# Data-driven Modeling of Adaptive Occupant Thermostat Behavior Dynamics

**Michael Kane, PhD**
*Member ASHRAE*

**Kunind Sharma**
*Student Member ASHRAE*

## ABSTRACT HEADING

*In the U.S., building HVAC systems consume nearly half of the electricity produced. Demand response (DR) programs aim to reduce the peaks of this demand, thereby reducing emissions from inefficient 'peaking' power plants and avoiding costly infrastructure improvements to serve increasing peak demand. The future of such load flexibility programs, grid-interactive efficient buildings (GEBs), will not only shed but shift and modulate the load to provide higher fidelity and faster grid services, such as avoiding renewable curtailment, ramping services, and frequency regulation.*

*Current DR Implementations are primitive: emails, texts, phone calls, and direct load control lack fidelity required to achieve accurate and reliabile load shedding, shifting, and modulation. Even when smart devices are used, the same settings are used for all participants, regardless of their sensitivities to reduced quality of service and economics. Such strategies result in occupant overrides, over 30% of 8-hour DR events in one study, resulting in loss of consumer trust and millions of dollars in penalties nationwide. This drives a need to personalize DR controls to maximize demand flexibility and reliability.*

*This paper presents a preliminary exploration of approaches to develop personalized predictive models of manual overrides of thermostat setpoints, from sparse data on individuals yet a plethora of users. These data come from the ecobee Donate Your Data dataset with ~1259 users and ~285 events per user. The study's methods explore the effectiveness of decision trees and artificial neural network on various features, with a focus of looking at the dynamic nature of this behavior in the context of GEBs. These personalized modeling approaches are compared in terms of their accuracy, computational complexity, and outlier management.*

## INTRODUCTION

The electric grid that provides 40% of energy consumed by buildings is increasingly vulnerable. New stresses on the grid are arising from a growing demand for low global-warming potential energy, and the resulting electrification of space heating and transportation. Increasing variable renewable electricity generation combined with transmission and distribution (T&D) infrastructure constraints exacerbate these issues. To alleviate peak demand and to help balance supply and demand on the grid, utilities are increasing looking to control the load on the grid (Neukomm et al., 2019).

Currently, building load flexibility is primarily leveraged through demand response (DR) programs. Some of these programs temporarily reduce building energy use by increasing the air conditioning thermostat setpoint during peak events on the grid. While such programs can reduce peak loads by ~30%, costumer quality of service is significantly affected. In one study, upto 30% of participants overrode or opted-out of DR events that lasted for as long as 8-hours

**Michael Kane** is an Assistant Professor in the Department of Civil and Environmental Engineering, Northeastern University, Boston, Massachusetts. **Kunind Sharma** is a graduate research assistant at Northeastern University, Boston, Massachusetts.

(Seiden et al., 2017). The discomfort from these energy saving setpoint changes is exacerbated for customers that are already trying to save energy — operating at the edge of their comfort zone (Cappers et al., 2016). In addition to upsetting customers, these overrides cost utilities millions of dollars per year in penalties for not providing the promised load curtailment to the wholesale grid marketplace (PJM Interconnection LLC, 2018).

DR programs that aim to maximize flexibility *and* reliability must therefore choose thermostat setback magnitudes and durations which are customized for the user and adapted to the context. The most common personalized model, *the adaptive model of thermal comfort* (de Dear and Schiller Brager, 2001), is insufficient for this purpose as it depends on a steady-state physiologic model of comfort that is adapted for slowly varying context (e.g., geographic location, season) and provides no information as to how long an occupant will tolerate a temperature. Only two *dynamic* thermal comfort models have been identified in recent reviews of personalized thermal comfort (Kim et al., 2018; Mirakhorli and Dong, 2016). These existing dynamic models were trained in laboratory settings with only 13 (Chen et al., 2015) and 109 (Zhang et al., 2010) subjects, insufficient for capturing the breadth of human diversity.

This paper aims to use data from 1,259 thermostat to understand if machine learning (ML) models can capture the dynamics of occupant behavior, adapting to personal and contextual factors. The Methods section introduces the data used for this study and the preprocessing techniques used to extract relevant data features. Additionally, two ML approaches, decision trees and neural networks, used to build adaptive personalized behavior models are explained. The Results and Discussion section presents the resulting behavior models including their accuracy and the impact of model parameters and data features. This manuscript closes with a summary and conclusions relevant to realizing these methods in industry and defining future work in research on occupant centric controls and GEBs.

## METHODS

The ecobee Donate your Data (DyD) dataset is used to train the proposed ML models of occupant thermostat-use behavior. Before training models with the DyD data, it is first preprocessed, and relevant features are extracted. A framework is developed for analyzing occupant thermal comfort behavior dynamics, yielding an initial population level model. Then, two ML methods, decision trees and artificial neural network (ANN), are proposed for their ability to predict behavior based on contextual and personal factors. Assessments are defined to compare these algorithms in terms of accuracy, reliability, and computational efficiency.

### Dataset Curation and Pre-processing

The DyD dataset consists of years of data from ~70k thermostats across the world (Ecobee, Inc., 2018), sampled every 5-min. These data are provided by ecobee customers as part of an opt-in program, where the data is provided to researchers anonymized to the city level. The data variables relevant to this study are show in Table 1. This study considers the 1,259 homes which are self-reported as have only a single occupant. This data subset enables this study to focus on personal drivers of thermal comfort behavior dynamics, removing confounding effects of social factors.

The data was preprocessed to identify the units (i.e., $°F$ or $°C$) used by the occupants, correct setpoint sampling errors, and reduce noise and bias in the passive infrared (PIR) occupancy sensor data. All the data is reported in $°F$; however, occupants choose to see and change their thermostat in either Fahrenheit or Celsius, in increments of $1°F$ or $0.5°C$ respectively. The setpoint change data, modulo these increments are used to determine the users working units.

Users can manually change the setpoint at any time, either through the thermostats interface or through a smartphone app, yet the data are only sampled every 5-min. The resulting data are the average setpoint over the sampling

**Table 1. Variables in the DyD Dataset Relevant to this Study**

| | | | |
|---|---|---|---|
| Time stamp | Cooling/heating setpoint | Duty-cycle of equipment | Setpoint temperature |
| PIR motion sensor data | Indoor temperature | Outdoor temperature | Setpoint change events |

**Setpoint Change Event Types**

| | | | |
|---|---|---|---|
| Demand response | Manual | Scheduled: home, away, … | Smart: geofence, PIR, … |

**Table 2. Data Features Used in Model and Analysis**

| **Variable Name:** Description | **Variable Name:** Description |
|---|---|
| **minutes/ weekday/month**: min. from midnight/day of week/month of yr. when the SC occurred | **usr_heat/cool_pct_up**: Percent times user increases thermostat heating/cooling (H/C) setpoint |
| **temp/humIn/Out**: Indoor and outdoor temperature and relative humidity | **usr_heat/cool_up_mean**: User's harmonic mean value of H/C setpoint decreases |
| **prevOccDt**: Duration occupancy was continuously detected backwards from the current SC | **usr_heat/cool_down_mean**: User's harmonic mean value of H/C setpoint decreases |
| **prevEvent**: Event at previous SC | **usr_heat/cool_ttd_pct_inst**: Percent times user's H/C MSCs occurred within 10min of the previous SC |
| **prevT_ctrl/out**: Average indoor and outdoor temperature at previous SC | **usr_heat/cool_ttd_delay**: User's mean time from the H/C SC to the MSC, for durations >15 min. |
| **prevT_stp_heat/cool**: H/C setpoint at previous SC | **usr_avg_MSC_yr**: User's mean of number of overrides in a year |
| **HVACmode**: Is the equipment operating in heating or cooling mode | **usr_mode_TOD_wDay**: Mode time of day of the user's MSC during the weekday |
| **DOD_est**: Estimated degree of discomfort, the change from the previous setpoint to the current | **usr_mode_TOD_wEnd**: Mode time of day of the users' MSCs during the weekend |
| **DOD**: Degree of discomfort, the change from the current setpoint to the one after the following MSC | **TTD**: Time to discomfort, the time between the SC and the following MSC |

period, which could result in under, or over-estimates. Therefore, when a single setpoint change (SC) is observed, the following datapoint is projected forward. When multiple SCs are observed in successive sampling periods, the setpoint in the middle increments is estimated assuming the change occurred 2.5-minutes into each sampling period.

Occupancy is measured with anywhere between 1 to 14 PIR sensors; most homes have between 1 to 3. Because parts of the home are likely occluded from the sensors view, true positive detection rates are as low as 33% (Pedersen et al., 2017). To correct for this, the union of all the sensor measurements are filtered to fill in any gaps less than 30-minutes. This method is adapted from (Pedersen et al., 2017), and is based on the heuristic that occupants may be occluded from the sensors for only a short period of time. Regardless, the filtered data is likely biased toward false negatives, leading to the framing of the methods below to minimize the effect of this bias.

## Feature Extraction

Occupant thermostat-use behavior dynamics result from a feedback loop of factors that includes the building physics, the HVAC system, and the person's physiology, cognition, and decision making (D'Oca et al., 2017; Kane, 2018). Inputs into each of these subsystems can be characterized as environmental, contextual, and personal (Schweiker et al., 2018). To study the affect of contextual and personal factors on whether an occupant will override a previous thermostat setpoint change (SC), every SC is exctracted from the DyD dataset. For each SC, the contextual and personal features in Table 2 are calculated. The "`usr_*`" variables are intended to act as proxies for personal factors such as perceived barriers/support, skills/understanding, and internal cues (Kane, 2018).

To understand manual overrides, all such manual setpoint changes (MSCs) were identified in the SC data. Only those MSCs that occurred within 4-hours of a previous SC, and occupied continuously in that period, were considered for further analysis to focus the study on capturing the affect of thermal comfort. It is also assumed that the MSC are intended to make the occupants more comfortable. Therefore, the delay from the SC to the MSC is termed the "time to discomfort" (TTD) and the amount the thermostat is changed is called the "degree of discomfort" (DOD).

## Population Model

For each MSC in the DyD dataset, the TTD can be estimated as the time from the previous SC, and the DOD estimated as the difference between the MSC setpoint and the previous SC setpoint. This acknowledges that the building dynamics affect the amount of time it would take a user to notice the change and eventually override it. Additionally, it assumes that the MSC will return the temperature to one which is comfortable for the user. These values were extracted for every MSC. An exponential function fits these data for eaching season, mapping TTD to DOD or vise versa.

### Decision Tree

The population level model estimates the TTD only considering the DOD, i.e., the magnitude of the preceeding thermostat change. Since thermal comfort behavior may be driven by other contextual and personal factors (Schweiker et al., 2018), context aware and personalized models are explored using a decision tree ML approach. All the variables in Table 2 are considered, first just the contextual, then the value of including the personal factors. First, a decision tree is trained to predict whether or not an MSC occurs. If one occurs, a second tree predicts the TTD, binned into 5-minute discrete categories between 5 to 240 min. Upto 100 nodes are allowed in both trees; otherwise, the models are consistent with the defaults of (Mathworks, 2019). The classification model was trained using 60% of the data with the remaining 40% for final testing. The TTD decision tree was trained and validated using the same 60/40 ratio using only MSC data.

### Neural Network

Thermal comfort behavior is nonlinear (Kane, 2018), which decision trees may not capture well. In contrast, artificial neural networks (ANNs) can fit arbitrary input-output relationships with a sufficiently sized network and diverse dataset. In pursuit of a more accurate TTD predictor, and on the heals of prior success of using ANNs to model steady-sate thermal comfort behavior (Deng and Chen, 2018), the Matlab Deep Learning Toolbox (MathWorks, 2019) can train a shallow ANN for predicting TTD based on contextual and personal factors. Like the decision tree approach, an initial network predicts if an MSC will occur, if so the second network predicts the value of the TTD.

The two ANNs are described by Equation 1 and Equation 2 respectively, with weights **w** and biases **b** optimized to minimize the crossentropy classification error in Equation 1 and the root mean squared error (RMSE) in Equation 2. The initial classification ANN uses two layers, the hidden layer consisting of nodes using the tan-sigmoid transfer function (Equation 3) and a softmax transfer function (Equation 4) in the output node. If the first ANN predicts an MSC within 240 min., the second network re-processes the features with a tan-sigmoid hidden layer and a linear output node to predict a regression on TTD. Both ANNs use a single output node and the same number of hidden layer nodes as features; N, for contextual or contextual+personal features from Table 2 respectively. The models were trained through 5-fold validation, using 60% of the data for training, 20% for validation, and the remaining 20% for testing. The TTD network was trained and validated using the same 60/20/20 ratio using only MSC data.

$$Y = \text{softmax}\left(\boldsymbol{w}_{output} \cdot \sum_{n=1}^{N} \text{tansig}\left(\boldsymbol{w}_{hidden,n}\boldsymbol{X} + \boldsymbol{b}_{hidden,n}\right) + \boldsymbol{b}_{output}\right) \tag{1}$$

$$Y = \boldsymbol{w}_{output} \cdot \sum_{n=1}^{N} \text{tansig}\left(\boldsymbol{w}_{hidden,n}\boldsymbol{X} + \boldsymbol{b}_{hidden,n}\right) + \boldsymbol{b}_{output} \tag{2}$$

$$\text{tansig}(n) = (2 \cdot (1 + \exp(-2n))^{-1}) \tag{3}$$

$$softmax(n) \ = \ exp(n)/sum(exp(n)) \tag{4}$$

### Method Assessments

The RMSE prediction error is used to assess the ability of the exponential population model, tree classifier, and the ANN to predict the TTD for a given MSC. The population model is unable to predict if an MSC will occur; however, both the tree and ANN have an initial classifier to predict if an MSC will occur. For a fair comparison, the RMSE was computed for only the SC that were followed by an MSC (and thus have a valid TTD) and that the tree/ANN predicted would have an MSC. To compare the ability of the first layer of the tree and ANN model to predict if an MSC will occur, the classification accuracy and false/true positive/negative rates are calculated for all the SCs. The affect of including personal features during training on classification and regression accuracy is also considered.

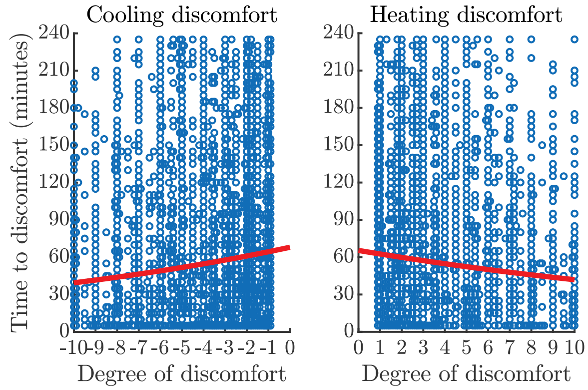


**Figure 2** An exponential fit DoD vs. TTD

## RESULTS AND DISCUSSION

### Data Curation and Feature Extraction

Filtering ecobee's DyD single occupancy dataset for instances where the space was occupied for 2-hrs, or until an MSC, after the SC resulted in ~50% data reduction. Of this 50%, ~6% of MSC data was not included due to HVAC mode transition from heating to cooling or cooling to heating season (shoulder season) as only heating and cooling seasons are within the scope of this paper. This preprocessing resulted in approximately ~367k SCs, including ~111k MSCs.

### Population Model

A scatter plot of all the MSCs' DOD and TTD (within 240 min.) is shown in Figure 1. An exponential curve defined by Equation 5 was fit to these data. The resulting models predict TTD (in min.) given DOD (in $°F$) with paramters $\alpha_1 = 68$ and $65$ and $\alpha_2 = 0.055$ and $-0.045$ for cooling and heating season respectively. The high spread of the data and the relatively close magnitude of these coefficients indicate thermal comfort behavior dynamics may not be statistically significantly different between heating and cooling seasons.

$$TTD = \alpha_1 \, e^{\alpha_2 \, DOD} \quad (5)$$

The quality of this model is expressed in the Figure 2 histogram of true vs. predicted TTDs . It can be seen that the population model captures well the 45% of overrides that occur within 60 minutes, but the accuracy of predicting TTD is 0 for longer time periods. It should also be noted that this model was only trained on the MSC data. Thus, it cannot predict if an MSC never occurs within the 4-hour window after the SC.

### Decision Tree

The first layer of the decision tree's accuracy in predicting if an MSC will occur is illustrated in the confusion chart of Table 3. Adding personal factors to the contextual only improves accuracy of correctly predicting an MSC from 37.6% to 38.6%. The first branch in the tree, indicitve of the primary contributing factor is based on the "event" feature, followed by two nodes for cooling setpoint and indoor temperature. The second tree model also poorly predicts the TTD, rarely predicting the MSC will be greater than 10-min. regardless of the true value as seen in Figure 3. This large error can likely be attributed to using a classifier to predict the TTD which lies on a continuous domain. For context, approximately 30% of the SCs do not result in an MSC, and ~20% of the SCs result in an MSC within 10-min. As such, the second layer focuses on mizimizing the error of the 20% of the MSCs that occur within 10-min. Since the classifier has no sense of "closeness" between incorrect TTD estimates it rarely chooses bins within this continuous domain.

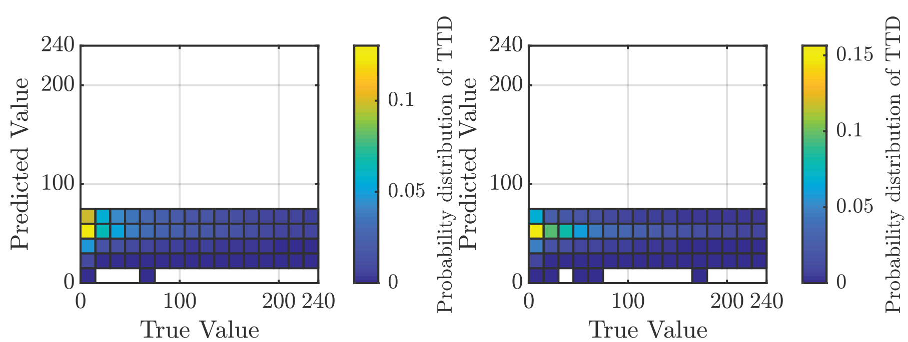


**Figure 1** Population model histogram of true vs. predicted TTD values in minutes (Left: cooling season; and Right: heating season)

## Neural Networks

Table 3 also illustrates the ANN's classification accuracy in predicting if an MSC will or will not occur. The accuracy improves from 67.3% to 70.6% with the addition of the personal factors, reinforcing previous findings that personal factors partially drive behavior. The Figure 4 histogram comparing the true TTDs to those predicted by the second ANN show nearly a Guassian distribution that consistiently over-estimates by ~15min. The reason for this bias is unclear, yet similar to the bias in the population model. While adding personal factors improved the ANN classifier's accuracy, it did not do so for the regression ANN. This suggests that (the chosen) personal factors may be important for determining if an override occurs, but not when.

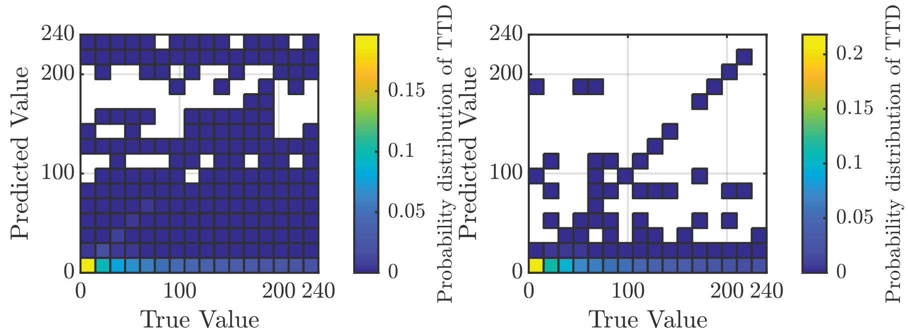


**Figure 3** Fine decision tree model histogram of true vs. predicted TTD values in minutes (Left: cooling season; and Right: heating season)

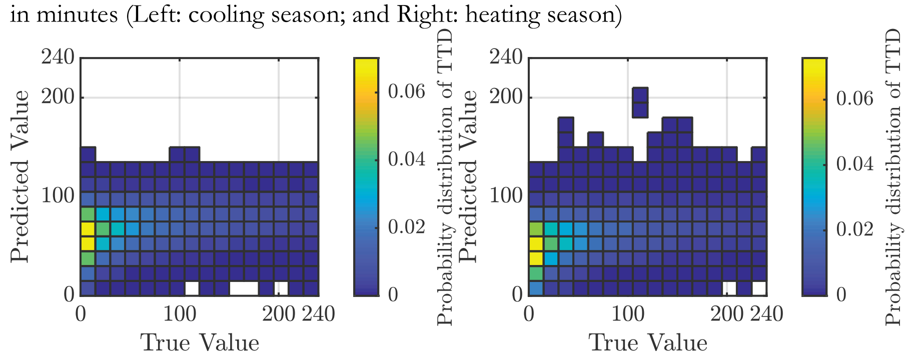


**Figure 4** ANN model histogram of true vs. predicted TTD values in minutes (Left: cooling season; and Right: heating season)

## Method Comparisons

Table 3 and Table 4 compare the accuracy of the different models in predicting if an MSC will occur and if so the resulting TTD. The ANN consistiently outperforms the tree in predicting both if an MSC will occur (70% vs. 39%

**Table 3: Confusion Chart for Classifying MSC vs. No MSC**

| Predicted Value | Model | True Value: No MSC | True Value: MSC | Row Summary |
|---|---|---|---|---|
| No MSC | NN | 50.4% 103,538 | 15.5% 31,753 | 76.5% 23.5% 135,291 |
| No MSC | Personal NN | 45.9% 94,368 | 13.9% 28,598 | 76.7% 23.3% 122,966 |
| No MSC | Tree | 46.7% 96,036 | 29.5% 60,658 | 61.3% 38.7% 156,694 |
| No MSC | Personal Tree | 46.3% 95,198 | 29.0% 59,678 | 61.5% 38.5% 154,876 |
| MSC | NN | 2.3% 4,774 | 31.8% 65,399 | 93.2% 6.8% 70,173 |
| MSC | Personal NN | 6.8% 13,944 | 33.4% 68,554 | 83.1% 16.9% 82,498 |
| MSC | Tree | 6.0% 12,276 | 17.8% 36,494 | 74.8% 25.2% 48,770 |
| MSC | Personal Tree | 6.4% 13,114 | 18.2% 37,474 | 74.1% 25.9% 50,588 |
| Column Summary | NN | 95.6% 4.4% 108,312 | 67.3% 32.7% 97,152 | |
| Column Summary | Personal NN | 99.8% 0.2% 108,312 | 70.6% 29.4% 97,152 | |
| Column Summary | Tree | 88.7% 11.3% 108,312 | 37.6% 62.4% 97,152 | |
| Column Summary | Personal Tree | 88.0% 12.0% 108,312 | 38.6% 61.4% 97,152 | |

**Table 4: RMSE of TTD Estimates when MSCs Occur**

| Method | Population Model | | Decision Tree | | Artificial Neural Network | |
|---|---|---|---|---|---|---|
| | Cooling | Heating | Cooling | Heating | Cooling | Heating |
| Without personal variables | 150 | 152 | 201 | 208 | 108 | 113 |
| With personal variables | | | 198 | 209 | 112 | 112 |

accuracy) and if it will not occur (~100% vs. 89%), the later being an important factor for DR programs to predict overrides. In terms of RMSE of TTD predictions, the tree prediction underforms even the population level model. However, the ANN does outperform the population model: ~110 min. vs ~150 min. of RMSE. Overall, there is no appreciable difference in behavior between heating and cooling seasons.

## SUMMARY AND CONCLUSIONS

This paper discusses the dynamics of occupant thermostat-use behavior, and the ability of ML algorithms to predict if an occupant will override a thermostat setpoint change, and if so how long will it take for the occupant to feel so uncomfortable as to get up and change the thermostat, i.e. TTD. A population level model was developed to predict TTD based solely on the DOD, i.e. the amount the thermostat was initially changed. Since contextual and personal factors are known to affect thermal comfort behavior, such features were extracted from the DyD dataset of ~367k setpoint changes which includes ~111k manual setpoint changes (MSC). Fine decision tree and ANN models were both explored to map these features to estimate if an MSC occurs, and if so, to estimate the TTD.

Overrall the ANN model performed best, able to predict if an MSC occurs with 240 min. of an SC with 71% accuracy and estimate the TTD with an RMSE of ~110 min. While these errors seem large, they are a ~30% improvement in the population level model. The chosen contextual information drive both if an MSC will occur, and if so, the TTD; the chosen personal factors appear to have a marginal affect on if an MSC will occur, and no affect on TTD. The tree models poorly estimate both if an MSC will occur, and if so, the TTD. Given the current trends in DR programs, using contextual and personalized ANN models can help them reduce overrides compared to the current practice of issuing the same setpoint changes (driving DOD) across the entire population

## ACKNOWLEDGEMENTS

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 DyD dataset.

## CONFLICT OF INTEREST STATEMENT

This work is support in part through in-kind support from ecobee Inc. by furnishing the dataset to the authors under a non-exclusive no-cost agreement. No other competing financial interests or personal relationships could appear to have influenced the work reported in this paper.

## REFERENCES

Breiman, L., 2017. Classification and regression trees. Routledge.

Cappers, P., MacDonald, J., Page, J., Potter, J., Stewart, E., 2016. Future Opportunities and Challenges with Using Demand Response as a Resource in Distribution System Operation and Planning Activities (No. LBNL--1003951, 1333622). https://doi.org/10.2172/1333622

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. https://doi.org/10.1016/j.enbuild.2015.01.038

de Dear, R., Schiller Brager, G., 2001. The adaptive model of thermal comfort and energy conservation in the built environment. Int. J. Biometeorol. 45, 100–108. https://doi.org/10.1007/s004840100093

Deng, Z., Chen, Q., 2018. Artificial neural network models using thermal sensations and occupants' behavior for predicting thermal comfort. Energy Build. 174, 587–602. https://doi.org/10.1016/j.enbuild.2018.06.060

D'Oca, S., Chen, C.-F., Hong, T., Belafi, Z., 2017. Synthesizing building physics with social psychology: An interdisciplinary framework for context and occupant behavior in office buildings. Energy Res. Soc. Sci. 34, 240–251. https://doi.org/10.1016/j.erss.2017.08.002

Ecobee, Inc., 2018. Donate your Data [WWW Document]. Smart WiFi Thermostats Ecobee. URL https://www.ecobee.com/donateyourdata/ (accessed 2.10.18).

Fit data by training a two-layer feed-forward network - MATLAB [WWW Document], n.d. URL https://www.mathworks.com/help/deeplearning/ref/neuralnetfitting-app.html (accessed 12.2.19).

Kane, M.B., 2018. Modeling Human-in-the-Loop Behavior and Interactions with HVAC Systems, in: 2018 Annual American Control Conference (ACC). Presented at the 2018 Annual American Control Conference (ACC), pp. 4628–4633. https://doi.org/10.23919/ACC.2018.8431913

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. https://doi.org/10.1016/j.buildenv.2018.01.023

Marie-Andrée, L., Ahmed, D.A., Célyn, L.B., 2011. DEVELOPING WINTER RESIDENTIAL DEMAND RESPONSE STRATEGIES FOR ELECTRIC SPACE HEATING, in: Proceedings of Building Simulation 2011. Presented at the 12th Conference of International Building Performance Simulation Association, International Building Performance Simulation Association, Sydney, AU, p. 8.

Mathworks, 2019. Decision Trees [WWW Document]. MathWorks Doc. URL https://www.mathworks.com/help/stats/decision-trees.html?searchHighlight=decision%20trees&s_tid=doc_srchtitle (accessed 12.1.19).

Mathworks, n.d. Fit data by training a two-layer feed-forward network [WWW Document]. MathWorks Doc. URL https://www.mathworks.com/help/deeplearning/ref/neuralnetfitting-app.html (accessed 12.1.19).

Mirakhorli, A., Dong, B., 2016. Occupancy behavior based model predictive control for building indoor climate—A critical review. Energy Build. 129, 499–513. https://doi.org/10.1016/j.enbuild.2016.07.036

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. https://doi.org/10.1016/j.buildenv.2017.01.023

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

Schweiker, M., Huebner, G.M., Kingma, B.R.M., Kramer, R., Pallubinsky, H., 2018. Drivers of diversity in human thermal perception – A review for holistic comfort models. Temperature 5, 308–342. https://doi.org/10.1080/23328940.2018.1534490

Seiden, K., Olig, C., Max, D., Sherman, M., 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 IEEE Third International Conference on Smart Grid Communications (SmartGridComm). Presented at the 2012 IEEE Third International Conference on Smart Grid Communications (SmartGridComm), IEEE, Tainan, Taiwan, pp. 103–108. https://doi.org/10.1109/SmartGridComm.2012.6485967

Zhang, H., Arens, E., Huizenga, C., Han, T., 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. https://doi.org/10.1016/j.buildenv.2009.06.018