Predicting Individual Thermal Comfort using Machine Learning Algorithms

Asma Ahmad Farhan[1], Krishna Pattipati [2], Bing Wang [1], and Peter Luh[2]

***Abstract*— Human thermal sensation in an environment may be delayed, which may lead to life threatening conditions, such as hypothermia and hyperthermia. This is especially true for senior citizens, as aging alters the thermal perception in humans. We envision a decision support system that predicts human thermal comfort in real-time using various environmental conditions as well psychological and physiological features, and suggest corresponding actions, which can significantly improve overall thermal comfort and health of individuals, especially senior citizens. The key to realize this vision is an accurate thermal comfort model. We propose a novel machine learning based approach to learn an individual's thermal comfort model. This approach identifies the best set of features, and then learns a classifier that takes a feature vector as input and outputs a corresponding thermal sensation class (i.e. "feeling cold", "neutral" and "feeling warm"). Evaluation using a large-scale publicly available data demonstrates that when using Support Vector Machines (SVM) classifiers, the accuracy of our approach is 76.7%, over two times higher than that of the widely adopted Fanger's model (which only achieves accuracy of 35.4%). In addition, our study indicates that two factors, a person's age and outdoor temperature that are not included in Fanger's model, play an important role in thermal comfort, which is a finding interesting in its own right.**

I. INTRODUCTION

Human thermal comfort, i.e., their condition of mind that expresses their satisfaction with the thermal environment [1], has tremendous impact on their health and productivity. Perception of thermal comfort in an environment, i.e., feeling too warm, too cold, or comfortable, however, may exhibit delays. This is particularly true for senior citizens, who do not sense the heat or cold in overly warm or cold environments until after an extended time period. This delayed perception makes them vulnerable to hypothermia [10] or hyperthermia [21] conditions, which can be life threatening in severe cases.

We envision a decision support system that automatically predicts an individual's thermal comfort in an environment in real-time and suggests the corresponding recommended actions. For instance, this decision support system may suggest to a person that the current condition is too cold for him and suggests him to put on an additional layer of clothing. It may also suggest changes to the set point of the HVAC (heating, ventilating, and air conditioning) inside a building, e.g., increase the temperature by two degrees. An individual's thermal comfort can be predicted by the decision support system well ahead of the actual perception of the thermal comfort settles in. The suggested actions can improve the thermal comfort and prevent harmful conditions (e.g., hypothermia or hyperthermia) from happening, which is tremendously helpful to people, particularly for senior citizens with delayed perception of thermal sensations.

The key to the decision support system is a thermal comfort model that accurately predicts the thermal sensation for a person. Human thermal perception has been shown to be a complex process, depending on various environmental conditions as well psychological and physiological attributes of a person. Fanger's model [13] is the most widely accepted thermal comfort model. It has been adopted as part of ASHRAE (American Society of Heating, Refrigerating, and Air-Conditioning Engineers) 55 standard. On the other hand, despite being widely accepted, Fanger's model has been shown to have various drawbacks (see [11], [22] and the references within).

In this paper, we propose a novel machine learning based approach to learn an individual's thermal comfort model. Specifically, our approach identifies the best set of features (including various environmental, psychological and physiological attributes), and then learns a classifier that takes a feature vector as input and outputs a corresponding thermal sensation class (i.e. "feeling cold", "neutral" and "feeling warm"). Our approach differs radically from existing regression-based approaches [12], [13], [18], [19]. Evaluation using a publicly available large-scale data set demonstrates that when using Support Vector Machines (SVM) classifiers, the accuracy of our approach is 76.7%, over two times higher than that of Fanger's model (which only achieves accuracy of 35.4%). In addition, hypothesis tests demonstrate that our approach outperforms Fanger's model with strong statistical significance. Last, our study indicates that two factors, a person's age and outdoor temperature that are not included in Fanger's model, play an important role in thermal comfort, which is a finding interesting in its own right.

The rest of the paper is organized as follows. Section II briefly reviews related work. Section III presents our approach. Section IV describes the dataset. Section V evaluates the performance of our approach. Last, Section VI concludes the paper and suggested future work.

II. RELATED WORK

Broadly, two approaches are used to define thermal comfort, viz., heat balance approach and the adaptive approach.

*This work was supported by National Science Foundation grant 1331850.

[1]Asma Ahmad Farhan and Bing Wang are with the Department of Computer Science & Engineering, University of Connecticut, USA. (asma.farhan@engr.uconn.edu, bing@engr.uconn.edu)

[2]Krishna Pattipati and Peter Luh are with the Department of Electrical and Computer Engineering, University of Connecticut, USA. (krishna@engr.uconn.edu,luh@engr.uconn.edu)

Fanger's model [13] is a widely accepted model that is is based on Fanger's comfort equation [13], a heat balance equation that works on the principle of balance between metabolic heat production and the overall heat losses to the environment. The equation is derived using the thermal sensation data collected from subjects in a thermally controlled experimental chamber. The thermal sensation of the subjects is monitored for three hours, during which the subjects' responses are collected with the help of a questionnaire. The model maps physiological and indoor environmental parameters to a seven point ASHRAE thermal sensation scale named Predicted Mean Vote (PMV). PMV value is then used to propose the percentage of Dissatisfied (PPD) index, which finds the percentage of dissatisfied persons in certain settings. The PMV/PPD model is incorporated in ASHRAE guidelines for the thermal comfort of occupants in buildings, and ASHRAE suggests that at least 80% of the building occupants should be satisfied by the thermal conditions within the building. Despite success and widespread use of the Fanger's model, many researchers found significant discrepancies between the PMV and the Actual Mean Vote (AMV) [11], [22] .

Adaptive approach considers physiological, behavioral and psychological factors in a thermal comfort model. According to [4], adaptation is an important factor to consider when people interact with their environment. Adaptive approaches for thermal comfort are based on data collected from field studies. The goal of such studies is to understand the effects of "real" environment on thermal comfort. The adaptive model is primarily proposed for the naturally ventilated buildings. The model is part of ASHRAE (2004) Standard 55 as an optional method for naturally ventilated buildings. Various field studies are conducted across the globe to develop adaptive models. For example, [9] conducted a field study in Australia during the summer season and found a temperature difference of 1.7 Kelvin to -1.3 Kelvin between HVAC and naturally ventilated buildings. For further information, one may refer to [7].

Recently, the prevalence of smartphones have inspired many researchers to use participatory sensing to develop thermal comfort models. Several recent studies [12], [18], [19] have used participatory sensing to collect individual user response on thermal perception and proposed the concomitant thermal comfort models. These models range from individualized to thermal models for groups. Authors in [19] proposed "CarryEn" that combines the advantages of both PMV and adaptive thermal comfort to find user's comfort temperature. A weighing factor is used to incorporate the effect of outdoor temperature. Their idea is applicable to both single users and groups. In [12], authors proposed "Thermovote", a mobile application to improve occupant's thermal comfort. They used PMV and added humans as sensors to improve the thermal comfort. Humans' thermal sensation data are collected through smartphone app and are used in real-time to control the room temperature. The authors in [18] proposed SPOT+ that provides a measure of thermal comfort by predicting the future room occupancy

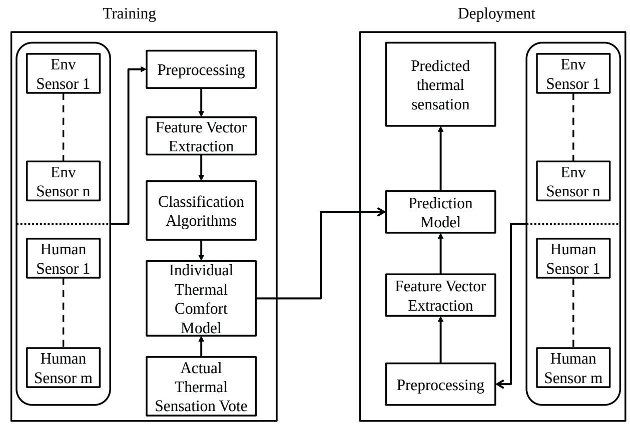

Fig. 1. High-level overview of our approach.

and optimal room temperature. Using the prediction model, SPOT+ finds a control schedule that optimizes both energy usage and thermal comfort.

Most existing thermal comfort models use a regression based approach, i.e., use a function to approximate the relationship between thermal sensation and various environmental, psychological and physiological factors. Our study differs from them in that we use a novel machine learn based approach. While machine learning based approach has also been used in prior studies [2] [14], they are motivated by the observation that PMV calculation in Fangers model involves an iterative process that is slow and incurs variable execution time. Therefore, their goal is to speed up PMV calculation so that it can be used for real-time control. In contrast, our study proposes a new thermal comfort model that significantly outperforms Fangers model.

III. Learning Thermal Model

Our goal is to derive a thermal comfort model that determines an individual's thermal sensation in real-time. Specifically, we consider three types of sensations, "uncomfortably warm (1)", "neutral (0)", and "uncomfortably cold (-1)". To learn an individual's thermal sensation, data are collected from various environmental and human (physical) sensors and are preprocessed. The feature vector is extracted from the raw data and a supervised classifier is trained using the features and the ground-truth thermal sensation as inputs. The trained classifier is then used for prediction and decision support. Figure 1 summarizes our approach, where the left part of the figure illustrates the training process, and the right part of the figure illustrates deploying the model for thermal comfort prediction in practice; n environmental sensors (e.g., sensing temperature, air turbulence) and m human sensors (e.g., related to metabolic rate, clothing) are being preprocessed to obtain a feature vector.

The following sub-sections first describes the feature selection and then the classification aspects of our approach.

A. Feature Selection

Human thermal sensation is influenced by both environmental and physiological factors. To understand their effects on thermal sensation, seven features are selected to train our thermal comfort model. Broadly, these features can be divided into three classes, viz., physiological features, and indoor and outdoor environmental features. Physiological features are used because thermal sensation is a subjective measure; different individuals perceive the same environment differently. Similarly, indoor environment (especially in HVAC settings) directly affects the occupant's comfort. Lastly, outdoor air temperature is used to include the psychological effects of seasons on thermal sensation and clothing preferences of individuals. To summarize, the following features are used to classify thermal sensation:

$$F = [A, C, M, T_a, T_m, A_t, T_o] \tag{1}$$

Here A represents age of the individual, C is the clothing insulation of the person, and M represents the metabolic rate. T_a, T_m and A_t represent the indoor environmental features corresponding to air temperature, mean radiant temperature and air turbulence, respectively. To consider the effect of outdoor environment, outdoor air temperature (T_o) is taken as the seventh features.

Our intuition to select age as one of the features is inspired by the fact that human thermoregulation and sensation change with aging. There is significant literature available that highlight these changes [10], [21]. Clothing is considered as one of the features in our model because it directly influences thermal sensation of a person. Additionally, the first step in dealing with thermal discomfort is change in clothing. Metabolic rate changes with the activity level and affects the thermal sensation; a physically active person has a different sensation than a sedentary one. Similarly, to understand the effect of indoor environment on human thermal sensation, indoor air temperature, mean radiant temperature and air turbulence are used as features for training the classifier. Since seasons have psychological effects on thermal perception, to understand this, outdoor air temperature is used as one of the features in our prediction model.

In Section V-A.1, we quantitatively evaluate the performance of using various feature sets using a publicly available data set. As we shall see, the feature set that provides the best performance conforms to our intuition. The only difference is that mean radiant temperature, T_m, does not help significantly to the data set we chose (it might be helpful in other data sets), and is omitted in our final selection of feature sets.

B. Machine Learning Algorithms

We experiment with a number of state-of-the-art machine learning classifiers. An overview of each of the classification algorithms is discussed below.

1) Support Vector Machine (SVM): SVM has a wide range of applications in machine learning and human behavior prediction [6]. It is a supervised learning algorithm that transforms training examples to a higher dimensional space and builds a linear model. This model is then used to classify the new training examples. In our work, we used LIBSVM [5] primarily because it can deal with multiple classes. To fine tune the SVM performance, we used C-support vector classification (C-SVC) with Radial Basis Function (RBF) as the kernel function. For tunable parameters, viz., cost and gamma (γ), we searched the space to find their optimal values.

2) Random Forest: It uses decision trees as a base classifier. The classifier combines the output of several decision trees and uses voting to predict the class for a given input data. It builds several classification trees, with each tree voting for a class. Random forest chooses the class label to be the one with the maximum votes. For more information on Random Forest, the reader is referred to [3].

3) Adaboost algorithm: Adaboost algorithm uses pseudo-loss as an error measure for training N observations with c classes. Since we consider three thermal comfort classes, we use a multiclass extension of Adaboost, i.e., Adaboost.M2 algorithm, in this paper. Adaboost.M2 uses decision stumps as the base classifiers. The input to the algorithm is training tuples (x_n, y_n), where x_n represents the features extracted from the data and y_n is a value representing the class of x_n. This algorithm, like other boosting algorithms, converts the weak classifiers to the strong ones. For details on Adaboost.M2, one may refer to [17].

IV. Dataset

We used the publicly available RP-884 database [7] to evaluate the performance of our approach. RP-884 contains data relevant to human thermal comfort and was developed as part of ASHRAE RP-884 project "developing an adaptive model of thermal comfort and preference". Specifically, ASHRAE funded a series of field studies that cover four different climatic zones, mainly focusing on office settings, following standardization as suggested by ASHRAE. The resulting database contains approximately 21,000 rows of data, collected from 160 different buildings in four different continents [8].

Since our goal is to analyze data collected from HVAC buildings only, the first step in our analysis was to extract data corresponding to HVAC buildings. Approximately half of the dataset in [7] contains HVAC data, comprising approximately $12,000$ rows of raw data. These data are obtained from 16 individual studies, covering four seasons (summer, winter, dry and wet), across 4 continents. Each row of the database contains physiological parameters (e.g., age, clothing and metabolic rate), indoor climate observation, user thermal sensation feedback, outdoor climate observation (collected at various intervals), and several thermal indices (including PMV). In addition, it also contains the AMV (actual mean vote), which is used as the ground-truth thermal comfort. For further information, one may refer to [7], [8].

A. Preprocessing

After feature extraction, the following preprocessing steps are taken to prepare data for input to the classifier.

- Some rows do not contain the complete set of features. All such rows are removed so that the classification algorithm only uses the rows with the complete set of features.

- The data for each feature is normalized to have zero mean and unit standard deviation. This allowed us to convert the data to the same scale, yet retains the shape properties of the data.
- The AMV in each row is a real value in $[-3, 3]$. Following [11] (and the references with in), we map the AMV into three classes: "Uncomfortably warm (1)" when AMV is larger than α, "Neutral (0)" when AMV is in $[-\alpha, \alpha]$, and "Uncomfortably cold (-1)" when AMV is smaller than α. We refer to α as *comfort threshold*, and set $\alpha = 0.5$ or 1.0 in the rest of the paper.
- The original data set contains an unbalanced number of votes for each of the classes. To resolve any over training issues, equal number of instances are extracted from each class. Specifically, when $\alpha = 0.5$, there are respectively 944, 1124, and 1178 instances of classes 1, 0, -1 in the original data set. To maintain the same number of instances in each class, 944 instances are extracted for each class, leading to a total of 2832 instances. When $\alpha = 1.0$, there are respectively 352, 2460, and 434 instances of classes 1, 0, -1 in the original data set. To maintain the same number of instances in each class, 352 instances are extracted for each class, leading to a total of 1056 instances.

V. Performance Evaluation

In the following, we first present the evaluation results using ASHRAE's recommended range of the comfort threshold i.e. $\alpha = 0.5$. Additionally, as $\alpha = 1.0$ represents thermally acceptable range, we also present the results when $\alpha = 1.0$. In each case, we first present the results on feature set optimization, and then present the classification results under the three machine learning algorithms to find the algorithm that provides the best performance. We further compare the performance of the machine learning algorithm with that of Fanger's model. At the end, we present the overall accuracy of the various algorithms using 10-fold cross validation.

A. Performance Results When $\alpha = 0.5$

When $\alpha = 0.5$ and using feature set F in (1), preprocessing (see Section IV-A) results in a data set with 944 instances for each class and a total of 2832 instances over all the three classes. We randomly choose 70% of the instances as training set and the remaining 30% as testing set. In the following, we first investigate what subset of F leads to the best performance, and then present the evaluation results using the best feature set. We only report the results under SVM and Random Forest since both of them outperform Adaboost; we briefly report the accuracy of Adaboost in Section V-C.

1) Feature Set Optimization: Let us consider all the subsets, $F' \subseteq F$, $F' \neq \emptyset$. For each F', we use the corresponding training set to obtain the thermal comfort model for each machine learning algorithm. We then apply the model to the testing set and choose F' that provides the best accuracy.

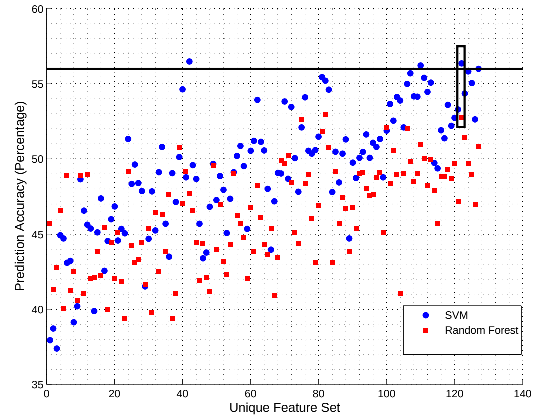

Fig. 2. Feature set optimization results, $\alpha = 0.5$.

Fig. 2 presents the prediction accuracy versus the choice of the feature sets (i.e., we index F' from 1 to $2^7 - 1 = 127$ since there are 7 features in F and we exclude the null set) when using SVM. The subset F' that provides the best accuracy is

$$F^* = [A, C, M, T_a, A_t, T_o] = F \setminus \{T_m\} \tag{2}$$

That is, the optimal feature set does not include mean radiant temperature T_m in F. We verified this results using multiple simulation runs (by choosing the training and testing sets differently). We also verify that F^* to be the optimal feature set when using Random Forest, as is shown in Fig. 2. The rest of the paper uses F^* as the feature set for the machine learning algorithms.

2) Accuracy Results: We next compare the performance of the various machine learning algorithms to identify the best machine learning algorithm for our study. Table I lists the confusion matrix for SVM and Random Forest, where 1, 0 and -1 represent the three classes, "uncomfortably warm", "comfortable", and "uncomfortably cold", respectively. The diagonal of the matrix represent the correct classification results. The overall accuracy of these two classifiers over the three classes is approximately 56% and 52%, respectively. We observe that SVM outperforms Random Forest for classes -1 and 1, while performs worse than Random Forest for class 0. To compare the accuracy between SVM and Random Forest across all the three classes, we use macro-average modified test [15]. This test treats all the classes equally, and computes the geometric mean and average correct classification rate over the classes. Specifically, it calculates the following test value

$$0.75 \left(\sum_{i=1}^{c} P_i P_{ii} \right) + 0.25c \left(\sqrt[c]{\prod_{i=1}^{c} P_i P_{ii}} \right),$$

where P_i is the probability of class i and P_{ii} is the probability of correct classification, i.e., the fraction of correct classification results for prior class i, and c is the number of classes (i.e., $c = 3$ in our context). Using SVM confusion matrix data from Table I, the macro-average modified test value is 0.54

for SVM and 0.50 for Random Forest, thus demonstrating that SVM slightly outperforms Random Forest.

Classifier			**Predicted Label**		
			-1	**0**	**1**
SVM	**Actual Label**	**-1**	235	40	14
		0	111	62	88
		1	55	61	184
Random Forest		**-1**	196	68	25
		0	90	84	87
		1	47	98	155

TABLE I

CONFUSION MATRIX FOR SVM AND RANDOM FOREST, $\alpha = 0.5$.

We next compare the performance of SVM and Random Forest with that of Fanger's model using McNemar test [20]. This test is applied to a 2×2 contingency table to determine whether the classification errors of two algorithms are statistically different. Table II lists the 2×2 contingency table between the SVM classifier and Fanger's model, where 1 represents the number of "correctly classified" instances and 0 represents the number of "incorrectly classified" instances.

Fanger's Model / **SVM**	**1**	**0**	**Total**
1	127	354	481
0	156	213	369
Total	283	567	850

TABLE II

CONTINGENCY TABLE BETWEEN THE SVM CLASSIFIER AND FANGER'S MODEL, $\alpha = 0.5$.

Let n_A denote the number of classification errors made by Fanger's model but not by SVM classifier, and let n_B denote the number of errors made by SVM classifier but not by Fanger's model. From Table II, $n_A = 354$, $n_B = 156$, and n_B is significantly smaller than n_A. We next use McNemar test to determine the statistical significance that n_B is smaller than n_A. Specifically, the null hypothesis, H_o, is

$$H_o : n_A = n_B, \tag{3}$$

whereas the alternate hypothesis, H_a, is

$$H_a : n_B \ll n_A. \tag{4}$$

McNemar's test [20] is given as:

$$\chi^2 = \left(\frac{|n_A - n_B| - 1}{\sqrt{n_A + n_B}} \right)^2 \tag{5}$$

Substituting the values in the "discordant cells" in Table II yields

$$\chi^2 = \left(\frac{|354 - 156| - 1}{\sqrt{354 + 156}} \right)^2 \approx (8.7)^2 \approx 75.7 \tag{6}$$

When the number of discordants (i.e., n_A and n_B) is sufficiently large, χ^2 has a chi-squared distribution with 1 degree of freedom. We therefore determine the statistical significance by evaluating the p-value for the χ^2 distribution. In this case, using table from [16], for 1 degree of freedom, the p-value is found to be < 0.01. Therefore, we reject the null hypothesis and accepts the alternate hypothesis that $n_B \ll n_A$.

Fanger's Model / **Random Forest**	**1**	**0**	**Total**
1	121	314	435
0	162	253	415
Total	283	567	850

TABLE III

CONTINGENCY TABLE BETWEEN THE RANDOM FOREST BASED CLASSIFIER AND FANGER'S MODEL, $\alpha = 0.5$.

We next compare the classification errors of the Random Forest classifier and Fanger's model. Table III lists contingency table between these two approaches. Substituting data from "discordant cells" in the table into equation (5) yields

$$\chi^2 = \left(\frac{|314 - 162| - 1}{\sqrt{314 + 162}} \right)^2 \approx 6.9^2 \approx 47.6 \tag{7}$$

Again, for one degree of freedom, the p-value is < 0.01, indicating a strong statistically significant result to reject the null hypothesis and accept the alternate hypothesis.

B. *Performance Results When $\alpha = 1.0$*

When $\alpha = 1.0$, preprocessing (see Section IV-A) results in a data set with 352 instances for each class and a total of 1056 instances over all the three classes. We randomly choose 70% of the instances as training set and the remaining 30% as testing set. In this case, we again find that the feature set F^* defined in (2) provides the best performance. In addition, macro-average modified test again indicates that SVM slightly outperforms Random Forest. In the following, we only present the results obtained using the SVM classifier. Table IV lists the confusion matrix for matrix for the SVM classifier. The overall accuracy over the three classes is approximately 75%. Table V lists the 2×2 contingency table between the SVM classifier and Fanger's model. Using McNemar's test, we have

$$\chi^2 = \left(\frac{|150 - 10| - 1}{\sqrt{150 + 10}} \right)^2 \approx 10.9^2 \approx 118.8 \tag{8}$$

Again the p-value indicates a statistically significant result to reject the null hypothesis and accept the alternate hypothesis.

Predicted Label / **Actual Label**	**-1**	**0**	**+1**
-1	70	37	5
0	10	88	0
+1	11	15	81

TABLE IV

CONFUSION MATRIX FOR THE SVM CLASSIFIER, $\alpha = 1.0$.

Fanger's Model / SVM	1	0	Total
1	89	150	239
0	10	68	78
Total	99	218	317

TABLE V

CONTINGENCY TABLE BETWEEN THE SVM CLASSIFIER AND FANGER'S MODEL, $\alpha = 1.0$.

C. 10-Fold Cross Validation

We now use 10-fold cross validation to evaluate the prediction accuracy of the various approaches. Specifically, for a given α, we divide the corresponding data set into 10 equal-size subsets, and construct ten training and testing sets, each training set uses 90% of the data and the testing set uses the remaining 10% of the data. Table VI summarizes the results for both $\alpha = 0.5$ and $\alpha = 1.0$. We observe that all the machine learning based algorithms significantly outperforms the Fanger's model, and the SVM based classifier outperforms the Random Forest and Adaboost based classifiers.

Threshold / Classifier	$\alpha = 0.5$	$\alpha = 1.0$
SVM	56.7%	76.7%
Random Forest	52.1%	74.1%
Adaboost.M2	51.1%	61.4%
Fanger's Model	33.2%	35.4%

TABLE VI

10-FOLD CROSS VALIDATION

VI. CONCLUSION AND FUTURE WORK

We have described a vision of a decision support system that predicts an individual's thermal comfort in real-time and suggests corresponding actions. The real-time prediction and the corresponding suggested action can significantly improve the overall thermal comfort and health of individuals, especially of senior citizens, as aging often leads to less acute and delayed thermal sensation. To realize this vision, we proposed a novel machine learning based approach to learn an individual's thermal comfort model. Specifically, this approach identifies the best set of features, and then learns a classifier that takes a feature vector as input and outputs a corresponding thermal sensation class. Evaluation using a publicly available large-scale data set demonstrates that when using SVM classifiers, the accuracy of our approach is 76.7% versus 35.4% for Fanger's model. In addition, hypothesis tests demonstrate that our approach outperforms Fanger's model with strong statistical significance. Last, our feature selection study indicated that two factors, a person's age and outdoor temperature that are not included in Fanger's model, play an important role in thermal comfort, which is a finding interesting in its own right.

As future work, we plan to test our approach in a retirement community. We are also planing to improve the proposed technique so that it can predict the missing values from the data. Additionally, we plan to use feature selection methods, such as regularization or discriminative based methods, to identify the optimal feature set for our approach and compare the results with other existing methods.

REFERENCES

[1] ASHRAE. Thermal comfort. *in: ASHRAE Fundamentals, American Society of Heating, Refrigerating and Air Conditioning Engineers, Atlanta*, 2009.

[2] S. Bin and H. Ke. Indoor thermal comfort pmv index prediction based on particle swarm algorithm and least square support vector machine. In *Intelligent System Design and Engineering Application (ISDEA), 2010 International Conference on*, volume 1, pages 857–860. IEEE, 2010.

[3] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[4] C. Candido and R. d. Dear. From thermal boredom to thermal pleasure: a brief literature review. *Ambiente Construído*, 12(1):81–90, 2012.

[5] C.-C. Chang and C.-J. Lin. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.

[6] C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[7] R. De Dear and G. S. Brager. Developing an adaptive model of thermal comfort and preference. *Center for the Built Environment*, 1998.

[8] R. J. de Dear and G. S. Brager. Thermal comfort in naturally ventilated buildings: revisions to ASHRAE Standard 55. *Energy and buildings*, 34(6):549–561, 2002.

[9] R. Dedear and A. Auliciems. Validation of the predicted mean vote model of thermal comfort in six australian field studies. *ASHRAE transactions*, 91(2B):452–468, 1985.

[10] D. W. DeGroot and W. L. Kenney. Impaired defense of core temperature in aged humans during mild cold stress. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 292(1):R103–R108, 2007.

[11] N. Djongyang, R. Tchinda, and D. Njomo. Thermal comfort: A review paper. *Renewable and Sustainable Energy Reviews*, 14(9):2626–2640, 2010.

[12] V. L. Erickson and A. E. Cerpa. Thermovote: participatory sensing for efficient building hvac conditioning. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 9–16. ACM, 2012.

[13] P. O. Fanger et al. Thermal comfort. Analysis and applications in environmental engineering. *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.

[14] P. M. Ferreira, S. M. Silva, A. E. Ruano, A. T. Negrier, and E. Z. Conceiçao. Neural network pmv estimation for model-based predictive control of hvac systems. In *Neural Networks (IJCNN), The 2012 International Joint Conference on*, pages 1–8. IEEE, 2012.

[15] C. Ferri, J. Hernández-Orallo, and M. A. Salido. Volume under the ROC surface for multi-class problems. In *Machine Learning: ECML 2003*, pages 108–120. Springer, 2003.

[16] R. A. Fisher, F. Yates, et al. Statistical tables for biological, agricultural and medical research. *Statistical tables for biological, agricultural and medical research.*, (Ed. 3.), 1949.

[17] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.

[18] P. X. Gao and S. Keshav. Optimal personal comfort management using spot+. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, pages 1–8. ACM, 2013.

[19] A. H.-y. Lam and D. Wang. Carrying my environment with me: A participatory-sensing approach to enhance thermal comfort. *Proc. ACM BuildSys' 13*.

[20] Q. McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947.

[21] J. C. Semenza, C. H. Rubin, K. H. Falter, J. D. Selanikio, W. D. Flanders, H. L. Howe, and J. L. Wilhelm. Heat-related deaths during the July 1995 heat wave in Chicago. *New England journal of medicine*, 335(2):84–90, 1996.

[22] J. Van Hoof. Forty years of Fangers model of thermal comfort: comfort for all? *Indoor Air*, 18(3):182–201, 2008.