A Heating Agent using a Personalised Thermal Comfort Model to Save Energy

(Extended Abstract)

Frederik Auffenberg, Sebastian Stein and Alex Rogers
Electronics and Computer Science
University of Southampton, UK
{fa1c12,ss2,acr}@ecs.soton.ac.uk

ABSTRACT

We present a novel, *personalised* thermal comfort model and a heating agent using this model to reduce energy consumption with minimal comfort loss. At present, heating agents typically use simple models of user comfort when deciding on a set point temperature for the heating or cooling system. These models however generally fail to adapt to an individual user's preferences, resulting in poor performance. To address this issue, we propose a *personalised* thermal comfort model using 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 17.5-23.5% more accurate than current models, regardless of environmental conditions and type of heating system. Further, our model has several additional outputs such as expected user feedback, optimal comfort temperature and thermal sensitivity that allow it to save between 18-20% of energy while still maintaining comfort.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence*

Keywords

Smart Heating; Thermal Comfort; Bayesian Networks; Agent-Based Control; Human-Agent Collectives

1. INTRODUCTION

Domestic heating, accounting for 12% of the worldwide energy consumption, 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, typically by applying intelligent heating schedules. The key component of such a system is a *smart thermostat.* Smart thermostats often act as an agent controlling the heating on behalf of the user. To satisfy the user's needs, such agents can utilise thermal comfort models to decide the set point temperature or learn their heating schedules. However, existing thermostats such as the Nest learning thermostat often fail to accurately predict the temperature at which the user will feel most comfortable [3]. Due to these inaccurate predictions and the consequential manual adjustments by the user, those thermostats often fail to save significant amounts of energy [3].

To address these shortcomings, we present and evaluate a comfort model that combines existing models into a single, more general model that only requires easily obtainable inputs. Using a Bayesian network, we add learning capabilities that allow the model to learn individual user's preferences from minimal feedback. We further extend our model to allow inference of various parameters such as the current expected user feedback for any temperature, the user's current thermal sensitivity as well as the user's current optimal comfort temperature. We then build a heating agent that utilises these outputs to reduce energy consumption of the heating system with minimal comfort loss for the user.

2. PERSONALISED THERMAL COMFORT

Our personalised thermal comfort model, shown in Figure 1, is a combination of Fanger's static model and the standard adaptive model [1]. These models have been modified to work with easily obtainable inputs and extended with user specific scaling factors. Further, we add a new component accounting for seasonal adaptations. The main 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.

We benchmark our model against Fanger's static model and the adaptive model, both taken from [1]. We use real world data from longitudinal studies of the ASHRAE RP-884 data set. The data contains information for 553 different individuals in 10 cities in 3 different locations (4 cities in Pakistan, the city of Athens and 6 cities in the San Francisco Bay area), covering domestic and office spaces with both naturally ventilated (NV) and heating, ventilation and air conditioning (HVAC) systems. For single individuals, the number of subsequent observations varies between 3 and 150 values. In general, our model converges after 10 observations and gives 17.5-23.5% more accurate predictions of the expected vote, T_{vote}, of a user.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

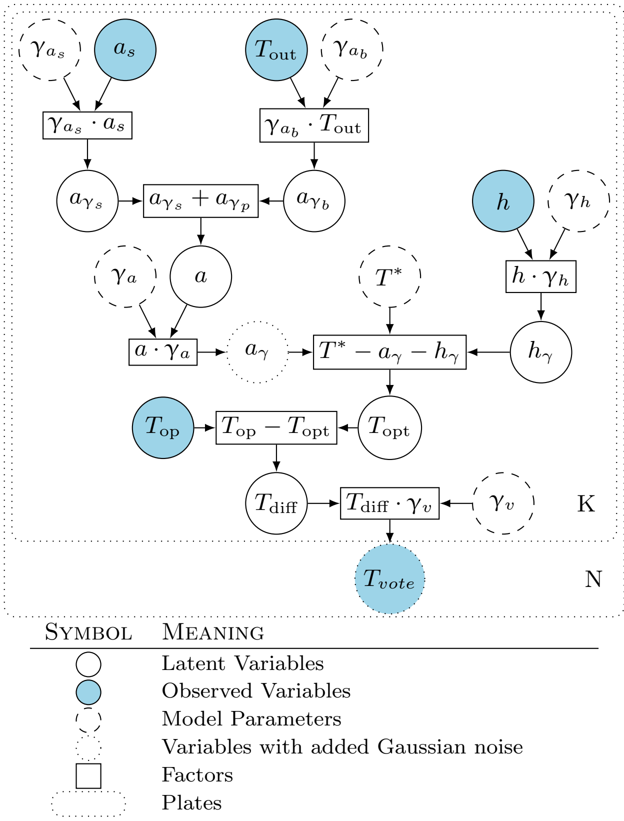

γ_{var}	User-specific scaling for variable var
var_γ	User-specific scaled value of var
a_s	Seasonal adaptations
T_{out}	Outside temperature
h	Relative humidity inside
a	Sum of all adaptations
T^*	Comfort temperature in neutral conditions
T_{op}	Operative temperature
T_{opt}	User's current optimal temperature
T_{diff}	Deviation of T_{op} from T_{opt}
γ_v	Thermal sensitivity
T_{vote}	User's current vote

Figure 1: The personalised thermal comfort model as a Bayesian network (with nomenclature)

3. THE HEATING AGENT

To evaluate the potential energy savings a personalised thermal comfort offers, we implement a smart heating agent running in simulated households using simulated users based on the Pakistan and Athens data sets. Due to the lack of observations per individual, the San Francisco data set was not used in the simulations.

Users are simulated using an instance of our comfort model that has been trained with a randomly chosen user from the data sets. Feedback was provided using the ASHRAE 7-point scale. To simulate manual adjustments, users constantly assess whether their T_{vote} falls within their acceptance range $[-0.75, 0.75]$. After at max an hour of discomfort where T_{vote} is outside $[-0.75, 0.75]$, the user reports back to the system and manually adjusts the set point to the current T_{opt}. With higher discomfort, the shorter the time it takes the user to complain.

The heating agent regularly reassesses the user's current comfort temperature range ($-0.4 \leq T_{\mathrm{vote}} \leq 0.4$) based on our model and adjusts the set point accordingly. Instead of a simple set point, the agent defines an acceptable temperature range for the heating system. The heating system only turns the AC or heating on when the inside temperature is about to leave this comfort range. Further, between 1$^{\text{st}}$ of June and the 30$^{\text{th}}$ the heating was generally disabled. For the rest of the time, the AC was disabled.

To model the thermal properties of the buildings, we used a thermal model which includes a leakage rate, Φ, (1/hr), and heater output, R, (°C/hr) [2]. To include air conditioning, a cooling rate, C, (°C/hr) is also used. Together, temperature changes of the air temperature T_{air} between times t and $t+1$ within the building is given by:

$$T_{\mathrm{air}}^{t+1} = T_{\mathrm{air}}^{t} + \left[R\, h_{\mathrm{on}}^{t} - C\, \mathrm{ac}_{\mathrm{on}}^{t} - \Phi(T_{\mathrm{air}}^{t} - T_{\mathrm{out}}^{t})\right] \Delta t \qquad (1)$$

The variables h_{on}^{t} and $\mathrm{ac}_{\mathrm{on}}^{t}$ describe whether the heater (h) or air conditioning (ac) are turned on at time t.

In our simulations, we compare how often and for how long users are dissatisfied, and how long the heater and AC are running over the course of one year, for heating systems using our agent and manually controlled heating systems. Overall, the comfort model reduced heating times by about 18.5% (123 hours) and AC times by about 20% (63 hours). In a real application we expect these values to be a bit lower as our model of manual adjustments by the user might be too simplistic at the current stage. Without the comfort model, users on average adjusted the temperature on their thermostat 10 times per year, with 12.5 times a year. This value includes values from the initial setup. Due to the fairly quick adjustments by the user, dissatisfaction times were similar with and without the comfort model.

4. CONCLUSION AND FUTURE WORK

In this paper, we introduced a novel, personalised thermal comfort model and show through an empirical evaluation that it is on average 17.5-23.5% more accurate than established models and provides additional outputs, enabling us to create a heating agent to reduce energy consumption by 18.5-20%. In future work, we will address satisfying multiple occupants simultaneously. Another extension could be to include an individual cost- vs. comfort payoff in the model.

REFERENCES

[1] Thermal Environmental Conditions for Human Occupancy (ANSI Approved), 2010.

[2] A. Rogers, S. Ghosh, R. Wilcock, and N. R. Jennings. 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*, pages 1–8, 2013.

[3] 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.