

A Comfort-Based Approach to Smart Heating and Air Conditioning

FREDERIK AUFFENBERG, University of Southampton
STEPHEN SNOW, University of Southampton
SEBASTIAN STEIN, University of Southampton
ALEX ROGERS, University of Oxford

In this paper, we address the interrelated challenges of predicting user comfort and using this to reduce energy consumption in smart heating, ventilation and air conditioning (HVAC) systems. At present, such systems use simple models of user comfort when deciding on a set point temperature. Being built using broad population statistics, these models generally fail to represent individual users' preferences, resulting in poor estimates of the users' preferred temperatures. To address this issue, we propose the Bayesian Comfort Model (BCM). This *personalised* thermal comfort model uses a Bayesian network to learn from a user's feedback, allowing it to adapt to the users' individual preferences over time. We further propose an alternative to the ASHRAE 7-point scale used to assess user comfort. Using this model, we create an optimal HVAC control algorithm that minimizes energy consumption while preserving user comfort. Through an empirical evaluation based on the ASHRAE RP-884 data set and data collected in a separate deployment by us, we show that our model is consistently 13.2 to 25.8% more accurate than current models and how using our alternative comfort scale can increase our model's accuracy. Through simulations we show that using this model, our HVAC control algorithm can reduce energy consumption by 7.3% to 13.5% while decreasing user discomfort by 24.8% simultaneously.

CCS Concepts: • **Human-centered computing** → *Ubiquitous and mobile computing systems and tools*;

Additional Key Words and Phrases: User behaviour modeling and learning, Computational Sustainability, Agents, Machine Learning

ACM Reference Format:

Frederik Auffenberg, Steve Snow, Sebastian Stein, and Alex Rogers, 2016. A comfort-based approach to smart heating and air conditioning. *ACM Trans. Intell. Syst. Technol.* V, N, Article A (January YYYY), 20 pages.
DOI: 0000001.0000001

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 energy consumption. This has led to the development of smart energy systems, which aim to reduce energy consumption of heating, ventilation and air conditioning (HVAC), by either simplifying the interaction between the user and the HVAC system or by automation. The key component of such a system is a *smart thermostat*, which allows the definition of detailed HVAC schedules, offers modern user interfaces and additional features, such as mechanisms to predict occupancy and learn about the thermal environment, the users' preferences and their schedules. This information enables the smart thermostat to autonomously decide on an optimal temperature set point and when to switch the HVAC system on and off, taking account of both the user's comfort and their preference for energy savings.

However, existing smart thermostats such as the Nest Learning Thermostat or Honeywell's smart thermostat often fail to accurately learn a user's personal preferences and as a result fail to save significant amounts of energy [Yang and Newman 2012; 2013]. This is because 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 function of the outside temperature. A key shortcoming of both is that they have been created based on large population statistics, representing only average preferences over the whole population. However, as stated by Chappells et al., user comfort is fluid and can not be expressed by static numbers [Chappells and Shove 2005]. Thermal preferences can vary greatly between different people and change over time. Thus, when applied to individual users or small groups of people, these existing models often fail to give accurate estimates of these users' thermal comfort.

There are two main ways this issue has been addressed recently: by modelling and predicting individual users' comfort and by deploying participatory approaches that react in real-time to user feedback. Regarding the first, emerging work in the area of artificial intelligence has started to look at modeling an 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 or similar portable devices, allows the development of more sophisticated models that utilise machine learning techniques to adapt to an individual user's 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. Such 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. Reactive, participatory approaches in contrast rely on real-time feedback from occupants about their thermal environment [Purdon et al. 2013; Shetty et al. 2015]. In such approaches, occupants are given the opportunity to report their current comfort, desired changes to the set-point or similar metrics directly to the smart heating system. Based on this feedback, the heating system decides on a new set point.

However, both approaches have a number of drawbacks (discussed in more detail in Section 2). Predictive approaches usually either fail to model the user's comfort preferences accurately or are impractical. Specifically, while approaches based on Fanger's model accurately represent heat gain and loss in the human body, they require very specific, hard to obtain input variables such as metabolic rate, clothing level and air velocity. In contrast, approaches using the adaptive model tend to oversimplify the problem by modelling the comfort temperature as a linear function of only the outside temperature. Factors such as humidity [ASHRAE 55 2010], acclimatisation [Auliciems 1981] or seasonal adaptations [Liu et al. 2012; de Dear and Brager 2002], proven to have effects on thermal comfort, are neglected. For example, 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 the user's preferences. In the case of participatory approaches, the main drawback is the reliance on user continuous feedback. Research on user engagement in smart home systems suggests that users quickly lose interest in such systems and participation drops [Hargreaves et al. 2013; Snow et al. 2013]. In practice, this means that the system can be expected to perform worse over time as user engagement diminishes.

To address these shortcomings, in this paper, we develop and evaluate the Bayesian Comfort Model (BCM), a comfort model that is capable of learning the user's prefer-

ences from minimal feedback using only data that is easily obtainable from unobtrusive sensors. Based on the learned preferences, the model is able to accurately predict a user's individual comfort level (using scales such as the ASHRAE 7-point scale or the requested change scale introduced in Section 5.1) as well as the user's optimal comfort temperature at any given point in time in arbitrary climate conditions. Using this model, we further develop an HVAC control algorithm that minimizes energy consumption while retaining user comfort.

In more detail, we combine principles of the human physiology-based approach of static comfort models with principals of manual adaptation from adaptive models to create a more general comfort model. We identify and extract user-specific variables from classic approaches and parameterise the model with these. We translate the model into a Bayesian network, adding learning capabilities which allow it to adapt to individual user's preferences. We then propose an HVAC control algorithm that minimises energy consumption while using the BCM to minimise users' discomfort. Using weather forecasts and learned thermal properties of the house, the algorithm creates an HVAC schedule that minimizes energy consumption while keeping a comfortable indoor temperature for the occupants. Through an empirical evaluation using data from the ASHRAE RP-884 data set and data collected by us in a deployment assessing office workers' and library occupants' thermal comfort, we show that our personalised model predicts user comfort more accurately than existing models. We also investigate different feedback scales and their effect on the model's accuracy. Lastly, by simulating different households, we evaluate the energy saving potential of our heating algorithm when using our comfort model. In summary, this work advances the state of the art in the following ways:

- (1) We introduce a novel thermal comfort model with learning capabilities that enable it to adapt to individual users' preferences. The model is able to infer the users' comfort temperatures and ranges, allowing a smart thermostat to make more informed decisions about the set point.
- (2) We propose an intelligent heating algorithm that utilises the comfort model to minimise energy consumption without sacrificing user comfort.
- (3) We empirically evaluate the model using the ASHRAE RP-884 data set and data collected by us in a separate deployment. We show that the model generally gives 13.2% to 25.8% (in some cases up to 55%) more accurate estimates of a user's comfort level and can save up to 13.5% of energy while reducing discomfort by 24.8% using our heating algorithm.
- (4) Based on the empirical evaluation, we suggest the *requested change scale*, a scale to assess user comfort, that doubles our comfort model's performance as compared to the state of the art.

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 heating algorithm in Section 4. In Section 5 we empirically evaluate our model's accuracy, how it compares to existing models and the energy saving potential the model offers for our heating algorithm. Section 6 concludes and discusses future work.

2. RELATED WORK

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

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

Table I: 7-point thermal comfort scale

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 I. The PPD describes the percentage of people dissatisfied with the thermal environment.

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 try to take behavioural adaptive measures, such as opening windows, turning on a fan or changing of clothing, into account [de Dear and Brager 1998]. 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 factors defined in the static comfort model. To address the shortcomings of both approaches, we create a novel thermal comfort model that combines easily obtainable input variables from existing models into a single, more complete model.

Due to the lack of suitable comfort models, approaches automating heating and air conditioning (HVAC) tend to focus on predicting user occupancy while using either static comfort temperatures and ranges or variations of the models introduced above [Jo et al. 2013; Lu et al. 2010]. While these approaches can achieve significant energy savings, the inability to satisfy users’ individual preferences might limit acceptance among users. In addition, as we show in this paper, energy consumption can be further reduced by using personalised thermal comfort models.

One way to maximise user comfort despite the lack of suitable models is through participatory sensing. In such participatory approaches, users give live feedback about the thermal environment to the system, which in response calculates a new set point temperature based on the feedback. Usually, feedback is obtained through smartphone apps or computer interfaces. Shetty et al., for example, open a pop-up every 30 minutes asking their users about their current comfort levels [Shetty et al. 2015]. Using this feedback stream, individual users’ comfort ranges are constantly adjusted and the set

point changed accordingly. Purdon et al. provide their users with a smartphone app enabling them to submit their feedback while also providing them with information about past and current temperatures and other users' votes [Purdon et al. 2013]. The heating system adjusts the set point to achieve an equal balance of *hot* and *cold* votes.

A major drawback with such participatory approaches is their dependence on ongoing user feedback. Research into user engagement in smart energy systems suggests that users tend to lose interest in such systems and participation drops over time [Hargreaves et al. 2013; Snow et al. 2013]. For participatory approaches one can expect diminishing confidence in their decisions over time which is addressed by prompting users to provide feedback in some approaches [Shetty et al. 2015]. This however might be bothersome to the users and raises concerns about the real world applicability of this approach. In contrast, [Purdon et al. 2013] do not actively request feedback from their users. They assume that the prospect of increased thermal comfort is sufficient to keep users engaged. In practice however, as long as a reasonable level of comfort is sustained, user engagement is still likely to drop [Yang et al. 2014]. Using the example of the Nest learning thermostat, Yang et al. show that users become so disengaged that malfunctions of the system go unnoticed [Yang et al. 2014].

To overcome these issues, we propose a thermal comfort model that learns user preferences from minimal feedback. By learning an actual model of user comfort rather than a single comfort range, the model is able to accurately predict user comfort in varying surrounding conditions.

3. THE BAYESIAN COMFORT MODEL

We now introduce our personalised thermal comfort model. Our model uses a Bayesian network to learn users' preferences in order to predict their optimal comfort temperature and comfort vote, for example on the ASHRAE 7-point scale, at any given time. We combine the human-body centered approach of static models with the outdoor environment based approach of adaptive models into a more general and more precise model. In more detail, our model consists of three components: one to calculate the user's optimal comfort temperature based on a range of 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 transform 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 denoting the user's preferred temperature when all other factors are eliminated.

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







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 variables and their respective observations. The number of observations is denoted by the letter in its bottom right corner.
γ_{var}	Variables named γ_{var} describe the user-specific scaling for variable "var".
var_{γ}	The user-adjusted value of variable "var" scaled with its γ_{var} counterpart.
a_s	Describes seasonal adaptations (such as acclimatisation).
a_b	Describes behavioural adaptations (such as increasing clothing).
a	The sum of all adaptive measures taken by the user.
h	Relative humidity indoors.
T_{opt}	Optimal (= most comfortable) indoor temperature for a user.
T, T_{op}	Operative indoor temperature.
T_{out}	Outdoor temperature.
T_{diff}	Difference between T_{op} and the user's optimal temperature T_{opt} .
T_{vote}	The user's vote on the thermal environment.
ρ	Heating (ρ_r) or cooling (ρ_c) ratio.
ξ	Maximum energy consumption of heater (ξ^r) or cooler (ξ^c).
R	Heater output (in °C/hr).
C	Cooler output (in °C/hr).
Φ	Leakage rate (in 1/hr).
lb and ub	lower and upper bound of the comfort range.
γ_t	Energy consumption at time slot t (in kWh).

Table II: Model notations and nomenclature

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 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 practice, a training observation would be created as soon as the user provides feedback (for example by manually adjusting the set point). Plate K contains observations used for inferring other variables. These are created 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.

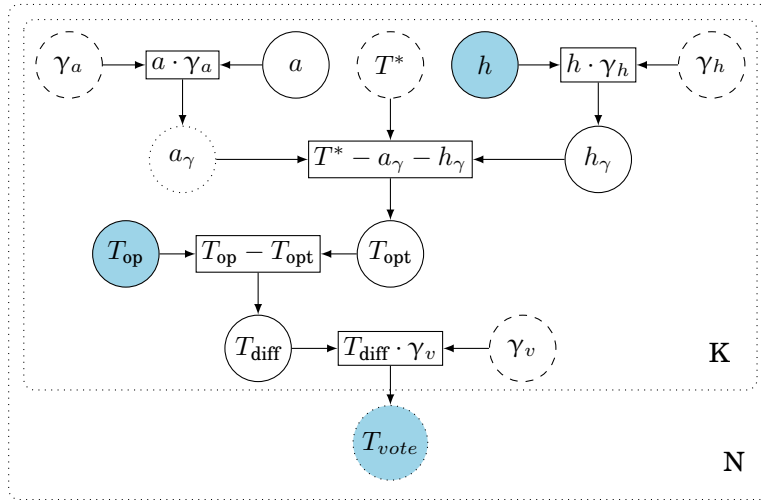


Fig. 1: The general comfort model

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} .

3.1.1. 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^* , describes the user's comfort temperature in neutral conditions where influences of other factors are either negligible 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.

3.1.2. Calculating the user's vote and comfort range. 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 operative temperature T_{op} from the optimal comfort temperature T_{opt} . The operative temperature is preferred in thermal comfort modelling as it incorporates radiated heat. It can be calculated as a combination of air temperature and mean radiant temperature. 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.

Depending on the scale used, the thermal sensitivity γ_v in conjunction with the optimal comfort temperature T_{opt} can be used to calculate the user's comfort range. The ability of inferring a comfort range rather than a single comfort temperature is crucial to our heating algorithm which utilizes said ranges to reduce energy consumption of the HVAC system. In case of the ASHRAE 7-point scale, ASHRAE standard 55 defines that the vote can deviate by 0.5 points from the neutral point 0 without causing noticeable discomfort to the user. This means that all temperatures with a predicted mean vote (equivalent to T_{vote}) between -0.5 and 0.5 can be considered comfortable for the user. Therefore, we base our calculations of a user's comfort range on these values. Rearranging the calculations shown in Figure 1 allows to calculate the operative temperature T_{op} depending on the desired user's vote:

$$T_{\text{op}} = \frac{T_{\text{vote}}}{\gamma_v} + T_{\text{opt}} \quad (2)$$

The lower bound lb of the comfort range can be obtained by setting $T_{\text{vote}} = -0.5$. Respectively the upper bound ub can be obtained by setting $T_{\text{vote}} = 0.5$.

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 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 occur when using our model to control the HVAC system. Other physiological factors (e.g. sweating) are already covered by the human-body centered approach of the static model.

Psychological adaptations are hard to quantify [Liu et al. 2012]. Because of this, we 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]. This is modelled by equation (3), which takes the current day of the year t_y as an argument:

$$a_s = \cos \frac{2\pi t_y}{365} \quad (3)$$

During colder seasons, the equation yields negative values up to -1 while during warmer seasons it yields positive values up to values of 1 . Since seasonal adaptations represent long-term adaptations throughout the year, a time resolution of one day is sufficient to accurately cover such adaptations. 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} . For example, in tropical regions where temperatures stay fairly constant throughout the year we expect this scaling factor to assume a low value.

Behavioural adaptations are modelled similar 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

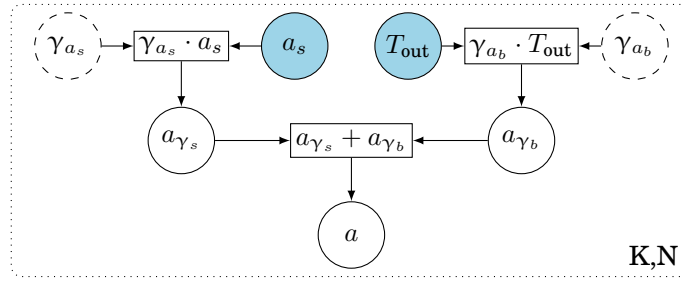


Fig. 2: Adaptive part of the comfort model

random variables and their conditional dependencies [Pearl 1986]. We implement the model using the free library Infer.NET [Minka et al. 2014], 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 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 up-to-date 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 phase and less frequently afterwards. On our test machine (Intel i5 3570), training the model took about 200ms for 50 observations. While smart thermostats have less powerful processors (ARM Cortex A8 in the Nest thermostat), training the model should still be in the range of only a few minutes at worst. 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. AN OPTIMAL HVAC CONTROL ALGORITHM

In this section, we introduce our optimal HVAC control algorithm that uses the Bayesian Comfort Model (BCM) introduced in Section 3 to minimise energy consumption of the HVAC system while maximising user comfort. As opposed to other approaches that try to reduce energy consumption by increasing the efficiency of the HVAC system [Wei et al. 2015], our approach focuses on optimising the actual temperature set point to reduce the temperature gradient between inside and outside. Using the comfort model, the algorithm is able to precisely assess how much it can let the indoor temperature follow the outdoor temperature without causing discomfort to the users. We model the HVAC scheduling problem as a linear programming problem of creating a profile of set-point temperatures for different time-slots of the day that minimises energy consumption. Using linear programming, we are able to calculate optimal HVAC schedules that require as little energy as possible. The set-point

temperatures are constrained to stay within the users' comfort ranges to guarantee a comfortable environment. The algorithm makes use of forecast weather data to predict comfort ranges and indoor temperatures.

4.1. Modelling Heating and Cooling Dynamics of the House

To accurately plan set-point temperatures, the HVAC control algorithm needs to be able to predict how switching on the heating or air conditioning will affect the room temperature. In this paper, we built upon a simple model introduced by [Rogers et al. 2013]. We chose this model over other more sophisticated models (for example [Tang et al. 2014]) due to its simplicity. Being a linear model it can be included into the optimisation with only minor increases in runtime. The model describes the HVAC system and room using a leakage rate $\Phi \in R^+$ (in 1/hr) and a heater output R (in $^\circ\text{C/hr}$). The leakage rate Φ describes the rate at which the indoor temperature adjusts to the outdoor temperature T_{ext} . The heater output describes by how many $^\circ\text{C}$ the indoor temperature increases per hour when the heater is running at full power. Rogers et al. limit their model to heating only. To support air conditioning, we have to add a third variable, the cooling rate C (in $^\circ\text{C/hr}$), which describes by how many $^\circ\text{C}$ the indoor temperature decreases per hour when the air conditioning is running at full power.

Most modern thermostats control the HVAC system by switching it on and off at varying frequencies to achieve different intensities. For more efficient optimisation, we simplify this behaviour and allow the HVAC system to run at fractions of its full power. We therefore introduce a heating ratio $\rho_r^t \in [0, 1]$ and cooling ratio $\rho_c^t \in [0, 1]$, describing at what fraction of their maximum output the heating and air conditioning are running. In practice, for a system where the heater or AC can either run or be switched off, a ratio of 0.5 in a time interval would mean that the heater/AC oscillates equally between on and off, effectively heating at half its maximum capacity. As a result, the indoor temperature T^{t+1} in the house at time $t + 1$ with a heater output R and cooler output C is calculated as follows:

$$T^{t+1} = T^t + \left[\rho_r^t R - \rho_c^t C - \Phi(T^t - T_{\text{out}}^t) \right] \Delta t \quad (4)$$

4.2. Formalization as a Linear Program

The main task of the HVAC control algorithm is to create an HVAC schedule that minimises the energy consumption of the HVAC system, while keeping the indoor temperature within the user's comfort range. For the formulation as a linear program, we add constraint 5 to the operative indoor temperature T_{op} , limiting its values to stay between the lower bound lb and upper bound ub of the comfort range.

$$lb \leq T_{\text{op}} \leq ub \quad (\text{constraint 5})$$

By including user comfort in the form of a constraint rather than adding it as another objective, the overall objective of the algorithm reduces to minimizing only the energy consumption. For that, the algorithm needs to be able to calculate the energy usage of the HVAC system. The energy consumption is mostly proportional to the heating ratio and cooling ratio. When heating at 50% of its maximum capacity, the energy usage can be expected to be close to 50% of the heater's maximum power consumption as well. We therefore model the energy consumption as the heating ratio or cooling ratio, multiplied by the maximum power consumption of the heater (ξ^r in kW) or air conditioner (ξ^c in kW). The overall energy consumption γ^t (in kWh) in a time-slot of length Δt ending at time t is calculated as follows:

$$\gamma^t = (\rho_r^t \xi^r + \rho_c^t \xi^c) \Delta t \quad (6)$$

In some cases, it can be beneficial or even necessary for the algorithm to plan ahead and look at several time slots simultaneously. For example, when dealing with large temperature gradients between inside and outside, the HVAC system might not be powerful enough to keep the indoor temperature within the user’s comfort range. This issue can be addressed by pre-heating or pre-cooling the house at an earlier time. Another example are variable hourly price rates for energy that are becoming more common due to the development of the *smart grid*. The algorithm should be able to pre-heat or pre-cool the house during times when energy is cheap. To allow the algorithm to plan ahead, we consider several time slots $t \in S$ simultaneously and minimise the aggregated energy consumption. S denotes the set of all time slots considered in the current computation. The resulting objective is shown in equation 7.

$$\min\left(\sum_{t=0}^S \gamma_t\right) \quad (7)$$

To factor in variable energy prices, it is sufficient to multiply each single energy consumption with the energy price p_t at that time slot. The objective function for minimizing cost is shown in equation 8.

$$\min\left(\sum_{t=0}^S \gamma_t p_t\right) \quad (8)$$

In the next section, we empirically evaluate how accurate predictions about user’s comfort ranges done by the BCM are. Using simulations, we assess the energy saving potential the BCM offers in combination with our HVAC control algorithm.

5. EVALUATION

In this section, we evaluate the accuracy of predictions made by the Bayesian Comfort Model (BCM) and the energy savings achieved using our HVAC control algorithm. The model’s accuracy is evaluated using data sets from the publicly available ASHRAE RP-884 database and data collected in a deployment carried out by us. Energy savings are assessed using simulations. The thermal properties of a typical house and HVAC system are chosen based on data presented by [Rogers et al. 2013].

5.1. Experimental Setup

To show the validity of our comfort model and emphasise the need for more personalised models, we empirically evaluate it and, using simulations, we demonstrate our heating algorithm’s energy saving potential when using the comfort model. We use data from existing longitudinal studies from the ASHRAE RP-884 project and from a deployment conducted by us. The ASHRAE RP-884 database is a standard database used to evaluate and create thermal comfort models and contains a number of different independent thermal comfort studies. In the studies we chose, users were asked to provide feedback on their thermal comfort using the ASHRAE 7-point scale. Since the BCM works with continuous feedback scales, in our own deployments, we use a modified version of the 7-point scale that uses continuous values ranging from -3 to 3 rather than discrete values. In an attempt to improve the interaction with the heating system for the user and get more accurate feedback, we test an alternative comfort scale. Rather than asking users how they feel, potentially disregarding that some users might prefer to feel slightly warm or cold, we use the *desired change scale* asking users how they would like the temperature to change.

In addition to information about indoor climate conditions and user votes, the model also requires data about outdoor weather conditions. To obtain this data, historical

	PAKISTAN	ATHENS	BAY AREA, CA	SOUTHAMPTON
Subjects (s, w)	16, 15	31, 0	271,220	11, 12
Observations	50-150	65	up to 7	up to 20
Time-span	1 week	10 - 60 days	5 days	14 - 28 days
Separate days	5 - 7	up to 10	up to 5	up to 5
Consecutive days	yes	some	yes	some
Feedback scale	{-3..3}	{-3..3}	[-3, 3]	[-3, 3]
Ventilation	NV	HVAC	NV & HVAC	NV & HVAC
Space type	both	office	office	office

Table III: 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).

records for the locations of the studies was downloaded from Weather Underground¹. If no historical records for the year and location of a study were present, 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 in the Pakistan data set where for most dates only data for every six hours was available.

We test the model’s accuracy with respect to the amount of training observations. Overall, we test our model on 576 different individuals in 10 different cities. The parameters of each data set are shown in Table III. Overall, these studies cover a wide range of scenarios like different seasons, ventilation systems and space types.

In the following sections we introduce the used data sets in more detail and explain the design and setup of our deployment.

5.1.1. The ASHRAE data sets. The Pakistan data set contains data for the cities of Karachi, Peshawar, Multan, Quetta and Saidu [Nicol et al. 1994]. We have omitted data from Saidu due to extreme values (e.g. indoor temperatures of 14°C during winter) that should not occur in practice when using an automated HVAC system. The 7-point scale used in the data set is not suitable for such extreme conditions since no fine-grained feedback can be given in such extreme conditions. The Bay Area data set contains data for five different cities in the area: San Francisco, Berkeley, Palo Alto, San Ramon 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. Due to the low observation count per individual but high number of different individuals, the Bay Area data set was mainly used to show the general applicability for a wide range of different users rather than to assess its final solution quality.

5.1.2. Thermal Comfort study in Southampton. To gather more data and test alternative feedback scales, we conducted a deployment at the University of Southampton, UK, to measure people’s thermal comfort levels. We developed posters inviting people to log how they feel about their thermal environment. On each poster further featured a temperature logger measuring the temperature every 4-10 minutes. We deployed a total of 172 of these posters at various locations around (1) a university library and (2) several offices on one floor of a naturally ventilated office building. Each poster featured the title “How’s the temperature?”, a large QR code and a URL address unique

¹Weather Underground - <http://wunderground.com>

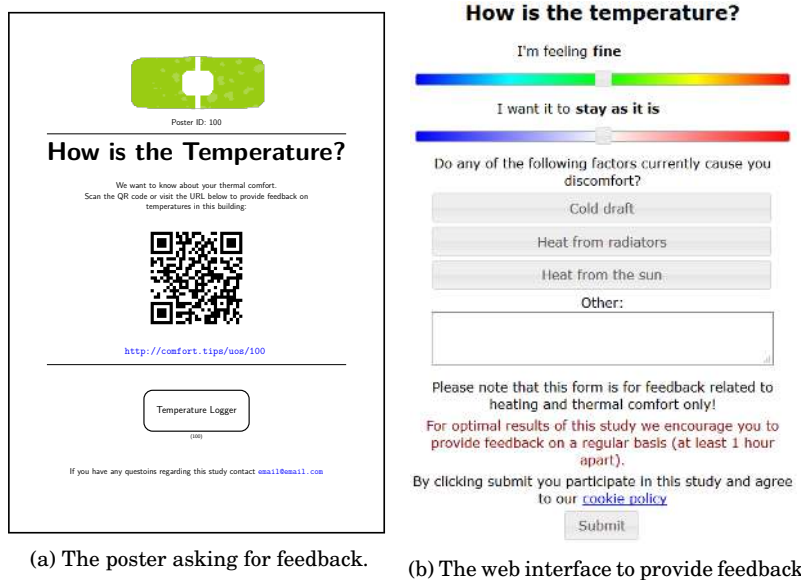


Fig. 3: The user interface of the study

to each poster (Figure 3a). Scanning the QR code or following the URL linked the user to a simple web-interface (Figure 3b) where they could log their thermal comfort.

There were two sliders on the user interface, the first representing the ASHRAE 7-point scale, the second our own *desired change scale*. Moving the sliders caused the label above to adjust to the respective value on the scale. For the ASHRAE scale, the labels were: I'm feeling [very cold, cold, slightly cool, neutral, slightly warm, warm, hot]. Under the assumption that some people might like it to be slightly warm or cold, we added the requested change scale, asking users how they want the temperature to change. This scale uses the following labels: I want it to [be much colder, be colder, be a bit colder, stay as it is, be a bit warmer, be warmer, be much warmer]. The labels correspond to values between -3 and 3. Compared to the ASHRAE scale, this scale is inverted, meaning that negative values correspond to positive values on the ASHRAE scale and vice versa. We also allowed users to mark whether one of the following things affected their thermal comfort: *cold draft*, *heat from radiators* and *heat from the sun*. To bottom of the page featured a comment box where users could list other things causing them discomfort. Feedback was linked to individual users by assigning each user an anonymous, unique ID that was stored in a cookie in the browser.

Office deployment. In the office deployment, we positioned 29 posters and three humidity sensors in 8 offices and 3 hallway locations around a single floor of a naturally ventilated office building with occupant-controlled windows and normal radiators. Some of the smaller offices further featured portable, manually controlled air conditioning units. Posters were deployed over a two-week period between February and March 2016. Occupant numbers in offices ranged from 2 people (1 office) to 4-10 people (4 offices) to more than 20 (3 large open plan offices). The offices were occupied by university administrative workers, who had been informed by email in advance of the deployment. Following the deployment, we further conducted semi-structured interviews with some of the occupants to find out more about their experience with thermal comfort, the deployment and its user interface.

	OFFICE	LIBRARY
Votes logged	167	990
Overall Users	57	688
Eligible Users	12	23
Processable Users	12	11
Processable Votes	98	61

Table IV: Statistics of votes and users in the deployments

Library deployment. In the library deployment, we positioned 143 posters over all 5 floors of the library. At the time of the deployment, the library was mostly ventilated with forced natural ventilation with some small areas being equipped with air conditioning. Since students were not approached directly with information about the study, the posters and loggers were deployed for 5 weeks between May and June 2016, giving students enough time to discover and familiarise themselves with the study. Posters were positioned to reach a representative geographic coverage for temperature and to be sufficiently visible and accessible to occupants in most parts of the library. Similar to the office deployment, three posters were equipped with humidity loggers. These posters were positioned on the first, third and fifth floor of the library. The library was mostly occupied by students, who due to wishes of the library were not approached directly or informed via email about the study. For the same reasons, no interviews were conducted following the library deployment.

Processing of results. In order to use the data obtained from these deployments for validating our comfort model, we had to process parts of it first. Since the model requires multiple observations per user, we discard all users that did not provide enough feedback to provide meaningful results. We set this threshold to be five observations or more. This value was decided based on evaluations on the ASHRAE data set, in which our model requires around four training observations to outperform the other models. Having five or more observations per user allows to train the model to that point while still leaving at least one additional observation for evaluation. Users who have provided five or more observations are considered to be *eligible users*.

In case of the library deployment, some additional filtering was required as it coincided with both a heat wave in the UK as well as exam periods. The unusually high outdoor temperatures together with high occupancy due to students preparing for their exams lead to very high temperatures (28°C and higher) in some parts of the library. As a result, a lot of users only provided extreme votes of 3 (ASHRAE scale) or -3 (our scale). Such situations should not occur in a system controlled using our heating algorithm and comfort model, since the model would suggest lower set-points after the first extreme vote. We therefore only considered users where extreme votes would make up 50% or less of the feedback. Some users further provided feedback too frequently, with votes just being minutes apart from each other, potentially leading the algorithm to overfit to these conditions. We therefore average all votes that are less than 15 minutes apart from each other. Users with five or more remaining votes after this reduction are considered to be *processable users*. Table IV gives an overview over how many users participated in each deployment, how many votes were logged and how much of the data was processable in the end.

5.2. Model Accuracy

We benchmark our model's prediction accuracy against the existing, standardised ASHRAE comfort models described in Section 2: 1. Static comfort model (PMV) by [Fanger 1970] 2. the adaptive model by [de Dear and Brager 1998] The approaches

DATA SET	PMV		ADAPTIVE	
Pakistan	27%	(0.39)	5%	(0.055)
Athens	31.2%	(0.363)	25.9%	(0.28)
Bay Area (s)	23.8%	(0.259)	29.2%	(0.342)
Bay Area (w)	18%	(0.216)	7.2%	(0.076)
Southampton (ASHRAE scale)	45.8%	(0.931)	27%	(0.406)
Southampton (our scale)	55.4%	(0.87)	43.1%	(0.529)
Overall	25.8%	(0.314)	13.2%	(0.15)

Table V: Accuracy gains (difference of the RMSE in parentheses) for the predicted T_{vote} of our model vs. PMV and adaptive models after up to 20 observations

were compared based on the RMSE of their predictions for T_{vote} . Using the RMSE allows to draw conclusions about both, absolute performance by looking at the absolute values as well as relative performance by comparing the RMSE of different approaches. While the PMV (similar to T_{vote}) for the static model was provided with the data sets, for our own deployment we manually computed it using the equations provided in the ASHRAE standard. We chose values of 1.1 for the metabolic rate, corresponding to the activity of “typing”. For the office deployment, the clothing level was set to 1.0 (typical winter indoor), for the library deployment to 0.5 (typical summer indoor). We use the model’s default value of 0.1 m/s for the air velocity, as this was not measured in our deployment. As the adaptive model only outputs an optimal comfort temperature T_{opt} instead of a vote T_{vote} , we need to calculate the vote. To do so, we multiply the difference between the operative temperature T_{op} and optimal comfort temperature T_{opt} by a thermal sensitivity (see equation 9).

$$T_{\text{vote}} = 0.29(T_{\text{op}} - T_{\text{opt}}) \quad (9)$$

We chose a value of 0.29 for the thermal sensitivity, which corresponds with values learned by our model as well as values found by [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 evaluation run, different amounts of training observations, increased in steps of 1, were tested. For data sets with many data points per individual (Pakistan and Athens), up to 20 training observations were used. For the Bay Area and Southampton, the amount was increased up to the maximum possible observation count of a user (number of observations - 1).

The evaluation for a single data point consisted of two steps. First, the model was trained using the training observations. Following this, feedback for different evaluation points was inferred and the squared error between the prediction and actual feedback was logged. From all single results, the RMSE and standard error σ were calculated, which will be discussed in the next section.

5.2.1. Empirical results. Table V shows accuracy gains for predictions of T_{vote} achieved by our model compared to the PMV and adaptive model. The poor performance when having less than three observations can be explained by inaccurate priors that were estimated manually by us after the first initial simulations (see Table VI for the values chosen for the priors). One can see that apart from the Pakistan data set (see Figure 4a), our model achieves significant accuracy gains (7-55% smaller prediction error for T_{vote}) in comparison to the other approaches (Figures 4b and 4c). We identify two main reasons for the accuracy gains: using a more complete model and adapting to individual users’ preferences. As opposed to the other approaches, the BCM includes

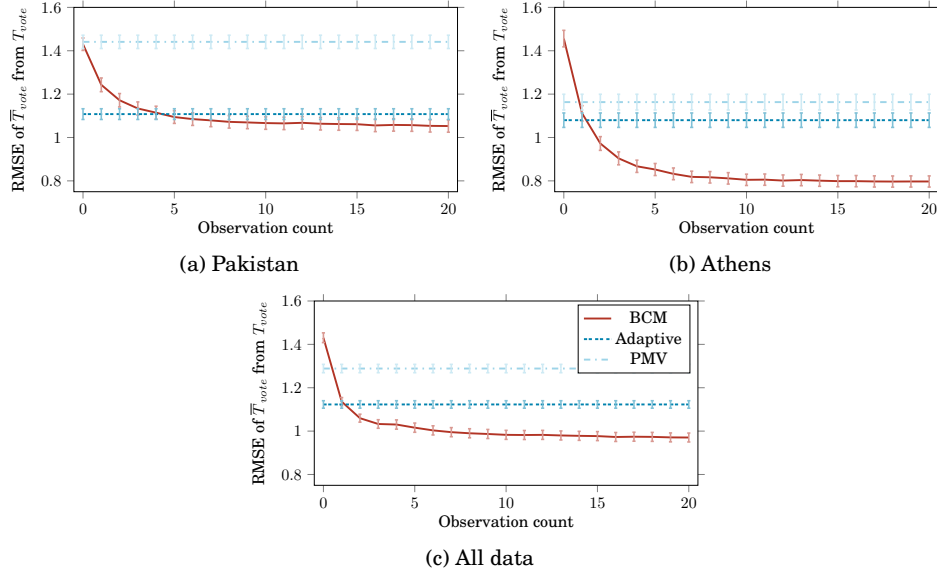



Fig. 4: RMSE of predicted vote vs. observation count (with 2σ confidence interval)

both static as well as adaptive components at the same time, yielding an overall more accurate representation of thermal comfort. However, even with the more complete model, individual users' preferences may still vary. By learning from user feedback, the BCM compensates for those varying preferences. The lower accuracy gains on the Pakistan data set can be explained by the fact that this data set contains many spurious 0 votes regardless of the thermal environment, 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 Bay Area 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 V). In general, our model typically converges after 10 observations (see Figure 4c), but starts outperforming the other models after 2 - 4 observations. We tried to incorporate measurements for draft and metabolic rates provided in the ASHRAE data sets into our model. Including these parameters however did not yield any notable improvements.

Table VI 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 have a greater impact on users' comfort levels. One can see that users' thermal sensitivities, γ_v , vary a lot. As we will show in section 5.3, our heating algorithm's ability to reduce energy consumption is highly dependent on the thermal sensitivity (as it determines the width of the comfort range), stressing the importance of learning this particular parameter.

Comparison of different temperature scales. In addition to the general evaluation of the prediction accuracy of our model, we also evaluate the effect different feedback scales have on the prediction's accuracy. Figure 5 shows the model's prediction accuracy using the standard ASHRAE 7-point scale (Figure 5a) as well as using the re-

PARAMETER	VALUE RANGE	μ	σ	PRIOR
T^*	[19.86, 25]	22.02	0.99	21
γ_v	[0.006, 0.96]	0.29	0.23	0.3
γ_h	[2.6, 3.37]	2.9	0.135	3
γ_a	[0.034, 0.91]	0.61	0.25	0.5
γ_{a_s}	[0.93, 1.32]	1.04	0.046	1
γ_{a_b}	[-0.43, 0.063]	-0.29	0.116	-0.3

Table VI: Learned parameter statistics (μ = average, σ = standard deviation)Fig. 5: Comparison of different feedback scales (with 2σ confidence interval)

requested change scale described in section 5.1.2 (Figure 5b). Despite the low number of participants and resulting large error, one can see that using our requested change scale seems to lead to much lower prediction errors in all comfort models. Our scale leads to reductions in prediction error of up to 22.7% lower for the PMV, 18.4% lower for the adaptive model and 36.4% lower for our comfort model.

A possible explanation for this is that, as explained in section 5.1.2, the ASHRAE scale does not correct for general users' preferences of finding slightly warm or cold temperatures preferable to a neutral environment. On average, the differences of users' votes between the scales was 0.54, suggesting that users often don't necessarily aim for a "neutral" environment.

5.3. Energy Savings

To assess the theoretical energy savings of our algorithm when using the BCM, we simulate households and their HVAC systems. We evaluate two main metrics: energy consumption and user discomfort. Energy consumption is calculated as the aggregate of the products of heating and cooling ratios with the maximum energy consumption of heater and AC (similar to equation 6). We chose values of $\xi^r = 8\text{kW}$ and $\xi^c = 12.5\text{kW}$ for the maximum energy consumption of the heater and AC which correspond to real values observed in the Pecan Street data set, a data set that containing disaggregated household energy data. Running times are calculated by multiplying the heating and cooling ratios (ρ_r^t and ρ_c^t) at time t with the length of a time step Δt . User discomfort is defined as the aggregated discomfort of each user over each time step of the simulation. For a time interval t , user discomfort is calculated as the product of the deviation of the indoor temperature from the user's comfort range and the length of the time interval Δt . We simulate 1000 households with either 2 or 3 occupants per household with an average of 2.4 occupants per household which is the average number of occupants in a household in OECD countries. Users are simulated using random configurations of our comfort model where parameters are drawn from the distributions shown in Table VI. To cover a variety of climates and seasons, households were simulated over the time




Fig. 6: Energy saving potential and discomfort

period between 1st of January 2014 to the 31st of December 2015 in the following cities: Austin, TX, USA; Brussels, Belgium; Cape Town, South Africa; Moscow, Russia; San Francisco, CA, USA; Seattle, WA, USA; Shanghai, China; Sydney, Australia.

We define two main benchmarks for our model: a fixed set-point heating system (FSP) and a fixed comfort range (FCR) system. The FSP emulates a typical thermostat that keeps the indoor temperature as close as possible to a user defined set point. As real thermostats usually fluctuate slightly around the set-point temperature, we allow the indoor temperature to deviate by up to $0.2^{\circ}C$ from the set point temperature. We choose the set point to be the average comfort temperature of all occupants throughout the simulations ($23^{\circ}C$). Using these set point temperature and deviation values, the algorithm keeps the indoor temperature between $22.8^{\circ}C$ and $23.2^{\circ}C$. This is realized by defining this as the lower and upper bounds (lb and ub) of the comfort range (see Section 4.2). We use the FCR to demonstrate the benefits of using comfort ranges over a single comfort temperature. Similar to the FSP, the FCR is implemented by setting lb and ub to static values. For the FCR, we allow $0.5^{\circ}C$ deviation from the base temperature, resulting in a temperature range from $22.5^{\circ}C$ to $23.5^{\circ}C$.

5.3.1. Dealing with multiple occupants. In practice, most houses will be occupied by multiple occupants simultaneously. Each of these occupants will have their own configuration of the comfort model, resulting in each occupant having their own personalised comfort range. As the algorithm only optimizes for a single comfort range, we need to merge all occupants' comfort ranges. We implement a simple *comfort compromiser* that aggregates comfort ranges by taking the maximum of all occupants' lower bounds and the minimum of all occupants' upper bounds. There can be cases when not all comfort ranges overlap, leading to the maximum lower bound being larger than the minimum upper bound. In that case there is no way to guarantee a comfortable environment for every user at all times. We resolve such conflicts by using the next highest lower bound or next lowest upper bound. This is repeated until an overlap is found.

5.3.2. Results. Figure 6a shows the resulting relative running times of heater and air conditioning units depending on the comfort model used. One can see that the FSP and FCR require more energy for both, heating and cooling. In the case of heating, the FCR uses about 6.4% more energy than the BCM, the FSP about 10.5%. For cooling, the FCR increases energy consumption by about 15.1%, the FSP by 39.4%. Cooling times were about 10 times lower than heating times, meaning that the great difference for cooling only has a minor effect on the overall energy consumption which was increased by 7.3% with the FCR and by 13.6% with the FSP. The FSP used about 5.8% more energy than the FCR, showing the positive effect that larger temperature ranges have on the algorithm's energy saving capabilities. Compared to a real thermostat, we expect the energy savings to be slightly higher as our algorithm is utilizing even the narrow comfort range in the FSP to reduce energy consumption. Figure 6b shows the

average discomfort experienced per user per day. The BCM leaves about $0.26^{\circ}Ch$ per day, meaning that on average, users experience the equivalent of an hour of $0.26^{\circ}C$ from their comfort range per day. One can see the BCM decreases discomfort by about 24.8%. The remaining discomfort in our model mainly results from non-overlapping comfort ranges of occupants in the same building which cannot be avoided.

6. CONCLUSIONS

In this work we presented the Bayesian Comfort Model (BCM), an intelligent HVAC control algorithm using the BCM and proposed the desired change scale to assess user comfort as an alternative to the ASHRAE 7-point scale. The BCM existing comfort models, simplifies them to only require easily obtainable input parameters and utilises a Bayesian network to learn individual users' preferences. The HVAC control algorithm utilizes linear programming to find optimal HVAC schedules that maximise user comfort while minimizing energy consumption.

Through empirical evaluations using standard data sets and data collected in a new deployment we showed that our model outperforms existing approaches by 13.2% - 25.8% after a short initial learning phase. We further showed how the model's performance also benefits from using the desired change scale instead of the standard ASHRAE 7-point scale. By using the ability to infer additional user parameters such as thermal sensitivity, our heating algorithm was able to reduce energy consumption for heating by 6.4% to 10.5%, energy consumption for air conditioning by 15.1% to 39.4% while at the same time reducing user discomfort by 24.8%.

ACKNOWLEDGMENTS

The authors would like to thank the Carbon Management Fund of the University of Southampton for providing funding for the deployments conducted as part of this paper.

REFERENCES

- ASHRAE 55 2010. Thermal Environmental Conditions for Human Occupancy (ANSI Approved). (2010).
- Auliciems. 1981. Towards a psycho-physiological model of thermal perception. *International Journal of Biometeorology* 25, 2 (1981), 109–122.
- Eric Bauer, Daphne Koller, and Y Singer. 1997. Update rules for parameter estimation in Bayesian networks. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*. Providence, Rhode Island, 3–13.
- Heather Chappells and Elizabeth Shove. 2005. Debating the future of comfort: environmental sustainability, energy consumption and the indoor environment. *Building Research & Information* 33, 1 (2005), 32–40.
- Richard de Dear and Gail Brager. 2002. Thermal comfort in naturally ventilated buildings: revisions to ASHRAE Standard 55. *Energy and Buildings* 34, 6 (July 2002), 549–561.
- Richard de Dear and Gail Schiller Brager. 1998. Developing an adaptive model of thermal comfort and preference. In *ASHRAE Transactions*, Vol. 104. ASHRAE, 145–167.
- Varick Erickson and Alberto Cerpa. 2012. Thermovote : Participatory Sensing for Efficient Building HVAC Conditioning. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. 9–16.
- Povl Ole Fanger. 1970. *Thermal comfort. Analysis and applications in environmental engineering*. R.E. Krieger Pub. Co. 244 pp. pages.
- Pierre Gadonneix, Younghoon David Kim, Kevin Meyers, Graham Ward, and Christoph Frei. 2013. *World Energy Resources 2013*. Technical Report. World Energy Council.
- Tom Hargreaves, Michael Nye, and Jacquelin Burgess. 2013. Keeping energy visible? Exploring how householders interact with feedback from smart energy monitors in the longer term. *Energy Policy* 52 (2013), 126–134.
- Hyung Chul Jo, Sangwon Kim, and Sung Kwan Joo. 2013. Smart heating and air conditioning scheduling method incorporating customer convenience for home energy management system. *IEEE Transactions on Consumer Electronics* 59, 2 (2013), 316–322.

- Abraham Hang-yat Lam, Yi Yuan, and Dan Wang. 2014. An occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings. In *Proceedings of the 5th international conference on Future energy systems - e-Energy '14*. 133–143.
- Jing Liu, Runming Yao, and Rachel McCloy. 2012. A method to weight three categories of adaptive thermal comfort. *Energy and Buildings* 47, 0 (2012), 312–320.
- Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. 2010. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. *Proceedings of ACM SenSys* 55 (2010), 211–224.
- TP Minka. 2001. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*. 362–369.
- T. Minka, J.M. Winn, J.P. Guiver, S. Webster, Y. Zaykov, B. Yangel, A. Spengler, and J. Bronskill. 2014. Infer.NET 2.6. (2014). Microsoft Research Cambridge. <http://research.microsoft.com/infonet>.
- F Nicol, G N Jamy, O Sykes, M Humphreys, S Roaf, and M Hancock. 1994. A survey of thermal comfort in Pakistan. *July, Final report, Oxford Brookes University: School of Architecture, Oxford* (1994).
- Judea Pearl. 1986. Fusion, propagation, and structuring in belief networks. *Artificial intelligence* 29, 3 (1986), 241–288.
- Leen Peeters, Richard De Dear, Jan Hensen, and William D'haeseleer. 2009. Thermal comfort in residential buildings: Comfort values and scales for building energy simulation. *Applied Energy* 86, 5 (2009), 772–780.
- Sean Purdon, Branislav Kusy, Raja Jurdak, and Geoffrey Challen. 2013. Model-free HVAC control using occupant feedback. In *Proceedings - Conference on Local Computer Networks, LCN*. 84–92.
- Alex Rogers, Siddhartha Ghosh, Reuben Wilcock, and Nicholas R. Jennings. 2013. A Scalable Low-Cost Solution to Provide Personalised Home Heating Advice to Households. *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings - BuildSys'13* (2013), 1–8.
- Alex Rogers, Sasan Maleki, Siddhartha Ghosh, and Nicholas R Jennings. 2011. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. In *Second International Workshop on Agent Technology for Energy Systems (ATES 2011)*. 71–78.
- Mike Shann and Sven Seuken. 2013. An Active Learning Approach to Home Heating in the Smart Grid. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*. AAAI Press, 2892–2899.
- Mike Shann and Sven Seuken. 2014. 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*. 821–828.
- S Shindu Shetty, Hoang Duc Chinh, and S K Panda. 2015. Strategies for Thermal Comfort Improvement and Energy Savings in existing Office Buildings using Occupant Feedback. In *2015 IEEE International Conference on Building Efficiency and Sustainable Technologies*. 23–27.
- Stephen Snow, Laurie Buys, Paul Roe, and Margot Brereton. 2013. Curiosity to Cupboard: Self Reported Disengagement with Energy Use Feedback over Time. In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration*. 245–254.
- Fan Tang, Andrew Kusiak, and Xiupeng Wei. 2014. Modeling and short-term prediction of HVAC system with a clustering algorithm. *Energy & Buildings* 82 (2014), 310–321.
- Xiupeng Wei, Andrew Kusiak, Mingyang Li, Fan Tang, and Yaohui Zeng. 2015. Multi-objective optimization of the HVAC (heating, ventilation, and air conditioning) system performance. *Energy* 83 (2015), 294–306.
- Rayoung Yang and Mark W Newman. 2012. Living with an Intelligent Thermostat : Advanced Control for Heating and Cooling Systems. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)* (2012), 1102–1107.
- R. Yang and M. W. Newman. 2013. 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)* (2013), 93–102.
- Rayoung Yang, Mark W. Newman, and Jodi Forlizzi. 2014. Making sustainability sustainable: Challenges in the Design of Eco-Interaction Technologies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. 823–832.