A Domain-Assisted Data Driven Model for Thermal Comfort Prediction in Buildings

Liang Yang
Sch. Computer Science, Beijing Institute of Technology

Zimu Zheng
Dept. Computing, The Hong Kong Polytechnic University

Jingting Sun
HVAC & Gas Institute, Tongji University; Dept. BSE, HKPolyU

Dan Wang
Dept. Computing, The Hong Kong Polytechnic University

Xin Li
Sch. Computer Science, Beijing Institute of Technology

ABSTRACT

Recent studies on thermal comfort often require feedback from occupants or additional devices installed. This often limits the scalability of these approaches. In this paper, we for the first time study thermal comfort prediction of an occupant by training a model from the data of not only the targeted occupant but also others, guided by domain knowledge. We demonstrate, using ASHRAE data, that this approach has potential, and is worth exploring.

CCS CONCEPTS

• **Information systems → Data analytics**; • **Hardware → Thermal issues**;

KEYWORDS

Thermal Comfort, Domain Knowledge, Applied Machine Learning

ACM Reference format:
Liang Yang, Zimu Zheng, Jingting Sun, Dan Wang, and Xin Li. 2018. A Domain-Assisted Data Driven Model for Thermal Comfort Prediction in Buildings. In *Proceedings of e-Energy '18: International Conference on Future Energy Systems, Karlsruhe, Germany, June 12–15, 2018 (e-Energy '18)*, 6 pages.
https://doi.org/10.1145/3208903.3208914

1 INTRODUCTION

Thermal comfort is defined as the condition of mind that expresses satisfaction with the thermal environment [1], or simply put, it shows whether an occupant feels cold or hot. Thermal comfort provides a quantitative assessment linking the setting of indoor thermal environment parameters to occupant's subjective evaluation. It has been widely applied for decision-making of building design alternatives, operation set-point management of the air-conditioning system, and it heavily influences the level of energy consumptions in buildings. As a matter of fact, building operators are reluctant to adopt new energy conservation technologies without having a clear knowledge of their impact on occupant comfort.

Thermal comfort has been an important topic in the domain of built environment. Studies emphasize on modeling an "average" person, in the sense that volunteers are first recruited and evaluated; then thermal comfort is drawn by some statistic average of the comfort of these volunteers. Recently, individualized thermal comfort environment is highly advocated. With the advances of smart devices and sensing technologies, computer researchers contributed many approaches on getting the knowledge of personalized thermal comfort. In these studies, certain participatory feedbacks or additional devices are needed for each occupant. This limits the scalability of these approaches.

In this paper, we propose a data-driven approach to conduct thermal comfort prediction for personalized thermal comfort and with scalability. More specifically, we train a model from an existing set of thermal sensation data of occupants, and then we use this model to predict the thermal comfort level of any individual occupant. The underlining assumption is that "similar people" may have similar thermal sensation. As such, for an occupant, the methodology changes from learning the comfort level from this occupant, to finding occupants that are similar to this occupant.

We present an initial study on the feasibility of this approach. We note that many data, though collected with separated effort and goals, have already been accumulated. For example, the data we use in this study are from ASHRAE, the standard body from built environment. The ASHRAE RP-884 data, with 22000 entries, were collected and have been publicly available since the late 1990s. These data were used for physiology/adaptive model development. There are other data set as well, e.g., De Dear et. al., is collecting a more comprehensive set of data to develop refined adaptive models [2]. Yet there is no previous study that uses data of an occupant to train a model to predict the thermal comfort of other occupants.

In this paper, we first show that a data-driven approach has a clear performance gain over the average-based PMV model. We apply the state-of-the-art ensemble model for training and learning. We found, however, that the features selected by the ensemble model conflict with domain knowledge. The reason is the noise introduced from purely data-driven model. While we see advances [3, 4] in collecting big data in the physical environment, the amount is still difficult to reach a scale of online social networks where one can purely rely on data-driven approaches to remove noise.

To this end, we propose a domain-assisted data driven model. The feature selection phase is jointly carried out with data and domain

e-Energy '18, June 12–15, 2018, Karlsruhe, Germany

ACM ISBN 978-1-4503-5767-8/18/06…$15.00
https://doi.org/10.1145/3208903.3208914

knowledge, where thermal comfort factors have been densely studied. Intrinsically, domain expertise provides affirmative knowledge, with which fewer data are needed to differentiate features.

In summary, the contributions of this paper are: 1) we for the first time study thermal comfort prediction, using a data-driven approach. This approach can estimate individualized thermal comfort prediction and has good scalability; 2) we observe that a direct application of the state-of-the-art data-driven machine learning models may not lead to explainable feature selection. We propose a domain-assisted feature selection process.

2 BACKGROUND

2.1 Related Work of Thermal Comfort

The **thermal comfort research in built environment** commonly study an "average" person as representative in the prediction of thermal comfort. Domain predictive methodologies could be generally divided into two categories: static models and adaptive models. Static models (as Fanger's PMV model [5]) view a person as a passive recipient of thermal stimulus and mediate the effect of ambient thermal environment exclusively through the heat and mass exchanges between the human body and the ambient environment. In contrast, adaptive models [1, 6] consider a human subject as an interactive participant in creating their own desired thermal environment.

Domain research has accumulated affirmative knowledge. For example, different scenarios can be partitioned to assist thermal comfort prediction. It is shown that comfort conditions over time could be subdivided with the season as the occupant's clothing and corresponding comfort zone are different between cooling season and heating season [5]. Building types (mechanically conditioned versus naturally ventilated) have also been classified to determine the selection of comfort models (PMV and adaptive, respectively) because the occupants within natural ventilation buildings are often given more "adaptive opportunities" [1, 6].

Domain studies have provided dense understanding on the influence of various parameters including indoor and outdoor climate factors on thermal comfort prediction during the past half century. For example, the temperature is shown to be the dominant determinant both interpreted with heat balance equation and convinced with field experiment analysis [5, 6].

All these methodologies, however, were used to predict the mean response of a group of people exposed to the same environment, i.e., to create an "average" person. Individual thermal preference could not be expressed with acceptable accuracy [7].

The **thermal comfort studies in computer research** applied various ubiquitous techniques to achieve personalized thermal comfort [8]. Thermovote let occupants vote in offices and then adjust the temperature [9]. A joint model of a physiology model and a voting model Temperature Comfort Correlation (TCC) was developed in [10]. With a physiology model as a baseline, the number of votes can be reduced. Devices were developed to infer human comfort levels [11] as well as adjust microenvironment (SPOT) [12–14]. All these studies either requires additional devices to be used or recurrent feedback from occupants. This limits the scalability of these approaches to arbitrary occupants. This challenge can be overcome by leveraging the historical data instead of run-time data to establish models. Data-driven thermal comfort models [15, 16] are then

Table 1: ASHRAE RP-884 Database

	Content	Variable Number	Detail of variables
Raw feature	Basic Identifiers	10	Information of building, subject (e.g., age, sex) and survey time
Raw feature	Thermal Questionnaire	8	Metabolic activity and insulation of subject
Raw feature	Personal Environmental Control	17	Perceived controls over the local thermal environment
Raw feature	Outdoor Climate Data	9	Outdoor air temperature (ta), relative humidity (rh) and effective temperature at 3pm and 6am on day of survey
Raw feature	Indoor Climate Physical Observation	14	Indoor thermal parameters including air temperature (taav), globe temperature, air speed, turbulence intensity, dewpoint temperature and plane radiant asymmetry temperature (trav)
Raw feature	Calculated Indexes	18	Average or maximum of indoor thermal parameters, effective temperature (et), standard et (set), PMV,PPD,two-node model index (tsens), percent Dissatisfied due to draft
Labels could be used for prediction purpose		7	Thermal sensation, thermal acceptability, thermal preference, air movement acceptability, air movement preference, general thermal comfort

developed for individuals and improve HVAC energy, e.g., with the setpoint control. However, they usually assume that different individuals prefer different thermal models and train each model with data from merely the targeted individual, which requires a large amount of data collected from each individual for consistent performance and incurs performance reduction when data are insufficient. Guided by domain knowledge, our approach overcome the challenge by using data not only from the targeted individual but also from others.

This paper for the first time leverage domain knowledge to examine whether we can use others' data to predict thermal comfort, which integrates scalability, interpretability and accuracy.

2.2 ASHRAE RP-884 Database

We use the data sets from ASHRAE RP-884 database. ASHRAE RP-884 database was originally collected to develop De Dear's adaptive comfort model. 25,248 observation sets from 160 buildings involving more than 80 variables were collected.

The data are shown in Table 1. We categorize the data with the variables that can be used as raw features and the variables that can be used for prediction. For raw features, we have basic identifiers, e.g., age and sex. These are personalized data, yet it is easier to collect these data. There is information on the physical environment, which can be collected by building automation systems. We have data which we call labels. These are related to occupant comfort and difficult to obtain. In this paper, we conduct prediction on direct comfort metrics of thermal sensation and preference which are for general thermal comfort process, while the unselected ones are for adaptive thermal comfort process and we leave the related investigation as future work.

3 THE DESIGN DETAILS

3.1 Problem Statement

To better assist decision makings on HVAC operations, it would be more helpful to directly obtain whether an occupant feels hot (need cooler), comfortable (need no change), or cold (need warmer).

From this point onwards, we directly infer such preference instead of numerical thermal comfort levels. We say a *task* to be a set of scenarios (which will be determined and justified in Section 3.2.1) for occupants, including the a set of features (e.g., temperature and humidity in RP-884 database) and the corresponding thermal comfort level to predict. For example, prediction for occupants in summer and nature ventilation buildings. Let $x_i^m \in \mathbb{R}$ be the value of the mth feature of the ith task, and $m \in [0, M]$, $i \in [0, I]$. Let $\boldsymbol{x}_i$ be the vector of x_i^m. Let y_i be the labeled comfort level (e.g., cold, comfortable, or hot) of the occupant of the ith task.

Problem MAPE-TC (***M****ulti-t****a****sk* ***P****r****e****diction on* ***T****hermal* ***C****omfort*): Given an observable full feature space $X = [x_i^m]$ and a label space $Y = [y_i]$. The goal is to find a classifier $F(\boldsymbol{x}_i)$ that estimates the corresponding label y_i, which minimizes the total number of prediction errors $E(X,Y) = \sum_i^I L(F(\boldsymbol{x}_i), y_i)$, where $L(\cdot)$ denotes a discriminant function which outputs whether a prediction error occurs or not:

$$L(F(\boldsymbol{x}_i), y_i) = \begin{cases} 0 & F(\boldsymbol{x}_i) = y_i \\ 1 & F(\boldsymbol{x}_i) \neq y_i \end{cases}$$

3.2 A Domain-assisted Data Driven Model

To solve the MAPE-TC Problem, a natural thinking is to leverage state-of-the-art data-driven model to minimize prediction errors. However, it may not be a good idea to follow all conclusions made by the purely data-driven model. Existing works [17] show that purely data-driven techniques introduce noise, which requires proper feature complexity and data size. Such a noise problem exists especially when using RP-884 for thermal comfort prediction, due to two reasons. First, because thermal comfort model is complex and related to many factors, the data-driven model leverages multiple features of different types of contents trying to capture the dynamic nonlinear relationship with the predicted thermal preference. As we known, in several cases, there are redundant features which bring noise to the training model and distract the feature selection. Second, though there are feature selection mechanisms, the size of the RP-884 dataset is still not big enough (tens of millions of records on Twitter in 2010 [18]) for the purely data-driven method to reliably remove the noise (see detailed discussion in Appendix A and B).

Both issues together make it difficult to solve the noise problem here by a purely data-driven method. To this end, we now propose a domain-assisted data driven model, where we leverage affirmative domain knowledge to assist feature selection, so as to remove noise when using the RP-884 database.

Our domain-assisted data driven model (Fig. 1) contains three phases. First, after data preprocessing as in [19], we have a raw feature set X in Table 1. The **Domain-based Task Definition** module partitions X according to scenarios. Second, the **Iterative Co-selection** module generates the feature with both domain knowledge and statistics. There are two sub-modules. The Domain-based Feature Engineering module iteratively generates new features X_{new} and then test them in Data-driven Testing until the performance cannot be improved. Finally, the **Prediction Model** module takes the final feature set X_{final} for training and prediction.

For the readers (e.g., application developer) who want to build the referred domain-assisted data driven model, what they need to do is to leverage the concluded features listed in Table 3 with a selected learning model, e.g., XGBoost. The ASHRAE dataset has

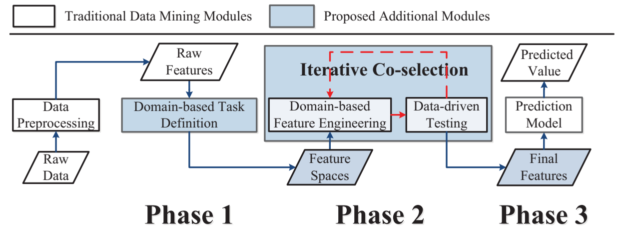

Figure 1: Domain-assisted data driven model.

nearly 80 variables and our reduced set could be used so that the data collection task becomes even easier.

3.2.1 Domain-based Task Definition. Domain knowledge has shown that thermal comfort has a wide range of application scenarios, e.g., in different types of buildings. In different scenarios, not only features but also applicable domain models differ greatly.

(1) As for our above TAAV temperature example, the importance of TAAV can be changed in different scenarios. In HVAC building, it can be more important than natural ventilation building because people are less adaptive [1, 6].
(2) For mechanically conditioned spaces, such as naturally ventilation building, the adaptive model shows better feasibility [5, 6]; while for centralized HVAC building, PMV model is commonly used.

Thus, we first partition these scenarios as *tasks*. Each task is associated with a thermal comfort model. Within the same task, occupants share the training data.

With domain knowledge, we study two types of information to determine tasks for the RP-884 data. The first is building types. As mentioned above, PMV and adaptive model can be used in different building types. The second is seasons. These different indoor thermal models are also desired in different seasons [1] (introduced in Section 2.1). As such, the given database is divided into different tasks, i.e., summer_hvac, winter_hvac, summer_nv, winter_nv. For different tasks, different domain models are used for feature engineering.

3.2.2 Iterative Co-selection. This module selects individual features to form the final feature set.

Domain-based Feature Engineering. As mentioned above, there are two possible reasons for the prediction error in purely data-driven model. The first is that a feature is not related and brings in noise itself. The second is that though a feature is important, the amount of data is too small in the database and thus introduces noise. We handle these through 1) Task-based Feature Selection and 2) Size-based Feature Filtering.

(1) *Task-based Feature Extraction.* Each task has its own model and thus model variable which can be taken as feature $\boldsymbol{x}^d$. For example, PMV model includes four environmental variables (air temperature *taav*, radiant temperature *trav*, air speed *velav* and humidity level *rh*) and two personal variables (clothing *insul* and physical activity *met*); while De Dear's adaptive comfort model, additionally includes the mean outdoor effective temperature *dayav_et*.
(2) *Size-based Feature Selection.* Besides, there are other contextual features with prediction potentials. Such features $\boldsymbol{x}^s$ are ordered by the size of their samples. The more samples they

Table 2: Datasets Summary (incomplete records removed).

Seasons		Summer		Winter	
Building Types		HVAC	NV	HVAC	NV
MAPE-TC Problem	#Training Records	1,127	3,387	320	1,467
	#Testing Records	376	1,129	107	489
# Total Records		1,503	4,516	427	1,956
#Feature Combinations		$\sum_{k=1}^{n=76}\binom{n}{k} = 2^{76}-1$			

have, the higher priority they get. Merely the top τ features are used in feature engineering. Based on our experience on the RP-884 data set, a suggested setting of τ is five.

With the features $\boldsymbol{x} := \boldsymbol{x}^d \cup \boldsymbol{x}^s$ selected and filtered above, this feature engineering module increasingly generates a new set X_{new}. It avoids introducing noises of training all existing features at the same time. Starting from an empty set, each time X_{new} tries to add a feature in $\boldsymbol{x}$ that brings the greatest accuracy increment.

Data-driven Testing. This module tests the performance of the final features. As justified in the Appendix B, we apply the state-of-the-art XGBoost as the testing data-driven model.

If the accuracy of a new feature set X'_{new} does not improve from the previous X_{new} for several v iterations, then the X'_{new} is recognized as the final subset X_{final}. Otherwise we continue the feature engineering, i.e., adding another features into $\boldsymbol{x}$ in each iteration considering both the targeted task and the size of samples associated with the new feature.

3.2.3 Prediction Model and Comparison. We choose to leverage general domain knowledge instead of merely maximize the accuracy on our small testing data and evaluate our final domain-assisted data driven model. The final features selected for two metrics are shown in Table 3. Note that though the data is not big enough to self-solve the noise problem, the sheer amount of data and features is significant. The numbers of data samples and possible feature combinations are summarized in Table 2. We establish a private cloud to run our experiments. The cloud has 12 cores, each with 2.6GHz, and a total memory of 64G.

We applied two standard metrics in built environment: thermal preference (*MCI*) and thermal sensation (*ASH*) [1, 6]. MCI has three categories of "want cooler" (1), "want no change" (2), and "want warmer" (3), and ASH has seven categories from "cold" (-3) to "hot" (3). In the four tasks, we count the number of errors in three types, i.e., false hot, false comfort and false cold. Fig. 2 shows the percentage of the number of errors reduced using our approach as against to the PMV model under MCI (Fig. 2 (a)) and ASH (Fig. 2 (b)). While confirming features to domain knowledge as in domain models shown in Table 3, our approach still maintains its performance and always outperforms the PMV model. For example, in buildings with HVAC on winter, our approach shows improvement of 33.36% and 51.49% on MCI and ASH, respectively.

4 CONCLUSION AND FUTURE WORK

In this paper, we proposed a data-driven model for thermal comfort prediction. To the best of our knowledge, this is the first model that predicts the thermal comfort of an occupant by using the data of other occupants with domain knowledge. As such, this model can estimate personalized thermal comfort that scales to any occupants.

We observed that a direct application of the state-of-the-art data learning model does not lead to explainable results due to noises.

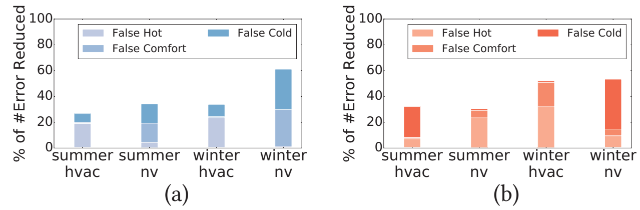

Figure 2: Percentage of reduced error using XGBoost instead of PMV, predicting (a) MCI and (b) ASH in four data sets.

Table 3: Final Features in four tasks for ASH and MCI

Feature	Summer_hvac	Summer_nv	Winter_hvac	Winter_nv
age	×	ASH, MCI	×	×
sex	×	ASH	MCI	ASH, MCI
met	×	ASH	×	ASH, MCI
insul	MCI	MCI	×	MCI
taav	ASH, MCI	×	MCI	ASH
trav	ASH	ASH	MCI	ASH, MCI
velav	×	MCI	ASH, MCI	ASH
rh	MCI	ASH	ASH	ASH, MCI
et	ASH	ASH, MCI	ASH	ASH, MCI
set	×	×	MCI	×
tsens	ASH	ASH	ASH, MCI	ASH
pcc_ag	MCI	×	×	×
day15_ta	ASH, MCI	ASH, MCI	ASH	ASH
day06_ta	×	ASH	MCI	ASH, MCI
dayav_rh	×	ASH	×	×
dayav_et	MCI	MCI	ASH, MCI	×

We showed that domain knowledge is important. Our model used such knowledge in two ways. First, on the macro level, domain knowledge can be used to partition scenarios. Second, on the micro level, domain knowledge can be used to select individual features because built environment has dense studies on the influence of some features in certain occasions.

We provided reduced sets of features for the four targeted tasks. The ASHRAE dataset has nearly 80 variables and our reduced set could be used so that the data collection becomes even easier.

Our model is preliminary but it has shown good potentials. We see that it leads to two future research directions. First, a knowledge base can be developed with domain experts for both the macro and micro level modules of our model. Rules and mechanisms on automatic domain knowledge selection should also be developed. Second, our problem is currently under the formulation with a 0/1 loss, which, admittedly, may not be the best formulation due to the lack of data. A person could be slightly uncomfortable, or extremely hot/cold. So a more graded prediction would be more useful in building management. Third, clearly, the more dimensions of data, the merrier. Yet not only personalized data (as in this paper) can be collected, but also the data related to the correlation between people in thermal comfort can be collected and analyzed, e.g., the information of zone type can benefit a more graded task definition. Fourth, it would also be useful if this test can be performed on real-world buildings to further validate the efficacy of this approach.

ACKNOWLEDGMENTS

Parts of the work were supported by the Hong Kong Polytechnic University under Grant No.: 1-BBYX, the National Key Research and Development Program of China under Grant No.: 2017YFC0702600, and the National Natural Science Foundation of China under Grant No.: 61772074

APPENDIX

A AN AVERAGE PERSON MODEL

To tackle **MAPE-TC**, we first study the average person model from built environment. We use Fanger's PMV model. Given the same inputs, i.e., occupant information and environment settings, PMV votes can be calculated from the RP-884 database. We compare it with the real comfort votes surveyed, also in the RP-884 database.

To compare the PMV results and real comfort votes, we show the prediction loss, i.e., $|y_{\text{real}} - y_{\text{PMV}}|$. Figure 3 shows the results. Fig. 3, a prediction result is regarded as acceptable (marked in green) if it is in the same class as the ground truth, otherwise unacceptable (marked in red). We see that there is substantial prediction loss of PMV. As a matter of fact, the number of unacceptable cases is 40.14%.

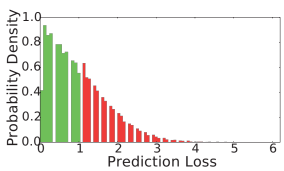

Figure 3: Probability density function of acceptable (in green) and unacceptable (in red) PMV loss.

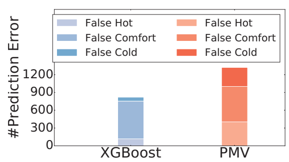

Figure 4: Number of the total prediction error using XGBoost and PMV model in RP-884 database.

B A DATA-DRIVEN MODEL

B.1 EXtreme Gradient Boosting (XGBoost)

We apply the state-of-the-art E**X**treme **G**radient **Boost**ing (XGBoost) Algorithm, which has proved its superior recently on major non-linear classification problems in real-world contests of Kaggle [20]. We compare XGBoost with the PMV model. In the comparison, a prediction error is counted when the model makes a wrong classification, e.g., predicting the user as cold when s/he is not. We show three types of errors including false hot (in deep color), false comfort (in normal color) and false cold (in light color).

Fig. 4 shows the total number of prediction errors. We see that the data-driven approach improves the PMV for 37.43%. That is because the data-driven approach better adapts to the real-world data of the local context, which reduces the prediction error coming from the general average-person model.

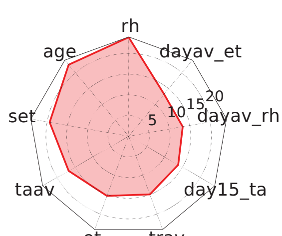

Figure 5: Features with top importance ($\times 10^3$).

B.2 Limitation of the Data-driven Approach

On the one hand, the data-driven model is effective in reducing errors. On the other hand, we need to justify whether such effectiveness is valid.

To do this, we choose to study whether the results can be interpreted. One way to evaluate the interpretability is to study the trained relationship between features and labels in the model, i.e., we can obtain the importance of features trained by XGBoost. We then request a domain expert to evaluate whether the importance ranking of the features can be interpreted by domain knowledge. The experiment and conclusion are as below.

In XGBoost, it has an embedded function of *f-score* to show the importance of all features and used to select feature. A feature containing more information for classification will be marked with a higher f-score. Figure 5 shows the top ranked features.

It may not be a good idea to follow all conclusions made by the purely data-driven model. We found that the feature importance rank of such model shows poor interpretability, i.e., it shows obviously unconvincing result against common sense and domain knowledge. As in Fig. 5, the f-score of age is marked as high as 22556 and ranks the second; while that of indoor air temperature average (TAAV) is marked 16744 and ranks the fourth. In other words, the data-driven model suggests that TAAV is less important even than age. This finding conflicts with domain knowledge. As we known, TAAV directly affects the thermal feedback, i.e., hot or cold. In the built environment, TAAV has been confirmed to be the most determinant of thermal comfort prediction both from comfort models [6, 21] and ISO standards [5]; while age is usually regarded to be more an indirect factor of thermal comfort than TAAV and usually not considered especially in HVAC building [1, 6].

An average model led by domain knowledge will not contain such a noise against common sense. The reason of such a noise is introduced by the purely data-driven techniques, which requires proper feature complexity and data size [17]. First, because thermal comfort model is complex and related to many factors, XGBoost leverages multiple features of different types of contents trying to capture the dynamic non-linear relationship with the predicted thermal preference. As we known, in several cases, there are redundant features which bring noise to the training model and distract the feature selection. Second, though XGBoost contains embedded feature selection mechanism using f-score, the size of the RP-884 dataset is still not big enough for the purely data-driven method to reliably remove the noise, which leads to the noise shown in Fig. 5.

We believe both issues together make it difficult to solve the noise problem here by a purely data-driven method, e.g., f-score based selection, SVD and PCA. For the first problem, it is not easy for a layman to know whether some features are noise or not. For the second problem, while advances are being made [3, 4] in collecting related data to the physical world, the size of dataset (tens of thousands of records) does not reach a scale of online social networks (tens of millions of records on Twitter in 2010 [18]) for the purely data-driven method to reliably remove the noise.

REFERENCES

[1] ASHRAE (Atlanta, Georgia). *ANSI/ASHRAE Standard 55-2013: Thermal Environmental Conditions for Human Occupancy.* 2013.

[2] Christhina Candido, Jungsoo Kim, Richard de Dear, et al. Bossa: a multidimensional post-occupancy evaluation tool. *Building Research & Information*, 2016.

[3] Olivia Guerra-Santin and Christopher Aidan Tweed. In-use monitoring of buildings: An overview of data collection methods. *Energy and Buildings*, 93, 2015.

[4] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta N Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. Farmbeats: An iot platform for data-driven agriculture. In *NSDI'17*, pages 515–529.

[5] ISO 7730: 2005. Ergonomics of the thermal environment – Analytical determination and interpretation of thermal comfort using calculation of the pmv and ppd indices and local thermal comfort criteria.

[6] Richard de Dear, Gail Brager, and Donna Cooper. Developing an adaptive model of thermal comfort and preference. *FINAL REPORT ASHRAE RP-884*, 1997.

[7] Niklas Fransson, Daniel Västfjäll, and Jennie Skoog. In search of the comfortable indoor environment: A comparison of the utility of objective and subjective indicators of indoor comfort. *Building and Environment*, 42(5):1886–1890, 2007.

[8] Adrian K Clear, Janine Morley, Mike Hazas, Adrian Friday, and Oliver Bates. Understanding adaptive thermal comfort: new directions for ubicomp. In *Proc. ACM UbiComp'13*, pages 113–122.

[9] Varick L Erickson and Alberto E Cerpa. Thermovote: participatory sensing for efficient building hvac conditioning. In *Proc. ACM BuildSys'12*, pages 9–16.

[10] Abraham Hang-yat Lam, Yi Yuan, and Dan Wang. An occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings. In *Proc. ACM e-Energy'14*, pages 133–143.

[11] Farrokh Jazizadeh and S Pradeep. Can computers visually quantify human thermal comfort?: Short paper. In *Proc. ACM BuildSys'16*, pages 95–98.

[12] Peter Xiang Gao and Srinivasan Keshav. Spot: a smart personalized office thermal control system. In *Proc. ACM e-Energy'13*, pages 237–246.

[13] Peter Xiang Gao and Srinivasan Keshav. Optimal personal comfort management using spot+. In *ACM BuildSys'13*, pages 1–8.

[14] Alimohammad Rabbani and Srinivasan Keshav. The spot* personal thermal comfort system. In *Proc. ACM BuildSys'16*, pages 75–84.

[15] Zheng Yang and Burcin Becerik-Gerber. The coupled effects of personalized occupancy profile based hvac schedules and room reassignment on building energy use. *Energy and Buildings*, 78:113–122, 2014.

[16] Ali Ghahramani, Farrokh Jazizadeh, and Burcin Becerik-Gerber. A knowledge based approach for selecting energy-aware and comfort-driven hvac temperature set points. *Energy and Buildings*, 85:536–548, 2014.

[17] Juha Reunanen. Overfitting in making comparisons between variable selection methods. *Journal of Machine Learning Research*, 3(Mar):1371–1382, 2003.

[18] Shaozhi Ye and S Felix Wu. Estimating the size of online social networks. *International Journal of Social Computing and Cyber-Physical Systems*, 1(2):160–179, 2011.

[19] Richard de Dear and Mark E Fountain. Thermal comfort in air conditioned office buildings in the tropics. *Center for the Built Environment*, 1994.

[20] Chen et al. XGBoost: A scalable tree boosting system. In *Proc. ACM KDD'16*.

[21] Poul O Fanger et al. Thermal comfort. Analysis and applications in environmental engineering. 1970.