\mathrm{i})}\right) \qquad (4)$$

$$\overline{U}_{SC}^{(\mathrm{i})} = \left(\frac{\beta_R^{(\mathrm{i})}}{\beta_Q^{(\mathrm{i})}}\right)\left(T_{sp_{SC}}^{(\mathrm{i})} - T_{out_{SC}}^{(\mathrm{i})}\right) \qquad (5)$$

$$\bar{U}_{MSC}^{(\mathrm{i})} = \left(\frac{\beta_R^{(\mathrm{i})}}{\beta_Q^{(\mathrm{i})}}\right)\left(T_{spMSC}^{(\mathrm{i})} - T_{outMSC}^{(\mathrm{i})}\right) \quad (6)$$

$$\Delta t_{trSC}^{(\mathrm{i})} = \left(-\beta_R^{(\mathrm{i})}\right)^{-1} \ln\left(\frac{T_{SP_{SC}}^{(\mathrm{i})} - T_{O_{SC}}^{(\mathrm{i})}}{T_{I_{SC}}^{(\mathrm{i})} - T_{O_{SC}}^{(\mathrm{i})}}\right) \quad (7)$$

$$\Delta t_{trMSC}^{(\mathrm{i})} = \left(-\beta_R^{(\mathrm{i})}\right)^{-1} \ln\left(\frac{T_{SP_{MSC}}^{(\mathrm{i})} - T_{O_{MSC}}^{(\mathrm{i})}}{T_{I_{MSC}}^{(\mathrm{i})} - T_{O_{MSC}}^{(\mathrm{i})}}\right) \quad (8)$$

Note that in some cases the values of the variables β, $\bar{U}$, and Δt may be negative, not a number, or infinity. In which case, the SC is removed from any study requiring these parameters.

Automated Curtailment (Q1)

An estimate for the baseline energy $E_{BL}^{(i)}$ over a DR event of length Δt_{DR} as if each SC_i never occurred is computed by (9) assuming steady state operation.

$$E_{BL}^{(i)}(\Delta t_{DR}) = \bar{U}_{prev}^{(\mathrm{i})} \cdot \Delta t_{DR} \quad (9)$$

The energy consumption that accounts for the SCs setback $E_{SC}^{(i)}$ can then be similarly estimated by the two phases of curtailment defined by (A8). Zero energy is consumed during the transient phase of length $\Delta t_{trSC}^{(i)}$. The steady state phase of length $\Delta t_{SS_{SC}}^{(i)}$, if one occurs, is the time remaining in the DR event after the transient phase.

$$E_{SC}^{(i)}(\Delta t_{DR}) = \bar{U}_{SC}^{(\mathrm{i})} \cdot \Delta t_{SS_{SC}}^{(i)} \quad (10)$$

The aggregate curtailment for a given integer setback ΔT_{SC} and duration Δt_{DR} is computed by first binning the SCs by their nearest integer of $\Delta T_{SC}^{(i)} = \left(T_{SP_{SC}}^{(i)} - T_{SP_{prev}}^{(i)}\right)$, noting that some setbacks may not be integer if operating in Celsius where thermostat settings are in $\pm 0.5°C$. The aggregate percent curtailment $C_{SC}(\Delta T_{SC}, \Delta t_{DR})$ for a given setback and duration across the population then follows (11).

$$C_{SC}(\Delta T_{SC}, \Delta t_{DR}) = 1 - \sum_{i}^{\Delta T_{SC}^{(i)} \in [\Delta T_{SC} \pm 0.5]} \frac{E_{SC}^{(i)}(\Delta t_{DR})}{E_{BL}^{(i)}(\Delta t_{DR})} \quad (11)$$

A plot of this curtailment versus setback and duration will provide insights into the balance of the two phases of curtailment.

Curtailment with Overrides (Q2 & Q3)

The previous curtailment study neglected the effect of occupant overrides. To account for these MSCs, up to four phases of energy consumption should be considered: the transient period $\Delta t_{trSC}^{(i)}$ after the SC with no energy consumed, the steady state period $\Delta t_{SS_{SC}}^{(i)}$ after the SC, the transient period of full power cooling $\Delta t_{trMSC}^{(i)}$ after the MSC, and the steady state period $\Delta t_{SS_{MSC}}^{(i)}$ after the MSC until the end of the DR event. The MSC can preempt the SC reaching steady state. Similarly, the end of the DR event can preempt the MSC transient, the MSC event, and the SC transient. The variables for these preempted periods are shortened or set to zero accordingly. The total energy consumed over the length of the DR for each SC_i then follows (12).

$$E_{MSC}^{(i)}(\Delta t_{DR}) = \bar{U}_{SC}^{(i)} \, \Delta t_{SS_{SS}}^{(i)} + 1 \, \Delta t_{trMSC}^{(i)} + \bar{U}_{MSC}^{(i)} \, \Delta t_{SS_{MSC}}^{(i)} \quad (12)$$

The aggregate curtailment accounting for the MSC for a given integer setback, binned as above, and duration follows (13). A DR program may deploy a uniform Δt_{DR} and ΔT_{SC} and anticipate no overrides yielding a curtailment of C_{SC}. However, overrides would occur, yielding C_{MSC}, and thus the DR programs curtailment error would be e_{MSC} defined by (14).

$$C_{MSC}(\Delta T_{SC}, \Delta t_{DR}) = 1 - \sum_{i}^{\Delta T_{SC}^{(i)} \in [\Delta T_{SC} \pm 0.5]} \frac{E_{MSC}^{(i)}(\Delta t_{DR})}{E_{BL}^{(i)}(\Delta t_{DR})} \quad (13)$$

$$e_{MSC}(\Delta T_{SC}, \Delta t_{DC}) = 1 - \frac{C_{MSC}(\Delta T_{SC}, \Delta t_{DR})}{C_{SC}(\Delta T_{SC}, \Delta t_{DR})} \quad (14)$$

Plotting these curtailments and errors versus setback and duration will provide insights into the effect of occupants overrides and help inform methods for reducing occupant dissatisfaction while maximizing curtailment.

Predicting Overrides

There are two key prediction tasks to consider when predicting the impact of MSCs on DR events. First is to predict if an MSC will occur during the DR event, i.e., before Δt_{DR}. The second, if an MSC is predicted to occur, predict when will it occur $\widehat{\Delta t}_{MSC}$. Two neural network (NN) models were developed for these tasks.

The first model, NN_1 consist of 28 neurons in first layer, corresponding to the 28 features, and one output neuron that returns a value between zero and one with the aim to discriminate between MSC vs. no MSC during the DR event. The second model, NN_2 also consists of 28 neurons but assumes that an MSC will occur and aims to estimate $\widehat{\Delta t}_{MSC}$.

First, NN_1 was trained with 75% of all the SCs, to predict which contained MSCs. Then, NN_2 was trained with just

75% of the data containing MSCs to predict when they occurred. Once trained, the two NNs were combined such that data is fed into NN_1, and if it predicts an MSC, the data is then fed into NN_2 to predict when it will occur.

The remaining 25% of the data was used for validation. If features were missing or nonsensical (e.g., negative β_R or complex Δt) the bad values of that feature were replaced with the mean of the good values in the rest of the dataset.

Personalized Curtailment (Q4)

The method for analyzing the ability of personalized and context-aware DR (PDR) durations $\widehat{\Delta t}_{PDR}^{(i)}$ to reduce curtailment error while maximizing curtailment are similar to the methods used for non-personalized DR in Q2 and Q3. If the predicted $\widehat{\Delta t}_{DR}^{(i)}$ is conservative, i.e., less than the actual $\Delta t_{MSC}^{(i)}$, then the actual DR duration $\Delta t_{DR}^{(i)} = \widehat{\Delta t}_{DR}^{(i)}$ and the predicted energy consumption $\hat{E}_{PDR}^{(i)}$ is equal to the true energy $E_{PDR}^{(i)}$, yet less than $E_{MSC}^{(i)}$. This energy is calculated by (15), with zero energy during the first transient period $\Delta t_{tr_{SC}}^{(i)}$, followed by the energy saving steady state power $\overline{U}_{SC}^{(i)}$ for the period $\Delta t_{SS_{SC}}^{(i)}$ followed by the transient recovery period $\Delta t_{tr_{PDR}}^{(i)}$ as defined by (16) from the indoor temperature $T_I^{(i)}(\widehat{\Delta t}_{DR}^{(i)})$ to the previous setpoint after the DR. This is followed by the remaining steady state period $\Delta t_{SS_{PDR}}^{(i)}$ when $\overline{U}_{prev}^{(i)}$ energy is consumed until the end of the maximum length of the DR event Δt_{DR}. As before, any of the Δt terms may be preempted.

$$\hat{E}_{PDR}^{(i)}(\Delta t_{DR}, \widehat{\Delta t}_{PDR}) = \overline{U}_{SC}^{(i)}\, \Delta t_{SS_{SC}}^{(i)} + 1\, \Delta t_{tr_{PDR}}^{(i)} + \overline{U}_{prev}^{(i)}\, \Delta t_{SS_{PDR}}^{(i)} \tag{15}$$

$$\Delta t_{tr_{PDR}}^{(i)} = \left(-\beta_R^{(i)}\right)^{-1} \ln\left(\frac{T_{SP_{SC}}^{(i)} - T_{O_{SC}}^{(i)} - \dfrac{\beta_Q^{(i)}}{\beta_R^{(i)}}}{T_{I_{SC}}^{(i)} - T_{O_{SC}}^{(i)} - \dfrac{\beta_Q^{(i)}}{\beta_R^{(i)}}}\right) \tag{16}$$

On the other hand, if $\widehat{\Delta t}_{PDR}^{(i)} > \Delta t_{MSC}^{(i)}$, then an override occurs and $\hat{E}_{PDR}^{(i)} > E_{PDR}^{(i)} = E_{MSC}^{(i)}$. The aggregate curtailment for the personalized DR with overrides is C_{PDR} and is defined by (17), while the predicted estimate of the PDR, assuming no overrides is $\hat{C}_{PDR}$ and follows (18). The curtailment error that a DR program would realize with this strategy is thus e_{MSC} as in (19).

$$C_{PDR}(\Delta T_{SC}, \Delta t_{DR}, \widehat{\Delta t}_{PDR}) = 1 - \sum_{i}^{\Delta T_{SC}^{(i)} \in [\Delta T_{SC} \pm 0.5]} \frac{E_{PDR}^{(i)}(\Delta t_{DR}, \widehat{\Delta t}_{PDR})}{E_{BL}^{(i)}(\Delta t_{DR})} \tag{17}$$

$$\hat{C}_{PDR}(\Delta T_{SC}, \Delta t_{DR}, \widehat{\Delta t}_{PDR}) = 1 - \sum_{i}^{\Delta T_{SC}^{(i)} \in [\Delta T_{SC} \pm 0.5]} \frac{\hat{E}_{PDR}^{(i)}(\Delta t_{DR}, \widehat{\Delta t}_{PDR})}{E_{BL}^{(i)}(\Delta t_{DR})} \tag{18}$$

$$e_{PDR}(\Delta T_{SC}, \Delta t_{DC}, \widehat{\Delta t}_{PDR}) = 1 - \frac{C_{PDR}(\Delta T_{SC}, \Delta t_{DR}, \widehat{\Delta t}_{PDR})}{\hat{C}_{PDR}(\Delta T_{SC}, \Delta t_{DR}, \widehat{\Delta t}_{PDR})} \tag{19}$$

Comparing C_{PDR} to C_{MSC} will yield insights into the increase or decrease in curtailment through personalized controls, while comparing e_{PDR} to e_{MSC} yields the improvement in prediction error.

RESULTS AND DISCUSSION

Dataset Curation and Feature Extraction

The DyD dataset of occupied energy-conserving cooling SCs in single occupancy homes yielded 151k SCs. For each of these SCs, 20 contextual features and 11 personal features were extracted to predict building energy consumption and when, and if, overrides would occur. Some of these features were missing for certain SCs, resulting in only 127k SCs fit for analysis. A further 8 features were computed characterizing the building physics, energy use, and transient periods for each SC. In some cases, these computed values were not usable (e.g. NAN), resulting in only ~1k SCs with every feature. However, not every feature is required for every analysis, so the maximum amount of data is used for each analysis.

Automated Curtailment (Q1)

The impact of the two phases of energy consumption can be clearly seen in Figure 1 plotting the curtailmentC_{SC}, without any MSCs, for different setbacks and durations of the DR event. Larger setbacks mean the cooling remains off while the indoor temperature climbs to the new setpoint. For example, a 10°F setback curtails the energy consumption by nearly 100% for the first 10-min. Gradually over time for every setback, the energy use increases as the HVAC system reaches steady state.

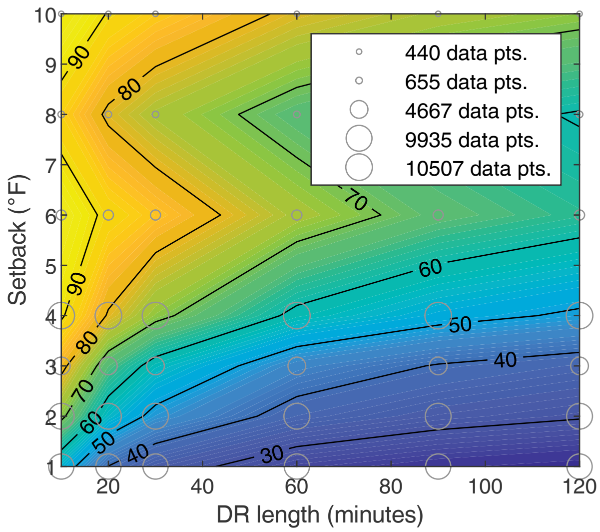

Figure 1: Energy curtailment C_{SC} without MSCs (The size of grey circles shows the number of data points used and the lines denote the respective percentages of curtailment)

Curtailment with Overrides (Q2 & Q3)

Manual overrides reveal discomfort as the occupant adjusts the setpoints to meet their thermal environment needs. The adjustments in discomfort may be made without consideration of its energy impact. Figure 2 shows the percent energy curtailed C_{MSC} with setbacks for various time frames. DR events of less than 10-min. with setbacks greater than 4°F result in ~80% curtailment. Smaller temperature setbacks increased energy use over baseline for DR events greater 30-min. This could be indicative of users overcorrecting for their discomfort, decreasing the setpoint more than the SC increased it

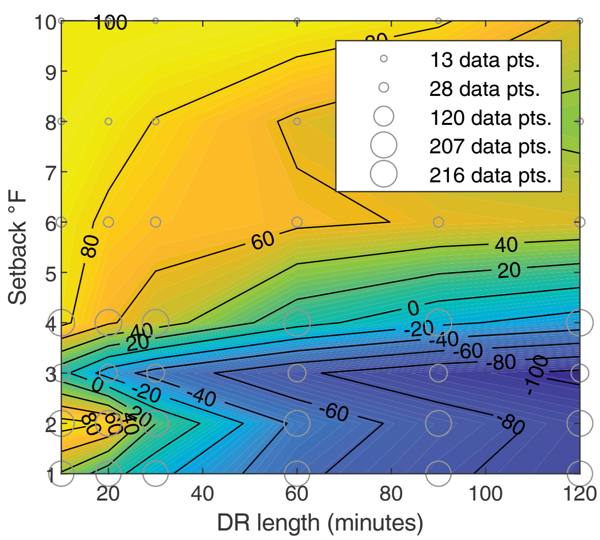

Figure 2: Energy curtailment with MSCs, C_{MSC} (The size of grey circles shows the number of data points used and the lines denote the respective percentages of curtailment)

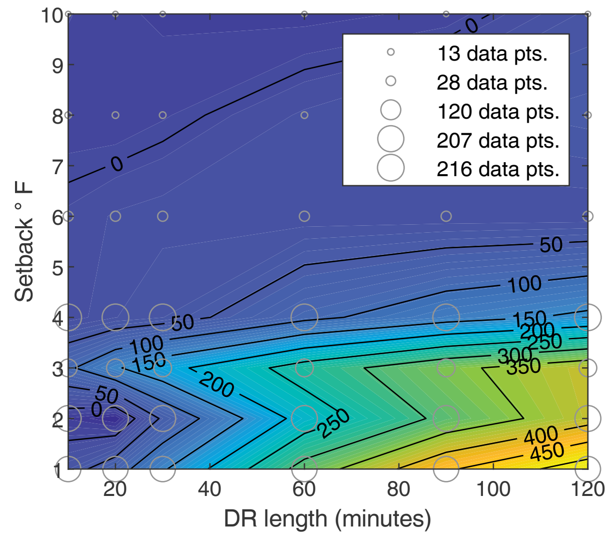

Figure 3: Energy curtailment error e_{MSC} from overrides (The size of grey circles shows the number of data points used and the lines denote the respective percentages of curtailment)

A DR program may consider using Figure 1 to predict curtailment for a given setback and duration. It is worth noting that DR programs should aim for greater than 25% accuracy in their curtailments. If this occurred, Figure 3 shows when overrides are permitted, this level of accuracy is achieved for 2°F setbacks less than 20min and for setbacks greater than 7°F. It should be noted that the data is limited for the latter case, only <1k points. For setbacks less than 4°F and greater than 30-min., 500% errors occur due to the large number of overrides over these long horizons.

Predicting Overrides

Two neural networks with 29 neurons each were used to classify manual overrides and predict the corresponding time to override after setpoint change. The first neural network or the pattern recognition neural network, resulted in 96% accuracy in classifying a manual override. 75% of the data (96k data points), composed of both SC with and without MSCs, was used to train the first network. The accuracy of the network is as follows: False Positives: 2%, False negatives 1.4%, True Positives 15.8% and True Negatives 80%.

The predicted override's data was used to predict the time an occupant might take to override the setback. Similar to the first network, 28 hidden neurons and 1 output neuron were used. The Root Mean Square Error (RMSE) of the NN model was found to be 48-min, admittedly not great but potentially better than no estimator. Figure 4, shows the 2d histogram of predicted

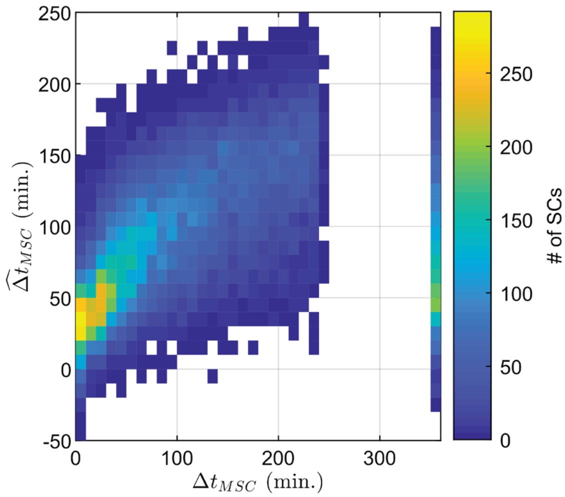

Figure 4: Histogram of actual time to override Δt_{MSC} and value predicted by NN, $\Delta\hat{t}_{MSC}$

and true values. No-overrides data points were assigned the value of 360 to visually show the misclassification of first neural net.

Personalized Curtailment (Q4)

Figure 5 shows the percent energy curtailed when the DR event length is limited to the time to override predicted by the two neural networks. The stricter assumptions for these calculations significantly reduced the amount of data available from ~100k down to ~1k.

DR event durations less than 20 min led to >80% curtailment for setbacks up to 4 °F. For higher setbacks, fewer data points led to inconclusive results due to limited data. As before, when DR events last longer than 20-min. curtailment deteriorates; although, not as low as even the baseline case.

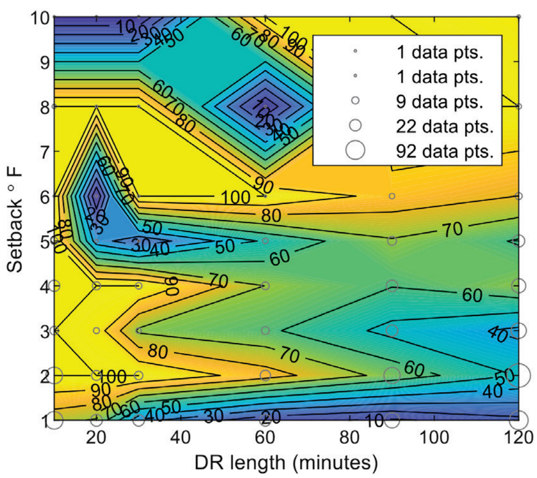

Figure 5: Energy curtailment C_{PDR} from personalized DR (The size of grey circles shows the number of data point used and the lines denote the respective percentages of curtailment)

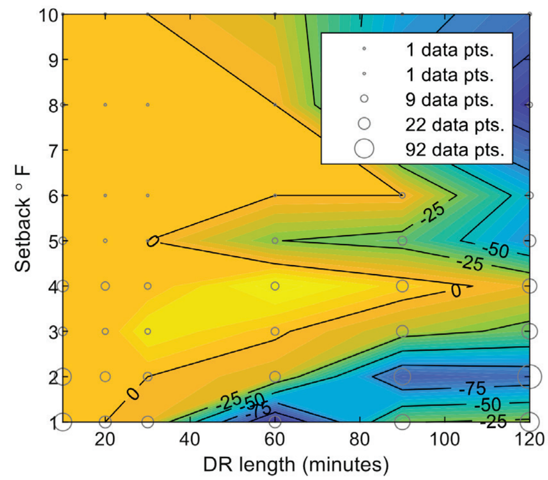

Figure 6: Energy curtailment error e_{PDR} from the personalized DR (The size of grey circles shows the number of data point used and the lines denote the respective percentages of curtailment.)

Comparing Figure 5 with Figure 6, desired errors of less than 25% are achieved for nearly all events shorter than 30-min. A significant improvement over the DR approach without personalization. However, longer events result in greater error. This is to be expected as modeling errors integrate.

SUMMARY AND CONCLUSION

This study's goal was to assess the impact of replacing the current approach of demand response, which uses uniform DR event durations for an entire community with adaptive and personalized setbacks. From a filtered dataset of about 151k data points, simulated as demand response events, our analysis answered 4 main questions.

First, how much do setback duration and magnitude effect energy curtailment during DR events? The results show curtailment increases with increasing setback. The effect of increasing setback diminishes with increasing duration, especially after ~20 minutes due to the device being off until the new setpoint is reached naturally.

Second, how do overrides affect DR events? Since these calculations regard estimation of building thermal parameters with stricter assumptions, much less data was available, increasing uncertainty in the analysis. Increasing setpoints, increases curtailment, with some exceptions. As the duration surpassed 20 minutes, the curtailment for setbacks less than 4°F went negative, meaning the DR events and subsequent override caused an increase in energy consumption. This change due to overrides raises the third question. How much error can operators expect due to overrides? Errors may exceed 500% , especially for events longer than 30-min. Yet, errors of less than the required 25% are possible for events less than 30-min. with $2°F$ setbacks.

Lastly, NNs can predict override occurrence and timing. The use of NNs to select DR event durations significantly improved curtailments compared to the general case with overrides. In addition to increasing curtailment percent, the errors were brought with to within the acceptable range for many situations.

It is important to note here that the results depicting energy curtailment with and without MSC arise from different dataset's, namely the setpoint changes that resulted in invalid β values or other corrupted parameters were ignored. This led to 100k+ data points in the analysis of Q1, 10k+ data point for Q2 and Q3, and only 1k+ for Q4. Thus, it is possible that the down-selected data could have different relevant characteristics that are unaccounted for. Even with potential data bias, this preliminary analysis demonstrates the potential for personalized demand response and occupant-aware grid-interactive efficient buildings and motivates a great deal more study in this area.

ACKNOWLEDGMENT

This work was conducted within the framework of IEA-EBC Annex 79. The authors would like to acknowledge their gratitude to ecobee's Social Impact Department for curating and sharing the Donate Your Data dataset that is foundational to this research.

REFERENCES

Almassalkhi, M., et al., 2018. Asynchronous Coordination of Distributed Energy Resources with Packetized Energy Management, in: Energy Markets and Responsive Grids: Modeling, Control, and Optimization. Springer, New York, NY, pp. 333–361.

Cappers, P., MacDonald, J., et al., 2016. Future opportunities and challenges with using demand response as a resource in distribution system operation and planning activities.

Chen, X., Wang, Q., Srebric, J., 2015. A data-driven state-space model of indoor thermal sensation using occupant feedback for low-energy buildings. Energy Build. 91, 187–198.

Ecobee, Inc., 2018. Donate your Data. Smart WiFi Thermostats Ecobee. URL https://www.ecobee.com/donateyourdata/.

Fanger, P.O., 1973. Assessment of man's thermal comfort in practice. Br. J. Ind. Med. 30, 313–324.

Hledik, R., Faruqui, A., et al., 2019. Brattle Study: Cost-Effective Load Flexibility Can Reduce Costs by More Than $15 Billion Annually.

Jin, X., Baker, K., et al., 2017. Foresee: A user-centric home energy management system for energy efficiency and demand response. Appl. Energy 205, 1583–1595.

Kane, M., 2018. Modeling Human-in-the-Loop Behavior and Interactions with HVAC Systems, Presented at the 2018 Annual ACC, IEEE, Milwaukee, WI, pp. 4628–4633.

Kim, J., Schiavon, S., Brager, G., 2018. Personal comfort models – A new paradigm in thermal comfort for occupant-centric environmental control. Build. Environ. 132, 114–124.

Marie-Andrée, L., et al., 2011. Developing Winter Residential Demand Response Strategies for Electric Space Heating, Presented at the 12th Conference of International Building Performance Simulation Association, Sydney, AU, p. 8.

Mirakhorli, A., Dong, B., 2016. Occupancy behavior based model predictive control for building indoor climate—A critical review. Energy Build. 129, 499–513.

Neukomm, M., Nubbe, V., Fares, R., 2019. Grid-interactive Efficient Buildings: Overview. U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Washington, D.C.

Pedersen, T.H., Nielsen, K.U., Petersen, S., 2017. Method for room occupancy detection based on trajectory of indoor climate sensor data. Build. Environ. 115, 147–156.

PJM, 2018. Load Management Performance Report – 2017/2018. Norristown, PA.

Schweiker, M., Ampatzi, E., et al., 2020. Review of multi-domain approaches to indoor environmental perception and behaviour. Build. Environ. 176, 106804.

Seiden, K., Olig, C., et al., 2017. National Grid Smart Energy Solutions Pilot - Final Evaluation Report, DPU 10-82. Massachusetts Electric Company and Nantucket Electric Company d/b/a National Grid, Worcester, MA.

Singla, S., Keshav, S., 2012. Demand response through a temperature setpoint market in Ontario, in: 2012 SmartGridComm. IEEE, Taiwan, pp. 103–108.

U.S. Energy Information Administration, 2015. 2015 RECS: Overview. URL https://www.eia.gov/consumption/residential/reports/2015/overview/.

Zhang, H., Arens, E., et al., 2010. Thermal sensation and comfort models for non-uniform and transient environments: Part I: Local sensation of individual body parts. Build. Environ., 1st International Symposium on Sustainable Healthy Buildings 45, 380–388.