

A Personalised Thermal Comfort Model using a Bayesian Network

Frederik Auffenberg, Sebastian Stein and Alex Rogers

Agents, Interaction and Complexity Research Group

Electronics and Computer Science

University of Southampton, UK

{fa1c12,ss2,acr}@ecs.soton.ac.uk

Abstract

In this paper, we address the challenge of predicting optimal comfort temperatures of individual users of a smart heating system. At present, such systems use simple models of user comfort when deciding on a set point temperature. These models generally fail to adapt to an individual user's preferences, resulting in poor estimates of a user's preferred temperature. To address this issue, we propose a *personalised* thermal comfort model that uses a Bayesian network to learn and adapt to a user's individual preferences. Through an empirical evaluation based on the ASHRAE RP-884 data set, we show that our model is consistently 17.5-23.5% more accurate than current models, regardless of environmental conditions and the type of heating system used. Our model is not limited to a single metric but can also infer information about expected user feedback, optimal comfort temperature and thermal sensitivity at the same time, which can be used to reduce energy used for heating with minimal comfort loss.

1 Introduction

Reducing energy consumption and emission of greenhouse gases to mitigate the adverse effects of global warming and dwindling supply of fossil fuels has been posed as one of the biggest challenges of the 21st century. Domestic heating, accounting for 12% of the worldwide energy consumption [Gadonneix *et al.*, 2013], offers great potential for reducing overall energy consumption. This has led to the development of smart heating systems, which aim to reduce energy consumption by simplifying the interaction between the user and the heating system. The key component of such a system is a *smart thermostat*, which allows detailed heating schedules, offers modern user interfaces and additional features, such as mechanisms to predict occupancy and learn about the thermal environment, users' preferences and schedules. This information enables the smart thermostat to autonomously decide the preferred temperature set point and when to switch on and off the heating system, taking account of both the user's comfort and their preference for energy savings.

However, existing smart thermostats such as the Nest learning thermostat often fail to accurately learn an individual user's personal preferences and as a result fail to save significant amounts of energy [Yang and Newman, 2012; 2013]. In more detail, current systems often use models from the widely applied thermal comfort modelling standard ASHRAE 55; specifically either Fanger's static comfort model [Fanger, 1970] or the adaptive comfort model [de Dear and Brager, 1998]. The static model is based on the balance of heat loss and gain in the human body and works on input variables such as a person's clothing level, metabolic rate and variables describing the thermal environment (humidity, air velocity and operative temperature). The adaptive model expresses a person's preferred temperature as a linear relationship with the outside temperature. A key shortcoming of both is that they have been created for shared spaces such as offices and public spaces and seek to satisfy a large number of people at the same time. As such, they fail to capture the personal preferences of individual users.

To address this, emerging work in the area of artificial intelligence has started to look at modeling individual's heating preferences [Shann and Seuken, 2013; 2014; Lam *et al.*, 2014]. In particular, the rising computing power of modern thermostats and other devices that can be utilised to control a smart heating system, such as smartphones, allows the development of more sophisticated models that utilise machine learning techniques to adapt to individual user preferences. Such models are usually aimed at domestic spaces where the aim is to satisfy each single individual instead of the majority of a large number of office workers. Existing machine learning approaches are usually based on the adaptive comfort model discussed above and, in order to adapt to an individual's preferences, add extra, user-specific parameters, such as an individual base temperature, thermal sensitivity or cost-comfort pay-off, to the model [Shann and Seuken, 2013; 2014; Lam *et al.*, 2014]. These variables are then learned using feedback on the heating provided by the user.

However, these approaches tend to be impractical or fail to accurately model thermal comfort. While Fanger's model accurately represents heat gain and loss in the human body, it requires very specific, hard to obtain input variables such as metabolic rates, clothing levels and air velocities. In contrast, adaptive models oversimplify the problem by modelling the comfort temperature as a linear equation of only the out-

side temperature. Other physical factors whose effect has been proven in experiments for static models such as humidity [ASHRAE 55, 2010] are neglected. The aforementioned models and their shortcomings will be explained in more detail in Section 2. Another factor that is neglected in both models are seasonal adaptations by the user. Expectations of colder or warmer seasons and repeated exposure to their respective thermal conditions may diminish the user’s thermal sensitivity [Liu *et al.*, 2012; de Dear and Brager, 2002]. In reality, the actual impact of each factor is likely to vary between individuals. Existing approaches either only consider single factors or do not provide means to adapt to user’s preferences. Further, these approaches have either not been benchmarked at all [Shann and Seuken, 2013; 2014] or only in very limited conditions [Lam *et al.*, 2014].

Lastly, existing models are limited to either estimating a user’s vote on the current thermal environment [Lam *et al.*, 2014] or on estimating the temperature at which the user will feel most comfortable (optimal comfort temperature). By combining both of these outputs, it would be possible to determine the range of acceptable temperatures for a user. Rather than heating/cooling exactly to the optimal temperature, the agent only needs to keep the temperature within the acceptable range. This would enable it to turn off the heater/AC more often, resulting in energy savings [Auffenberg *et al.*, 2015]. Further, comfort ranges of different users could be utilised when trying to satisfy multiple occupants by finding the overlap of their comfort ranges.

To address the shortcomings of the existing approaches, in this paper, we develop and evaluate a comfort model that is capable of learning the user’s preferences from minimal feedback using only easily obtainable data. Based on the learned preferences, the model is able to accurately predict a user’s individual comfort level as well as the user’s optimal comfort temperature at any given point in time in arbitrary climate conditions. This enables the model to determine a user’s current thermal sensitivity, which, in turn, determines the boundary temperatures that are still acceptable for the user.

In more detail, we combine principles of static comfort models with principles of adaptive models to create a general comfort model. In this general model, we identify and extract user-specific variables from classic approaches and parameterise the model with these. By translating the model into a Bayesian network, extensive learning capabilities are added. Through an empirical evaluation based on data taken from the ASHRAE RP-884 data set, we show that our personalised model outperforms existing approaches in most cases. Using this data set, the model was tested for naturally ventilated (NV) domestic spaces during summer and winter as well as heating, ventilation and air conditioning (HVAC) equipped office spaces during summer and winter. In summary, this work advances the state of the art in the following ways:

1. We combine existing static and adaptive comfort models to create a novel, general model that performs well in varying climate conditions.
2. We create a model that allows to infer optimal comfort temperature and a user’s vote at the same time, enabling to learn a user’s current thermal sensitivity

3. Using Bayesian networks, we add learning capabilities that enable the model to adapt to users’ individual preferences to give more accurate estimates of those users’ optimal comfort temperature and comfort level in different conditions
4. We empirically evaluate our model using data from the ASHRAE RP-884 data set and show that it gives 17.5-23.5% (up to 30% in some cases) more accurate estimates of a user’s comfort level

The remainder of this paper is structured as follows. First we discuss the existing comfort models that we build on in this work in Section 2. We then introduce our model in Section 3, followed by the empirical evaluation of the model in Section 4. Section 5 concludes and discusses future work.

2 Thermal Comfort Models

Fanger *et al.* introduced the first thermal comfort model in 1970 [Fanger, 1970], which, with slight modifications, is still used to this day as the static thermal comfort model defined in ASHRAE Standard 55 [ASHRAE 55, 2010]. Fanger’s model is built around heat balance in the human body. Thermal comfort is defined as the equilibrium of heat gain due to metabolism and heat loss of the body to the environment. The main measurements are the predicted mean vote (PMV) and the predicted percentage dissatisfied (PPD). The PMV denotes the expected mean vote of a group of people on the thermal environment based on the 7-point comfort scale shown in Table 1. The PPD describes the percentage of people dissatisfied with the thermal environment.

VOTE	THERMAL SENSATION
3	too hot
2	warm
1	slightly warm
0	comfortable
-1	slightly cool
-2	cool
-3	too cold

Table 1: 7-point thermal comfort scale

Fanger’s model considers five input variables: operative temperature, air speed, relative humidity, metabolic rate and clothing level. The operative temperature is preferred to simple air temperature, as it combines both air temperature and radiant temperature to give a more accurate estimate of the perceived temperature. While operative temperature and relative humidity are easy to obtain, the other variables pose problems in practice. Air speeds may vary within a single room [Erickson and Cerpa, 2012], so the air velocity at the exact position of the individual would have to be measured. This is impractical in all but very controlled environments. Metabolic rate and clothing level may be subject to irregular variations, making them hard to estimate [Peeters *et al.*, 2009]. In addition to these problems, Fanger’s model neglects possible adaptations by the user. Those adaptations are usually of physiological (acclimatisation), behavioural (modification

of clothing and other heating controls) and psychological (expectations of temperatures) nature [Liu *et al.*, 2012].

Adaptive models [de Dear and Brager, 1998] try to take behavioural adaptive measures such as opening windows, turning on a fan or changing of clothing into account. This is usually done by modelling the user’s optimal comfort temperature in relation to the outside temperature. Typically, the colder it is outside, the more adaptive measures a user will take to stay warm. Similarly, as it gets warmer outside, a user will take more measures to stay cool. As a result, the optimal comfort temperature T_{opt} can be modelled as a linear function of the outside temperature T_{out} shown in Equation 1 (taken from [de Dear and Brager, 1998]).

$$T_{opt} = 0.255 T_{out} + 18.9 \quad (1)$$

While accounting for possible adaptations by the user, adaptive models neglect the influence of other factors defined in the static comfort model. To address the shortcomings of both approaches, we create a general personalised thermal comfort model that combines easily obtainable input variables from existing models into a single, more complete model. This is further explained in the following section.

3 A Bayesian Network for Thermal Comfort

We now introduce our personalised thermal comfort model that uses a Bayesian network to learn a user’s preferences in order to predict their optimal comfort temperature and vote at any given time. We combine the human-body centered approach of static models with the outdoor environment based approach of adaptive models. Our model consists of three components: one to calculate the user’s optimal comfort temperature based on different factors, one to translate the comfort temperature into a vote on the current thermal environment and one that calculates the current influence of adaptations on the user’s optimal comfort temperature. The outputs of the model are the user’s optimal comfort temperature T_{opt} , describing the temperature at which the user feels most comfortable, the user’s vote T_{vote} , quantifying how dissatisfied a user is with the thermal environment and the user’s thermal sensitivity γ_v , describing how much the actual temperature can deviate from the user’s optimal comfort temperature.

Our model combines the static model, stripped down to reliable, easily obtainable inputs (namely the operative temperature and humidity), with an extension of the adaptive model to account for behavioural adaptations as well as seasonal adaptations [de Dear and Brager, 2002]. To transformation it into a Bayesian network we simplify relationships between variables to those that either increase or decrease the comfort temperature. As a result, the comfort temperature is calculated by adding and subtracting different factors from a neutral temperature of exposition.

For simplicity, the model has been broken down into two parts: the general comfort model (Figure 1) and its adaptive parts (Figure 2). The general comfort model contains the main equation for calculating the comfort temperature as well as the transformation of the comfort temperature into a user vote and will be discussed in detail in Section 3.1. The adaptive parts of the model show the detailed calculation of the influence of adaptive measures and are explained in Sec-







SYMBOL	MEANING
	Latent Variables are variables that cannot be observed directly and need to be inferred
	Observed Variables are variables that can either be observed directly or calculated using variables without further relevance for the model
	Model Parameters are variables that directly describe user preferences and are learned by the model. Model parameters are modelled as a Gaussian with a prior mean and precision. The priors for the mean have a Gaussian distribution, the priors for the precision a gamma distribution.
	Noisy Variables are expected to be noisy due to their user-centric nature. To compensate for such noise, Gaussian noise with a fixed precision is added to such variables.
	Factors define the operation which is used to calculate a variable (outgoing edge) based on the factor’s inputs (incoming edges).
	Plates denote sets of observations. The amount of observations is denoted by the letter in its bottom right corner.
γ_{var}	Variables named γ_{var} describe the user-specific scaling for another variable “var”.
var_{γ}	The user-adjusted value of variable “var” that has been scaled with its γ_{var} counterpart.

Table 2: Notions in the model

tion 3.2. Table 2 lists the different types of nodes and variables in the figures and explains their meanings.

3.1 The General Comfort Model

The general comfort model, shown as a factor graph in Figure 1, contains variables which directly influence the user’s optimal comfort temperature and the resulting votes. It consists of the calculation of the user’s optimal comfort temperature T_{opt} and the resulting comfort vote T_{vote} . The user’s optimal comfort temperature represents the temperature at which a user feels most comfortable and is comparable to the temperature calculated with adaptive models. The model has two different plates, K and N. Plate N contains all training observations that include user feedback. These are used to train the model and learn its parameters. In a real system, a training observation would be created as soon as the user provides feedback in some form, such as manually adjusting the set point. Plate K contains inference observations. These are triggered by the heating system itself when it has to decide on a set point temperature. As opposed to training observations, inference observations do not include user feedback.

The general model can be split up into two different parts: the part calculating the optimal comfort temperature and the part calculating the resulting vote by the user. The former consists of all variables and factors above T_{opt} , the latter consists of all variables on the same level or below T_{opt} .

Calculating the optimal comfort temperature

The optimal comfort temperature, T_{opt} , is calculated as a combination of the base temperature, T^* , adaptations by the user, a_{γ} , and effects of humidity, h_{γ}). The base temperature, T^* ,

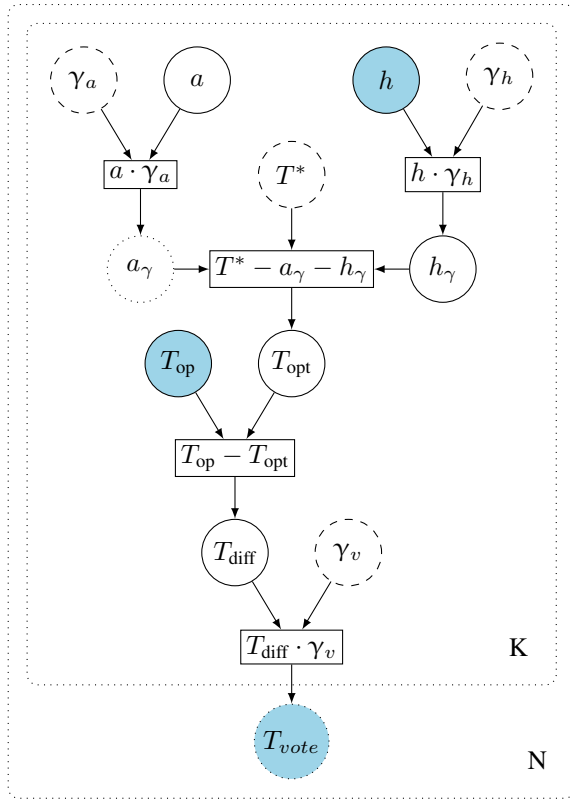


Figure 1: The general comfort model

describes the user’s theoretical comfort temperature in neutral conditions. Neutral conditions in this case means conditions during which other influences are either neglectable or cancel each other out. Humidity lowers the comfort temperature as the higher the humidity, the less efficiently the body’s natural cooling mechanism through evaporation of sweat works. As for adaptations, there are two cases: those to gain heat and those to lose heat. The former (e.g. adding clothing or increasing activity) allow a lower operative temperature. In contrast, the latter (e.g. turning on a fan) allow higher operative temperatures. The two different kinds of adaptations are represented by positive (heat gain) and negative (cooling) values of a_γ .

The three parameters, T^* , h_γ and a_γ , are user-specific. While T^* is a standalone variable, adaptation, a_γ , and humidity, h_γ , are scaled with user-specific scale factors (γ_a and γ_h respectively) of their observed or calculated counterparts (a and h respectively). The unscaled adaptation value a is based on a general adaptation formula that will be further described in Section 3.2. The unscaled humidity h describes the measured relative humidity inside the room.

Calculating the user’s vote

Thermal comfort is assessed as the deviation of the actual temperature from the user’s optimal comfort temperature as suggested by [Rogers *et al.*, 2011]. The vote T_{vote} on the current thermal environment is therefore based on the deviation T_{diff} of the actual operative temperature T_{op} from the optimal comfort temperature T_{opt} . The absolute deviation is translated

into a vote by multiplying it with a scaling factor describing the user’s thermal sensitivity γ_v , which can be learned from data. By manually setting the scaling factor, various common scales, such as the ASHRAE 7-point scale, can be supported by the model. By learning it, the model can compensate for different thermal sensitivities of users.

3.2 Adaptive Components

To cover a variety of adaptations by the user, the model includes a detailed section for adaptations (see Figure 2). As opposed to existing adaptive models, our model accounts for both psychological adaptations and behavioural adaptations. Physiological adaptations by the human body are not modelled separately. This is because some physiological adaptations like shivering are reactions to extreme conditions which should not be considered by the model. Further, other physiological factors (e.g. sweating) are already covered by the human-body centered approach of the static model.

Psychological adaptations are generally hard to quantify [Liu *et al.*, 2012]. Because of this, we currently restrict psychological adaptations to seasonal adaptations a_s , which reflect different expectations for the thermal environment by the user depending on the current season. For example during the colder seasons, people are expecting colder temperatures and are therefore more willing to accept them [de Dear and Brager, 2002]. To model this, we use Equation (2), which takes the current day of the year t_y as an argument:

$$a_s = \cos \frac{2\pi t_y}{365} \quad (2)$$

During colder seasons, the equation yields negative values up to -1 while during warmer seasons it yields positive values up to values of 1 . To adjust for conditions in the southern hemisphere, the result can be multiplied with -1 . As the amplitude of this effect might vary between different people and latitudes, the values are scaled with a learned factor γ_{a_s} .

Behavioural adaptations are modelled in a similar way to how existing adaptive models do this: as a linear relationship with the outside temperature T_{out} . In contrast to existing models, the slope γ_{a_b} of this relationship is learned from user feedback. Further, the base temperature is omitted, as it is already included in the core model as T^* .

The overall adaptation a is calculated by adding up both user-corrected parts for seasonal adaptations a_{γ_s} and behavioural adaptations a_{γ_b} .

3.3 Learning and Inference

As mentioned earlier, the model is implemented as a Bayesian network to add learning capabilities. Bayesian networks are directed acyclic graphs that represent sets of random variables and their conditional dependencies [Pearl, 1986]. We implement the model using the free library Infer.NET¹, and perform inference using expectation propagation (EP) [Minka, 2001]. EP is a more general version of belief propagation (also known as sum-product message passing) that works

¹Infer.NET - Microsoft Research (<http://research.microsoft.com/en-us/um/cambridge/projects/infernet/>)

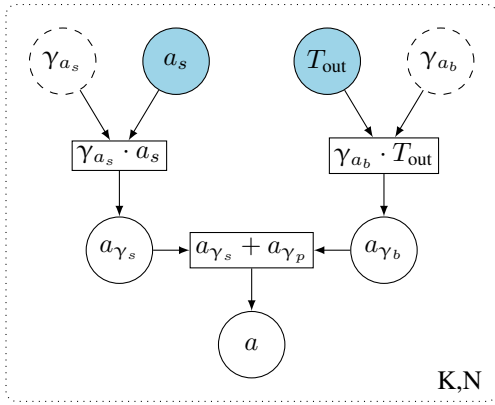


Figure 2: Adaptive part of the comfort model

with continuous variables of different probability distributions. The algorithm works on factor graphs where messages containing information about the current expected probability distribution of a node are sent between neighbouring nodes and factors.

Learning is performed by including the model parameters describing a user, namely T^* , γ_v , γ_h , γ_a , γ_{a_s} and γ_{a_b} , as additional nodes in the network. Performing inference for model parameters yields the updated, user-specific values for the given training data of the user. Replacing the priors of the model parameters with the learned values results in a user-specific model, based on which the optimal comfort temperature (T_{opt}) and user votes (T_{vote}) can be obtained.

To achieve best results, model parameters are fully relearned with every new training observation. As opposed to on-line learning where an updated model is kept and constantly updated based on the newest observation, relearning the entire model from scratch yields better results at the cost of speed [Bauer *et al.*, 1997]. The speed trade-off is acceptable as relearning model parameters is only triggered by new training observations, which we expect to happen about once a day during the initial training and less often afterwards. In addition, compared to how long the heating takes, the time to update the model is insignificant. Furthermore, single updates usually only result in small changes to the model, so updates to the model are not urgent.

4 Empirical Evaluation

To show the validity of our model and emphasize the need for more personalised models, we empirically evaluate it using existing longitudinal studies from the ASHRAE RP-884 project. In those studies users were asked to provide feedback on their thermal sensation using the ASHRAE 7-point scale. We test the model’s accuracy with respect to the amount of training observations. Overall, the model has been tested on 553 individuals in 10 different cities. The main parameters for each data set are shown in Table 3. Overall, the studies cover a wide range of scenarios, accounting for different seasons, ventilation systems and space types.

The Pakistan data set contains data for the cities of Karachi, Peshawar, Multan, Quetta and Saidu [Nicol *et al.*, 1994]. The

	PAKISTAN	ATHENS	SAN FRANCISCO
Subjects (s, w)	16, 15	31, 0	271, 220
Observations	50-150	65	up to 7
Time-span	1 week	10 - 60 days	5 days
Separate days	5 - 7	up to 10	up to 5
Consecutive days	yes	some	yes
Feedback scale	{-3..3}	{-3..3}	[-3, 3]
Ventilation	NV	HVAC	NV & HVAC
Space type	both	office	office

Table 3: Description of the different data sets. *Subjects* means the number of occupants during summer (s) and winter (w), *observations* the observation count for each occupant, *separate days* describes on how many separate days data was taken for each user and *space type* the usage of the building (office or domestic).

data for Saidu has been omitted due to extreme values (e.g. indoor temperatures of 14° during winter) which should not occur with automated heating controls. The San Francisco data set contains data for five locations in the San Francisco bay area (Berkeley, San Ramon, Palo Alto, San Francisco and Walnut Creek). In each data set, the indoor thermal environment is described by multiple values, of which we used the operative temperature, relative humidity inside the building, date and time. Note that due to the low observation count per individual, but high number of different individuals, the San Francisco data set was mainly used to show the general applicability of the model rather than its final solution quality.

As the ASHRAE RP-884 database only contains data about the highest and lowest temperature for each day, detailed weather data was obtained from Weather Underground². If no historical records for this particular year were present for a location, averages of other years were used. This was the case for most data points in the Pakistan and Athens data sets. For both, historical records from 2001 to 2014 were used. If no records for the exact hour were present, we performed a linear interpolation using the previous and next data point available. This was mainly the case for the city of Quetta, where for most dates only data for every six hours was available.

4.1 Benchmarks

To benchmark our model, we compared it to the existing, standardised approaches described in Section 2:

- Fanger’s static comfort model (PMV)
- The adaptive model

The approaches were compared based on the accuracy of their predictions for T_{vote} . While the PMV (similar to T_{vote}) for the static model was provided with the data sets, Equation 3 was used to translate the normal output T_{opt} into a vote T_{vote} for the adaptive model.

$$T_{vote} = 0.29(T_{op} - T_{opt}) \quad (3)$$

This equation multiplies the difference of the operative temperature T_{op} from the optimal comfort temperature T_{opt} with

²Weather Underground - <http://wunderground.com>

DATA SET	PMV		ADAPTIVE	
Pakistan (s)	27%	(0.4)	2.6%	(0.03)
Pakistan (w)	28%	(0.4)	9%	(0.1)
Athens	30%	(0.35)	25%	(0.27)
San Francisco (s)	18%	(0.2)	24%	(0.28)
San Francisco (w)	19%	(0.24)	30%	(0.45)
Overall	23.5%	(0.315)	17.5%	(0.21)

Table 4: Accuracy gains (absolute in parantheses) for the predicted T_{vote} of our model vs. PMV and adaptive models

the average learned thermal sensitivity of 0.29. This value corresponds with values that can be extracted from [de Dear and Brager, 1998].

The data was divided by single individuals into separate subsets. For those subsets, cross validation was performed using each single data point as an inference observation in separate evaluation runs, using random data points from the remaining data as training observations. For each single evaluation run, different amounts of training observations have been tested. For data sets with a lot of data points per individual (Pakistan and Athens), the amount of training observations was increased in steps of 2 for values between 2 and 30. For San Francisco, the amount was increased in steps of 1 between 1 and the maximum possible observation count (number of observations - 1).

The evaluation for a single data point consisted of two steps. First, the model was trained using the training observations. After that, the trained model was fed data from the evaluation data point excluding the user’s feedback. The feedback was inferred using the model and the squared error of the result was logged. From all single results, the root mean square error (RMSE) and standard error σ were calculated, which will be discussed in the next section.

4.2 Empirical Results

Table 4 shows accuracy gains for predictions of T_{vote} achieved by our model compared to the PMV and adaptive model. One can see that apart from the Pakistan data set (see Figure 3a), our model achieves significant accuracy gains (18-30% smaller prediction error for T_{vote}) in comparison to the other approaches (Figures 3b and 3c). A possible explanation for the low accuracy gains on the Pakistan data set is that it contains many spurious 0 votes regardless of the thermal en-

PARAMETER	VALUE RANGE	μ	σ
T^*	[19.86, 25]	22.02	0.99
γ_v	[0.006, 0.96]	0.29	0.23
γ_h	[2.6, 3.37]	2.9	0.135
γ_a	[0.034, 0.91]	0.61	0.25
γ_{a_s}	[0.93, 1.32]	1.04	0.046
γ_{a_b}	[-0.43, 0.063]	-0.29	0.116

Table 5: Learned parameter statistics (μ = average, σ = standard deviation)

vironment, possibly due to the participants misunderstanding the trial protocol, which hinders the learning process. Further, our model seems to benefit from the continuous scale used in the San Francisco data set, in which it reached a similar solution quality after only 4 observations as opposed to 6-8 in the other data sets (see Table 4). In general, our model typically converges after 10 observations (see Figure 3c).

Table 5 shows the value range, average μ and standard deviation σ of the learned parameters. Apart from seasonal adaptations, γ_{a_s} , one can see that all parameters are well spread out over their value ranges, indicating their importance to represent individual users accurately. The low variance in values for γ_{a_s} is a result of the data sets being limited to either winter or summer. In a data set spanning over longer times, we expect this parameter to gain importance. The large variety in thermal sensitivities, γ_v , of users shows that for some users significant energy savings can be achieved through mechanisms described in [Auffenberg *et al.*, 2015].

5 Conclusions and Future Work

In this work we presented a thermal comfort model that combines and simplifies existing models to only require easily obtained input parameters and utilises Bayesian networks to learn an individual user’s preferences.

Through empirical evaluations we showed that our model generally outperforms existing approaches by 17.5% - 23.5% after a short initial learning phase. Further, it enables inference of different information such as the expected feedback from the user about the heating, the optimal comfort temperature as well as the user’s thermal sensitivity which can be used by the heating system to reduce the energy used for heating and cooling [Auffenberg *et al.*, 2015].




Figure 3: RMSE of the predicted vote depending on observation count (with 2σ confidence interval)

In future work, we will use the learned information about single users to find compromises for environments with multiple occupants with different preferences. Besides, our current model could be extended to incorporate a cost-comfort trade-off function.

References

- [ASHRAE 55, 2010] Thermal Environmental Conditions for Human Occupancy (ANSI Approved), 2010.
- [Auffenberg *et al.*, 2015] Frederik Auffenberg, Sebastian Stein, and Alex Rogers. A Heating Agent using a Personalised Thermal Comfort Model to Save Energy. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pages 1799–1800, 2015.
- [Bauer *et al.*, 1997] Eric Bauer, Daphne Koller, and Y Singer. Update rules for parameter estimation in Bayesian networks. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, pages 3–13, Providence, Rhode Island, 1997.
- [de Dear and Brager, 1998] Richard J. de Dear and Gail Schiller Brager. Developing an adaptive model of thermal comfort and preference. In *ASHRAE Transactions*, volume 104, pages 145–167. ASHRAE, 1998.
- [de Dear and Brager, 2002] Richard J. de Dear and Gail S. Brager. Thermal comfort in naturally ventilated buildings: revisions to ASHRAE Standard 55. *Energy and Buildings*, 34(6):549–561, July 2002.
- [Erickson and Cerpa, 2012] Varick L Erickson and Alberto 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, 2012.
- [Fanger, 1970] Povl Ole Fanger. Thermal comfort. Analysis and applications in environmental engineering. In *Thermal comfort. Analysis and applications in environmental engineering*, page 244 pp. 1970.
- [Gadonneix *et al.*, 2013] Pierre Gadonneix, Younghoon David Kim, Kevin Meyers, Graham Ward, and Christoph Frei. World Energy Resources 2013. Technical report, World Energy Council, 2013.
- [Lam *et al.*, 2014] Abraham Hang-yat Lam, Yi Yuan, and Dan Wang. An Occupant-participatory Approach for Thermal Comfort Enhancement and Energy Conservation in Buildings Categories and Subject Descriptors. In *Proceedings of the fifth International Conference on Future Energy Systems*, pages 133–143, 2014.
- [Liu *et al.*, 2012] Jing Liu, Runming Yao, and Rachel McCloy. A method to weight three categories of adaptive thermal comfort. *Energy and Buildings*, 47(0):312–320, 2012.
- [Minka, 2001] TP Minka. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 362–369, 2001.
- [Nicol *et al.*, 1994] F Nicol, G N Jamy, O Sykes, M Humphreys, S Roaf, and M Hancock. A survey of thermal comfort in Pakistan. *July, Final report, Oxford Brookes University: School of Architecture, Oxford*, 1994.
- [Pearl, 1986] Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artificial intelligence*, 29(3):241–288, 1986.
- [Peeters *et al.*, 2009] Leen Peeters, Richard De Dear, Jan Hensen, and William D’haeseleer. Thermal comfort in residential buildings: Comfort values and scales for building energy simulation. *Applied Energy*, 86(5):772–780, 2009.
- [Rogers *et al.*, 2011] Alex Rogers, Sasan Maleki, Siddhartha Ghosh, and Nicholas R Jennings. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. In *Second International Workshop on Agent Technology for Energy Systems (ATES 2011)*, pages 71–78, 2011.
- [Shann and Seuken, 2013] Mike Shann and Sven Seuken. An Active Learning Approach to Home Heating in the Smart Grid. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pages 2892–2899. AAAI Press, 2013.
- [Shann and Seuken, 2014] Mike Shann and Sven Seuken. Adaptive Home Heating under Weather and Price Uncertainty using GPs and MDPs. In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems*, pages 821–828, 2014.
- [Yang and Newman, 2012] Rayoung Yang and Mark W Newman. Living with an Intelligent Thermostat : Advanced Control for Heating and Cooling Systems. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp ’12)*, pages 1102–1107, 2012.
- [Yang and Newman, 2013] R. Yang and M. W. Newman. Learning from a Learning Thermostat : Lessons for Intelligent Systems for the Home. *Proceedings of the 2013 ACM international joint conference on Pervasive and Ubiquitous computing (UbiComp 2013)*, pages 93–102, 2013.