UNIVERSITY OF SOUTHAMPTON

A Comfort-Based, Energy-Aware HVAC Agent and its Applications in the Smart Grid

by

Frederik Auffenberg

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the
Faculty of Physical Sciences and Engineering
Electronics and Computer Science

March 2017

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Frederik Auffenberg

In this thesis, we introduce a novel heating, ventilation and air conditioning (HVAC) agent that maintains a comfortable thermal environmant for its users while minimising energy consumption of the HVAC system and incorporating demand side management (DSM) signals to shift HVAC loads towards achieving more desirable overall load profiles. To do so, the agent needs to be able to accurately predict user comfort, for example by using a termal comfort model. Existing thermal comfort models are usually built using broad population statistics, meaning that they fail 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 using a Bayesian network learns 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. We extend this algorithm to incorporate DSM signals into its scheduling, allowing it to shift HVAC loads towards more desirable load profiles, reduce peaks or make better use of energy produced from renewable sources. 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 comfort model is consistently 13.2% to 25.8% more accurate than current models and that the alternative comfort scale can increase our model's accuracy. Through simulations we show that when using the comfort model instead of a fixed set point, our HVAC control algorithm can reduce energy consumption of the HVAC system by 11% while decreasing user discomfort by 17.5%, achieve a load profile 39.9% closer to a specified target profile and efficiently reduce peaks in the load profile.

Contents

List of Figures

List of Tables

Glossary

AC	Air conditioning
AMPC	Average model parameter compromiser
ANSI	American National Standards Institue
AVD	Actual vapour density
BCM	Bayesian Comfort Model
BN	Belief network / Bayesian network
CO_2	Carbon Dioxide
DLC	Direct load control
DR	Demand response
DSM	Demand side management
EM	Expectation maximisation
EP	Expectation propagation
FCR	Fixed comfort range
FSP	Fixed set point
GP	Gaussian Process
HEMS	Home energy management system
HVAC	Heating, ventilation and air conditioning
KL divergence	Kullback-Leibler divergence
MDP	Markov decision process
MIQP	Mixed-integer quadratic program
MRT	Mean radiant temperature
NV	Naturally ventilated
OCC	Overlap comfort compromiser
PMV	Predicted Mean Vote
power-EP	power Expectation propagation
PPD	Predicted Percentage Dissatisfied
RMSE	Root mean square error
SVD	Saturation vapour density

Nomenclature

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).

Declaration of Authorship

I, Frederik Auffenberg declare that the thesis entitled A Comfort-Based, Energy-Aware HVAC Agent and its Applications in the Smart Grid and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Either none of this work has been published before submission, or parts of this work have been published as:
    - Auffenberg et al. (2015b)
    - Auffenberg et al. (2015a)
    - Auffenberg et al. (2017, in press)
    - Snow et al. (2017)

Signed: ....

Date: ....

Acknowledgements

I would like to thank my supervisors Sebastian Stein and Alex Rogers for their support and encouragement throughout my PhD. Their continuous guidance and feedback greatly helped me to grow as a researcher. I could not have imagined better support throughout my PhD. I would also like to thank my lab mates for all their help and advice. In particular, I'd like to thank Matteo Venanzi, James Holyhead and Moody Alam for always being available to discuss different approaches and algorithms.

Thanks also to all the people who managed to make Southampton a somewhat enjoyable place to live in. Thanks to Tim Baarslag and Valerio Restocchi for their regular help with getting rid of excess knowledge over a beer or two after work. I would also like to express my gratitude towards Elliot Salisbury and James Holyhead for the frequent tea breaks, regular political discourse and their attempts to teach me about British culture. I was recently complimented on my ability to talk about the weather which would have not been possible without their support. I would also like to thank all my lab mates who — through regular shenanigans and tea-breaks — made the lab a social and enjoyable place to work in.

To my family: thank you for your unconditional support throughout my thesis and life so far.

Lastly, I would like to thank the Engineering and Physical Sciences Research Council (EPSRC) for providing funding for this PhD and PG Tips breakfast tea for not just fueling my research, but also all the procrastination tea-breaks.

Chapter 1

Introduction

One of the biggest challenges of the 21st century is to fulfil the growing demand for energy while at the same time reducing global emissions to mitigate the adverse effects of global warming and the dwindling supply of fossil fuels. Numerous governments have agreed on common targets as defined in the Kyoto Protocol or specified by the International Energy Agency. As part of the Climate Change Act (Department of Energy and Climate Change, 2008) the UK government for example agreed to reduce emissions in the UK by at least 80% by 2050 compared to the levels of 1990. In June 2017, the UK achieved a 42% reduction of CO_2 emissions as compared to levels of 1990, with the most recent reductions being achieved by a reduction of energy generated from coal (Commitee on Climate Change, 2017). With current efforts continued, the Commitee on Climate Change (2017) predicts that the UK will not meet its 2030 and 2050 targets and will only achieve half the reduction required to meet the goal for 2030. To be able to meet the 80% goal, electricity generation has to continue to move away from fossil fuels towards renewable energy sources such as wind energy or solar energy. Furthermore, other sectors heavily relying on fossil fuels such as heating and transport need to be electrified (Department of Energy and Climate Change, 2013). According to the UK Energy Consumption Report (Department of Energy and Climate Change, 2016), heating and transport account for the largest part of the overall end-use energy consumption in the UK. In 2015, both together accounted for about 82.7% of the overall end-use (domestic and industrial combined) energy consumption in the UK (heating: 42.8%, transport: 39.9%). Consequently, the electrification of these sectors will introduce a high rise in demand for electric energy in the future. Across industrialised countries heating alone accounts for 12% of the overall energy consumption (Gadonneix et al., 2013).

Moving towards renewable energy sources and the electrification of heating and transport will require significant changes to the electricity grid. The output of most renewable energy sources depends on the weather, making it hard to plan and introducing a lot of variation in output over time and location. This leads to peaks and troughs in production. During times of peak production, large amounts of energy need to be distributed through the grid and excess energy needs to be stored for times of low energy production. This variability poses a significant challenge for ageing electricity distribution networks (Strbac, 2008). In addition to that, the electrification of heating and transport will lead to a significant increase in electricity demand, especially during peak hours in the evening. Since the grid's capacity needs to be at least as big as the peak demand to guarantee a stable electricity supply (Strbac, 2008), the increased peak demand poses the risk of exceeding the grid's current capacity, requiring costly investments for increasing capacity. An alternative to these investments is to make better use of the existing infrastructure, for example by spreading out peaks into troughs or by controlling demand on the network to follow variations in energy production. The *smart grid* addresses this challenge by adding an additional communication layer to the network, enabling real-time communication between consumers and providers. This communication layer can be used for DSM to reschedule consumer demand to match instantaneous supply and reduce peak consumption. Transport, heating and existing loads from air conditioning (AC) offer great potential for DSM (Strbac, 2008; Ramchurn et al., 2011) due to their flexible nature. So far, research has focused on how to intelligently plan the charging of electric vehicles (Valogianni et al., 2015) or how to utilise their batteries as distributed energy storage (Galus and Andersson, 2008). In the field of heating and air conditioning, DSM technologies such as smart thermostats have only been utilised recently (Wernstedt et al., 2007; Erickson and Cerpa, 2010; Lu, 2012; Ramchurn et al., 2011) and there is still a lot of research to be done on how to incorporate typical DSM measures such as variable price rates for energy into heating and air conditioning.

Apart from its potential utilisation in DSM, smart heating and air conditioning also has the potential to reduce energy consumption as well as emissions in general. So far, most efforts to reduce energy consumption are focusing on improving the insulation of buildings or switching to more efficient or cleaner sources of heat, such as ground source heat pumps. While this approach has been proven to be very efficient, it usually is very expensive to improve insulation for a whole house or to change the whole heating and air conditioning system. A more accessible approach is to optimise heating and air conditioning schedules. This approach can be used regardless of the underlying heating or cooling system and the building's insulation. In non-domestic buildings where HVAC accounts for around 48% of the overall energy consumption (Pérez-Lombard et al., 2008), this is done to some extent already. In such buildings, the HVAC system is usually fully automated, keeping the temperature within narrow bounds and potentially switching the HVAC off outside of working hours. The domestic sector, which makes up about 30% of the total energy consumption in the UK (Department of Energy and Climate Change, 2016) has been ignored in this field for a long time. Accounting for 42% - 68% of domestic energy consumption (Pérez-Lombard et al., 2008; Palmer and Cooper, 2013), domestic heating and air conditioning offers great potential to reduce overall energy consumption.

One major limitation when trying to reduce energy consumption of heating and air conditioning is user comfort. People usually only have a very limited range of temperatures that they consider to be comfortable (Peeters et al., 2009). The HVAC system should constrain its set-point to lie within this range of temperatures. Determining this temperature range however is complicated as it depends on a multitude of factors such as activity levels and clothing of occupants or the thermal environment (ASHRAE 55). Based on that information, rich models to describe occupants' *thermal comfort* have been developed in the past, of which some have been standardised in ASHRAE 55 and approved by the American National Standards Institue (ANSI). Using these models, given the thermal environment, a set-point temperature satisfying the majority of a group of people can be calculated. Using broad population statistics and building just a single, static model however can be problematic. Thermal comfort tends to be a very individual metric with preferences varying strongly between different individuals (Chappells and Shove, 2005; Clear et al., 2013; Daum et al., 2011; Peeters et al., 2009). Research therefore has to move away from 'one size fits all' comfort models towards more flexible models taking individual preferences into account.

To be able to optimise energy consumption as well as user comfort, these objectives need to be quantified. While energy savings can easily be estimated and quantified (for example in kWh), thermal comfort is harder to assess. The most common quantifications for thermal comfort are the predicted mean vote (PMV) and predicted percentage dissatisfied (PPD) proposed by Fanger (1970), which have since been further improved and standardised as a measure for thermal comfort in ASHRAE 55. The main principle behind the PMV and PPD is to simulate how a group of people would rate the thermal environment using the ASHRAE 7-point scale. The 7-point scale asks users how they would rate their current thermal environment. Values range from -3 (cold) to 3 (warm) with 0 (neutral) being considered most comfortable. While this scale might provide valid ratings about how users perceive their thermal environment, it is not clear whether these values accurately represent users' preferences. For example, some users might have a general preference for warmer environments. For such users, a value of 1 (slightly warm) might represent the most comfortable environment.

In addition to maximising user comfort, smart heating systems offer significant promise in optimising grid load. As mentioned earlier, electrification of heating and existing electric AC offers great potential for DSM. Heating and cooling are generally slow and very predictable. While modern HVAC systems can heat up/cool down a place fairly quickly, even in houses with poor insulation it will take several hours for the temperature to get back to its starting point. This property can be utilised to shift and spread out the heating process in response to predicted or current demand. For example to shift the load forwards by an hour, the house can be heated up to a slightly higher temperature an hour earlier. Several approaches try to utilise HVAC systems for DSM. Some focus on using occupancy prediction to reduce the energy consumption of the HVAC system and reduce peaks (Erickson and Cerpa, 2010). Other approaches use direct load control, where a central controller controls the HVAC system in different houses (Lu, 2012). Wernstedt et al. (2007) model distributed heating systems as a multi-agent system and give an outlook on how DSM goals can be achieved with self-interested agents (Wernstedt et al., 2007). Ramchurn et al. (2011) build upon this idea, extending it with an adaptive stabilising mechanism preventing agents from causing new peaks through reactive crowd behaviour.

However, these approaches typically use oversimplified thermal comfort models which are static (Hubert and Grijalva, 2011; Lu, 2012; Ramchurn et al., 2011) or represent general population preferences rather than individual users (Erickson and Cerpa, 2010). Especially in the domestic sector, where the number of occupants of a building is small, individual users' preferences are likely to deviate from general population preferences. Often these approaches also rely on centralised controllers to perform calculations (Lu, 2012) which can become intractable when dealing with thousands of households. In addition, it requires households to send potentially sensitive information, such as heating schedules and comfort profiles revealing home occupancy as well as individual user's preferences to a central server. Preserving users' privacy and guaranteeing cybersecurity, however, is important for a successful implementation of the smart grid (Liu et al., 2012a).

A main factor in achieving aforementioned objectives is the HVAC control hardware at home or in the office. Recently, smart thermostats that interact closely with the occupants have received a lot of attention from industry and the research community. Such systems can be categorised into two systems: manually scheduled systems and heating agents. Manually scheduled systems rely on the user to specify to what temperature to keep at a given point in time. Heating agents introduce some degree of automation to the heating and cooling scheduling. Such agents could for example learn users' preferences and automatically decide on set point temperatures and when to switch the HVAC system off.

While the number of smart thermostats on the market is growing rapidly, most of them only implement manually controlled systems (for example the Hive active heating[1], Heat Genius thermostat [2] or different Honeywell thermostats) or offer only a very small degree of automation. The Nest learning thermostat[3] (shown in Figure 1.1) is an example of a smart thermostat supporting a higher degree of automation. It gives users the option to learn their heating schedules from previous temperature adjustments. Based on these learned schedules, the thermostat can then automatically adjusts the set point

[1]Hive Active Heating™ (`http://www.britishgas.co.uk/products-and-services/hive-active-heating.html`)

[2]Genius Thermostat (`https://www.geniushub.co.uk/`)

[3]Nest — The Learning Thermostat (`https://nest.com`)

Figure 1.1: An example of a smart thermostat (the Nest learning thermostat)

temperature over different times of the day and week. However, in practice, this mechanism often fails to provide satisfactory results, either leaving users confused about the set point chosen by the thermostat, creating uncomfortable conditions or by actually increasing energy consumption (Yang and Newman, 2012, 2013).

There is ongoing research on how to address these problems. Some approaches focus on maximising user comfort (Sarkar et al., 2016) while others aim to ensure a comfortable environment while reducing energy consumption simultaneously (Shetty et al., 2015; Zhou et al., 2016; Ghahramani et al., 2014; Purdon et al., 2013). Other approaches further try to optimise energy consumption with regards to the smart grid (Shann and Seuken, 2013, 2014; Alan et al., 2016; Shann et al., 2017; Hubert and Grijalva, 2011; Lu, 2012; Ramchurn et al., 2011; Erickson and Cerpa, 2010). While achieving good results in some areas, most approaches show some flaws in other areas. A number of approaches for example overly simplify their comfort models. In some cases, important thermal comfort indicators such as humidity (Sarkar et al., 2016; Shann and Seuken, 2013, 2014) or outdoor temperature (Sarkar et al., 2016) are neglected. Other approaches require a lot or frequent feedback (Shetty et al., 2015; Purdon et al., 2013), do not account for individual user's preferences (Zhou et al., 2016) or neglect that user preferences can change over time (Ghahramani et al., 2014).

Addressing these shortcomings, the aim of this work is to develop an individual-centric, comfort-based HVAC agent that is able to learn and adapt to individual users' preferences using only a minimum amount of feedback, minimize energy consumption while retaining a comfortable environment for its users and that offers capabilities for shifting HVAC loads for DSM. The main goal of the agent is to keep a comfortable environment for the user at any time. In addition to that, the agent aims to reduce energy consumption and shift HVAC loads as much as possible while staying within the comfort bounds if its users. The purpose of the agent is to create a semi-automated HVAC system that automatically optimises heating and cooling schedules on the user's behalf while retaining a comfortable environment for the user. To do so, the heating agent requires both a learning, personalised thermal comfort model as well as an algorithm to optimise the actual heating and cooling schedules with respect to different objectives.

1.1 Requirements

From the challenges described in the previous section, we can derive requirements for the HVAC agent and its underlying comfort model. A comfort-based HVAC agent should fulfill the following requirements:

1. **Accurate Comfort Model** To be able to automate parts of the HVAC system, the HVAC agent needs a precise way to model and predict individual user's thermal comfort. Deviations between predicted comfort temperatures and actual comfort temperatures need to be minimal. de Dear and Brager (1998) suggest that deviations by up to 2.5°C to either side are acceptable in naturally ventilated buildings, deviations of up to 1.2°C in HVAC-equipped buildings respectively. The agent's predictions should therefore on average deviate by less than 2.5°C with a preferrable deviation of less than 1.2°C.
2. **Comfort Constrained** User comfort needs to be the HVAC agent's highest priority. To maximise acceptance of the HVAC agent by the user and get them to voluntarily hand over some control to the agent, at no point in time should the agent compromise comfort for other goals, unless explicitely asked for by the user. This also means that in an environment with multiple occupants, the agent needs to be able to find compromises for different user's preferences and decide a set-point temperature that satisfies as many occupants as possible.
3. **Energy Aware** Next to user comfort, the agent should also aim to reduce energy consumption of the HVAC system. These reductions however can not go at the cost of user comfort as this would violate requirement 2.
4. **Unobtrusiveness** To further increase acceptance by the users, the HVAC agent needs to be as unobtrusive as possible. This means that it should be able to operate autonomously most of the time. The agent should operate on a minimum amount of feedback from the user and keep the information burden (such as varying energy price rates throughout the day) low for the user. However, to keep the system transparent to the user, user feedback should always override the set-point decided by the agent. Further, the agent should only require a minimum amount of unobtrusive sensors, with only one set of easily hideable sensors (such as temperature and humidity sensors) per room. Complex sensors that need to be placed openly in the room (such as anemometers) could potentially interfere with a user's daily routine and should therefore be avoided.

5. **Quick Adaptation** As user engagement can be expected to diminish over time (Hargreaves et al., 2013; Snow et al., 2013), the agent should be able to quickly learn and adapt to individual users' preferences. The initial training phase should take no longer than two weeks and should only require daily feedback at most. After this initial training period, the agent should be able to fulfill the aforementioned accuracy requirement.

6. **Scalability** Typical smart thermostats only feature small computing units (for example an ARM Cortex A8 in the Nest). Since the agent might have to evaluate the user's preferences frequently, this process should be very simple and take no longer than a few seconds on such processors. This requirement also holds for a scenario where calculations are carried out remotely in a datacenter.

7. **Smart Grid** The agent needs to be able to process typical DSM signals and react accordingly, for example by pre-heating or cooling to achieve more desirable load profiles. This means that the agent needs to be able to plan ahead and create HVAC schedules in advance.

8. **Decentralised** Especially with regards to its integration into the smart grid, to scale to a large number of buildings, the agent needs to be able to work as part of a decentralised collective of independent agents. To minimise computational complexity, each agent should be able to optimise its schedules independently from other agents. In aggregate, the sum of each individual agent's efforts should achieve results close to those achievable by a centralised agent controlling all buildings simultaneously.

9. **Stable and Predictable** With regards to the smart grid, the decentralised collective of HVAC agents needs to react in a stable, predictable manner to DSM signals, meaning that effects such as the *rebound effect* (Palensky and Dietrich, 2011) should are prevented.

1.2 Research Challenges

In this section we evaluate a number of existing HVAC agents against our research requirements. We give a brief overview over different approaches and their main focus areas. Based on the evaluation against our research requirements, we derive several research challenges.

The first HVAC agent is the auto scheduler of the Nest thermostat[4]. The auto scheduler does not utilise a thermal comfort model but tries to achieve thermal comfort by learning the schedule of a user's favoured set point temperatures at different times of the week

[4] Nest Auto Schedule (`http://support.nest.com/article/How-does-Auto-Schedule-learn`)

and day. Energy consumption is reduced by only turning on the heating or cooling when necessary. In conjunction with other Nest hardware, the thermostat is further able to measure home occupancy to save energy.

Another HVAC agent is the adaptive home heating agent by Shann and Seuken (2013, 2014). This agent incorporates a two-layer user model. The first, lower layer models the user's optimal comfort temperature. The optimal comfort temperature is defined by a positive linear relationship between comfort temperature and outside temperature following up on research by Peeters et al. (2009). The second layer models the user's personal trade-off function between energy savings and comfort. This layer enables the agent to incorporate variable energy price rates, a typical DSM measure.

There is a number of other HVAC agents that strongly focus on providing capabilities for incorporating smart grid signals. Ramchurn et al. (2011) for example build an adaptive HVAC agent that incorporates dynamic energy pricing signals The agent includes an adaptive mechanism that uses past load profiles to learn how to prevent rebound effects. Lu (2012) propose a centralised HVAC agent that remotely controls HVAC systems to follow load balancing signals. Both of these approaches use a static comfort model of a fixed set point temperature and acceptable deviations from it. Erickson and Cerpa (2010) build an HVAC agent that utilizes occupancy prediction to react to DSM signals and reduce energy consumption.

Another group of approaches strongly focuses on using thermal comfort modelling to reduce energy consumption. Zhou et al. (2016) for example investigate how humidity influences user comfort and build an agent that is able to reduce HVAC energy consumption by manipulating indoor humidity levels. Ghahramani et al. (2014) introduce an energy aware, comfort driven agent that learns users' comfort ranges and uses this information to minimise thermal discomfort. Participatory approaches utilise live feedback from occupants to adjust set point temperatures. Shetty et al. (2015) for example regularly ask occupants about their current comfort levels and based on this feedback identify the main factors influencing the thermal comfort of the occupants. Purdon et al. (2013) introduce an agent that tries to achieve an even balance between feedback stating that it is too hot and feedback stating that it is too cold.

Table 1.1 shows how the existing approaches meet the requirements specified in the previous section. In Section 2.2 of Chapter 2 of this work, the most relevant models are discussed in more detail.

As shown in Table 1.1, none of the existing agents match all requirements (see Section 2.1.2 for further discussion). Most of the existing agents lack a fitting thermal comfort model. Those who have accurate ways to represent user comfort do so by incorporating continuous, frequent user feedback, violating the unobtrusiveness requirement. A possible explanation for the lack of unobtrusive, accurate thermal comfort models is the amount of different factors that need to be taken into account. Usually, the more

Requirement	Nest	Shann and Seuken	Ramchurn et al.	Lu	Erickson and Cerpa	Zhou et al.	Shetty et al.	Ghahramani et al.	Purdon et al.
Comfort-constrained	yes	yes	partially	partially	yes	yes	yes	yes	yes
Energy Aware	partially	mostly	mostly	no	yes	yes	yes	yes	yes
Accurate Comfort Model	no	minmal	no	minimal	partially	minimal	yes	partially	N/A
Unobtrusive	yes	partially	yes	yes	yes	mostly	no	mostly	no
Quick adaptation	mostly	partially	N/A	N/A	N/A	N/A	yes	partially	yes
Scalable	yes	yes	yes	no	yes	yes	N/A	yes	yes
Smart Grid	no	yes	yes	yes	no	no	no	no	no
Decentralised	N/A	yes	yes	no	N/A	N/A	N/A	N/A	N/A
Stable % plannable	N/A	partially	mostly	yes	N/A	N/A	N/A	N/A	N/A

Table 1.1: Existing approaches for heating agents and how well they fulfill the requirements

complex a comfort model gets, the more data points are needed to let it adapt to an individual user's preferences. The need for many data points however conflicts with both, the unobtrusiveness requirement as well as the requirement for quick adaptation. While most HVAC agents fulfill the scalability requirement, the majority does so because of the simplicity of the comfort model used. Most agents are able to fulfill the scalability requirement by either using precomputed thermal comfort models that do not adapt to individual users (Zhou et al., 2016; Ramchurn et al., 2011; Lu, 2012) or stop learning from then early on (Ghahramani et al., 2014), by only relying on live feedback (Erickson and Cerpa, 2010; Shetty et al., 2015; Purdon et al., 2013) or by using simplified thermal comfort models (Nest, Shann and Seuken (2013)). A more accurate, unobtrusive, adaptive model can be expected to be more complex and therefore more costly to evaluate and train. From these points, the first research challenge can be derived:

> The creation of an HVAC agent that correctly, accurately and efficiently learns and predicts users' preferences using only a minimal amount of data that can be easily obtained from user feedback.

The majority of agents are energy aware, meaning that they aim to minimise energy consumption of the HVAC system. While most of them are comfort-constrained at the same time, only Shetty et al. feature an accurate comfort model. In addition, only some approaches leverage properties of their comfort models such as the influence of humidity on user comfort to save energy (Zhou et al., 2016). From that, a second research challenge can be derived:

> The creation of an HVAC agent that leverages properties of user comfort to reduce the HVAC system's energy consumption while retaining a comfortable environment for the occupants.

While some of the HVAC agents provide means to operate in the smart grid, due to simplistic comfort models, none of these agents provide satisfactory results with regards to user's comfort levels. Further, some approaches rely on dynamic pricing rates which can be hard to control or pose the risk of introducing mass reactive behaviour, i.e. rebound effects (Palensky and Dietrich, 2011). This raises a third research challenge:

> The creation of a HVAC agent that incorporates DSM signals into its scheduler to shift the HVAC system's energy consumption to more preferrable times without having to compromise user comfort.

The HVAC agent presented in this work fully addresses all three of these challenges. The HVAC agent present in this work however does not address challenges such as predicting room occupancy in order to switch off the HVAC system while a room is unoccupied or distinguishing between different types of spaces such as kitchens, bathrooms or bedrooms. Further, this work only provides simple means of aggregating preferences of multiple users.

1.3 Research Contribution

In order to improve smart thermostats and to address the challenges stated earlier, this work introduces a personalised thermal comfort model and an HVAC scheduling algorithm that together form a smart HVAC agent. To be able to evaluate the thermal comfort model, we further introduce two new studies examining peoples' thermal comfort levels. The agent presented in this work focuses on predicting user's thermal comfort and using this information optimises the HVAC schedule to minimise energy consumption and incorporate DSM signals. It currently does not incorporate signals such as occupancy, the type of room or specific times of day. Occupancy prediction could be used to strategically turn off the HVAC system when nobody is in the room to further reduce energy consumption. Incorporating the type of room would allow the agent to distinguish between different kindns of rooms and react accordingly. Such types of rooms could for example be rooms that tend to be occupied only for short amounts of time (i.e. kitchen or bathroom), rooms with longer occupancy (i.e. office or living room) and rooms where metabolic rate and insulation levels are better known (i.e. bedroom). When incorporating the time of day as well, this way the agent could for example decide to only adjust the temperature in the bedroom during the night.

The thermal comfort model is based on existing thermal comfort models that have been successfully applied in non-domestic environments in the past. We extend those models to incorporate additional, individual-centric factors. Our model also has been altered to use input parameters that do not require the installation of complicated, intrusive sensors. Using prior work done by Rogers et al. (2012), most thermal parameters of the space can be learned using only the interior temperature and the location of the house. Model inputs that can not be easily measured or estimated and need to be provided by users, such as feedback about the current thermal environments have been chosen to be easy to estimate by the user. We design the model to not only provide information about a single optimal comfort temperature, but to allow inference of multiple parameters that can be used to estimate a range of temperatures that would be considered comfortable by the user.

The model has been implemented as a belief network to allow easy integration of prior knowledge, for example obtained from existing thermal comfort models. Using the expectation propagation message passing algorithm, the network allows both: efficient learning of model parameters based on feedback of the user as well as quick inference of hidden variables. The network is evaluated against data from the free ASHRAE RP-884 data set, a standard thermal comfort modelling data set previously used for models standardised in ASHRAE 55, and data from our own thermal comfort studies. We show that the model generally outperforms existing approaches from ASHRAE 55 after about 5 observations or less and converges to its final solution quality after about 10 observations. On average, the comfort model provides 13.2% - 25.8% more accurate predictions of user's thermal thermal preferences than the existing ASHRAE 55 approaches. We further propose a novel feedback scale and show its potential to improve model accuracy.

The comfort model has been published at the following venues:

- "A Personalised Thermal Comfort Model using a Bayesian Network" at the International Joint Conference on Artifcial Intelligence (IJCAI 2015) (Auffenberg et al., 2015a). This publication introduces the comfort model described in Chapter 3 of this work and evaluates it against the ASHRAE RP-884 data set.
- "A Comfort-Based Approach to Smart Heating and Air Conditioning" in the Urban Intelligence special issue of ACM Transactions on Intelligent Systems and Technology (TIST) (Auffenberg et al., 2017, in press). This publication contains a more detailed discussion of the comfort model discussed in Chapter 3 and an evaluation of the model against the ASHRAE RP-884 data set and our own data set introduced in Chapter 3. Furthermore, this publication introduces our semi-autonomous HVAC agent discussed in Chapter 4.

The HVAC scheduling algorithm utilises the thermal comfort model and its outputs to (1) minimise energy consumption of the HVAC system and (2) shift HVAC loads to fit a specified target profile. The algorithm is constrained to always stay within the user's comfort range if the weather conditions allow it to. Within the bounds of the comfort range, the algorithm is able to achieve optimal results with regards to energy saving. We extend the algorithm to also incorporate DSM signals. To maximise control over the algorithm's behaviour, we use target profiles specifying the actual shape of the preferred load profiles rather than dynamic energy pricing profiles. Further, we design the algorithm to allow full control over how much extra energy can be used to achieve the DSM target.

The HVAC scheduling algorithm has been implemented as a Mixed Integer Quadratic Program (MIQP), which offers optimal results. To make the computation of HVAC schedules tractable for any number of households and to protect users' privacy, we provide a distributed version of the algorithm as well. Through simulations based on real data, we show that when using the BCM as opposed to a single fixed set point, the algorithm is able to achieve load profiles 29.8% to 39.9% closer a given target profile. We further show that the distributed, locally optimal version of the algorithm achieves results close to the globally optimal centralised version, which reduces the deviation from a given target profile by another 12.8% as compared to the distributed version.

1.4 Report Outline

The remainder of this report is structured as follows:

Chapter 2 reviews background literature on the topic. It explains what an HVAC agent is and what HVAC scheduling involves. Existing thermal comfort models introduced earlier are presented in detail and compared to each other. A brief introduction to DSM is given and common HVAC-based DSM approaches are explained. The chapter then gives an introduction to belief networks and expectation propagation, the core technology used for the thermal comfort model presented in this work. The chapter concludes with a summary and discussion of the described technologies.

In Chapter 3, the technical details of the thermal comfort model are presented (published in Auffenberg et al. (2015a) and Auffenberg et al. (2017, in press)). The model is split into two separate parts, one representing elements of classic static models and one containing elements related to adaptations by the user. The model is then evaluated using real world data. For the evaluation, we use data from the ASHRAE RP-884 database and data from two studies conducted as part of this research which are explained in detail in this chapter. This chapter addresses the first research challenge.

Chapter 4 introduces a simple HVAC agent that utilises the thermal comfort model introduced in Chapter 3 to minimise energy consumption for the heating system without compromising user comfort (published in Auffenberg et al. (2017, in press)). We introduce and evaluate different strategies for finding compromise comfort ranges for different users at the same time. This is followed by a simulation-based evaluation of the HVAC agent's energy saving potential based on the different comfort compromising strategies. This chapter addresses the second research challenge.

Building upon the HVAC agent introduced in Chapter 4, Chapter 5 introduces an algorithm to incorporate DSM signals into the HVAC agent. The algorithm incorporates target profiles defining the most desirable load shape into its HVAC scheduler. Since a centralised execution of this algorithm will quickly become intractable, we propose a distributed version of the algorithm that executes in constant time regardless of the number of participating households. Using real data-based simulations of users and households, the load shifting potential of this algorithm is evaluated. We evaluate how the algorithm's performance changes when using the thermal comfort model introduced in Chapter 3 as opposed to a fixed set point approach. In addition, we compare how the distributed version compares to the globally optimal centralised version of the algorithm. This chapter addresses the third research challenge.

Chapter 6 summarises the findings of this research and provides an outline of planned research following this work.

Chapter 2

Background

This chapter reviews existing literature and key technologies of this research. First, general challenges of space heating and cooling are described and existing approaches to address those challenges are introduced. This is followed by an analysis of existing thermal comfort models. We then introduce and discuss existing DSM approaches. After that, a short introduction to belief networks, a key tool that we will use in our work, and how to perform inference and parameter learning on them is given. The chapter concludes with a summary discussing the technologies presented in this chapter.

2.1 Smart Home Heating and Cooling

Smart home heating and cooling refers to automated HVAC systems that go beyond simple set points and minimalistic scheduling. The heating and cooling system includes every appliance that is part of the space climatisation process. This includes for example radiators, AC units, thermostats or boilers. A smart heating and cooling system usually offers some additional functionality to automate the process. Some systems might even act autonomously as an intelligent agent, making own, autonomous decisions without a user telling it to do so. In this section, we list the factors that influence smart home heating and cooling and present enabling or existing technologies.

Smart home heating and cooling systems consist of three main components: (i) a suitable infrastructure that allows automatic control of components of the HVAC system, (ii) the thermal environment and (iii) the control mechanism itself. As this work focuses on the control mechanism, which is mainly dependent on the thermal environment, only a brief summary of existing infrastructure solutions is given.

Smart heating and cooling systems normally require a special infrastructure. There are already several manufacturers that offer their own solutions. Often, these are part of complete home automation solutions. However, most of these products are proprietary and only allow interaction through products by the manufacturer or using software provided by the manufacturer. Z-Wave[1] is one of the more open approaches with open communication protocols available to the general public. Z-Wave primarily describes a wireless communication technology through which devices can communicate with each other. Besides the wireless communication technology, Z-Wave also defines standardised protocols that define how to interact with Z-Wave devices. These standardised protocols allow easy interaction with the devices from other sources.

In the following sections, the thermal environment component and control mechanisms are going to be described.

2.1.1 Thermal Environment

For a smart heating and cooling system to function, it needs to be aware of its thermal environment. For a smart thermostat that only performs simple scheduling with fixed set points and predefined heating intervals, Rogers et al. (2011) state that the thermal environment of a house is sufficiently described by the thermal output of the heater and the leakage rate depending on other parameters such as the thermal capacity of the home and mass of air inside it. Rogers et al. develop a simple model to represent the thermal environment, which gives a good approximation to the actual thermal environment using only heater output and leakage rate as parameters. They also describe how those parameters can be learned from a set of measurements of interior temperature and outside temperature. As warming and cooling follow similar rules in the simplified, learned model, the model can be extended to account for AC systems by adding a *cooler output* variable that acts similar to the heater output.

While leakage rate, heater output and cooler output may suffice to describe the thermal environment for a thermostat that only needs to be able to heat to specific temperatures at given times, thermal comfort models require additional parameters to describe the thermal environment. Instead of working with the air temperature (also called dry bulb temperature — the temperature usually measured by normal thermometers), thermal comfort models usually work with the operative temperature that is composed from the mean radiant temperature (MRT) and air temperature. The MRT describes the temperature experienced from radiation. Typical sources of radiant heat are radiators and the sun. The operative temperature gives a better estimate of the temperature experienced by an individual in the room. In addition, as for example defined in ASHRAE 55, factors such as humidity and draught need to be taken into account as well.

From that, the following list of parameters defining the thermal environment for a smart thermostat including a thermal comfort model can be derived:

[1] Z-Wave Home Control `http://www.z-wave.com`

- Heater output
- Cooler output
- Leakage rate of the house
- Operative temperature
- Inside humidity
- Draught

In the following, we define these variables in detail and describe how to measure or calculate them.

2.1.1.1 Heater Output, Cooler Output and Leakage Rate

Using only heater output and leakage rate, Rogers et al. (2011) are able to predict temperature changes in the house. Rogers et al. define the heater output as the expected increase of air temperature due to heating and leakage rate as the expected decrease of air temperature due to leakage. By combining a heater output $R \in \mathbb{R}^+$ (°C/hr), and leakage rate $\Phi \in \mathbb{R}^+$ (1/hr), Rogers et al. define a simple equation to predict changes of air temperature for a given time slot:

$$T^{t+1}_{\text{air}} = T^t_{\text{air}} + \Big[R\eta^t_{\text{on}} - \Phi(T^t_{\text{air}} - T^t_{\text{out}})\Big]\Delta t \tag{2.1}$$

T^{t+1}_{air} is the expected inside air temperature for time $t+1$, T^t_{air} the expected inside air temperature at time t and T^t_{out} the outside temperature at time t. The variable η^t_{on} defines whether the heating is switched on ($\eta_{\text{on}} = 1$) or off ($\eta_{\text{on}} = 0$) at a given time t.

To learn leakage rate and heater output, the error between predicted temperature and actual temperature is minimised. The error is calculated using the following equation:

$$\sum_{t \in T} (T^t_{\text{air}} - T^t_{\text{act}})^2 \tag{2.2}$$

where T^t_{air} is the expected and T^t_{act} is the actual measured air temperature at time t. As T is the same and therefore $|T|$ stays the same it is sufficient to compare the aggregate square error rather than the mean squared error.

As mentioned earlier, the model can easily be extended to account for AC. The cooling process achieved by the AC can be expressed by a cooler output $C \in R^-$ (°C/hr) that acts similar to the heating output. The cooling effect achieved by the AC can therefore be described by the following term:

$$\Delta T_{AC}(\Delta t) = C\zeta_{\text{on}}\Delta t \tag{2.3}$$

where C describes by how much the AC can reduce the temperature per hour and ζ_{on} describes whether the AC is switched on ($\zeta_{\text{on}} = 1$) or not ($\zeta_{\text{on}} = 0$). Adding this to equation 2.1 yields the full equation for predicting the indoor temperature for a time step $t+1$:

$$T_{\text{air}}^{t+1} = T_{\text{air}}^{t} + \Big[R\eta_{\text{on}}^{t} + C\zeta_{\text{on}}^{t} - \Phi(T_{\text{air}}^{t} - T_{\text{out}}^{t})\Big]\Delta t \tag{2.4}$$

2.1.1.2 Operative Temperature

The operative temperature is a temperature used to more accurately describe the perceived temperature by an individual. It combines different temperatures together, the most dominant being air temperature and radiant temperature. In addition, air movement can be taken into account as well, as it changes the input of air or radiant temperature on the actual temperature perception of a person. The air temperature describes the actual temperature of the air in the room. A regular thermometer — if used properly — usually measures the air temperature. The radiant temperature describes a temperature caused by radiated heat. Usually, hot surfaces do not only heat up the air touching them, but also radiate heat in the form of infrared radiation. Radiated heat has a dominant impact on a person's perception of temperature (ASHRAE 55) and therefore needs to be taken into account when calculating thermal comfort.

The operative temperature cannot be measured directly and the exact calculation can be very complex. Therefore, the operative temperature is often approximated. The ASHRAE 55 standard defines the following approximation of the operative temperature:

$$T_{\text{op}} = A\,T_{\text{air}} + (1 - A)\,T_{\text{mrt}} \tag{2.5}$$

where T_{air} represents the air temperature, T_{mrt} the mean radiant temperature and A a factor depending on the air velocity or draught. Table 2.1 shows how A is defined in ASHRAE 55.

Air velocity	A
< 0.2 m/s	0.5
0.2 to 0.6 m/s	0.6
0.6 to 1.0 m/s	0.7

Table 2.1: Factor A depending on Air velocity as defined by ASHRAE 55

This approximation however still requires a value for the radiant temperature, which is hard to measure in regular households. To represent the radiant temperature in a room, the mean radiant temperature of all surfaces is normally used. A simple approximation of the mean radiant temperature can be calculated by taking the average surface temperature relative to surface area as described in chapter 9 of the ASHRAE Fundamentals

Handbook (2009):

$$T_{\mathrm{mrt}} = \frac{\sum_{s \in S} T_s A_s}{\sum_{s \in S} A_s} \tag{2.6}$$

where T_{mrt} is the mean radiant temperature, T_s is the temperature of surface s, S is the set of all surfaces in the room and A_s is the surface area of surface s. To get a more precise estimate, the relative angle between surface and a person can be taken into consideration as well. However, for this the comfort model the exact dimensions of a room and the exact position of a person in the room would need to be known at any time. As a person's position can be expected to change frequently, the simple approximation which does not consider relative angles is used in practice. In Chapter 3 we will describe further how the radiant temperature is estimated in the model presented in this work.

2.1.1.3 Interior Humidity

Humidity describes the amount of water vapour in the air and is normally measured using a psychrometer or hygrometer. A common way to describe the relative humidity is to take the ratio of the actual vapor density (A_{vd}) and the saturation vapor density ($S_{\mathrm{vd}}(T)$). The A_{vd} describes the actual amount of water vapor in the air, the $S_{\mathrm{vd}}(T)$ describes the upper bound of how much water vapor the air can hold at air temperature T. The fraction of these two values yields the relative humidity h:

$$h = \frac{A_{\mathrm{vd}}}{S_{\mathrm{vd}}(T)} \tag{2.7}$$

2.1.1.4 Draught

In the terminology of thermal comfort, draught describes the cooling effect caused by air movement. Draught can result from multiple sources such as open or poorly closed windows or doors and from the so-called stack effect. The stack effect describes draught caused by temperature differences between the inside of the house and the outside (Klote (1991), The Engineering Toolbox (2014)). If the air inside the house is warmer and therefore the density of the air inside is lower, the air will move upwards inside of the house. If the air inside is colder than outside, the higher density of the air inside of the house causes the air to move downwards. The natural draught induced by the stack effect can be calculated using thermodynamics (Klote, 1991). These calculations however require several variables like hydraulic diameter or the minor loss coefficient that are complex to measure and hard to estimate for normal users. Hydraulic diameter for example also depends on how widely a door or window is opened.

In the scope of thermal comfort, ASHRAE 55 only defines a rough threshold for the values of draught. The standard suggests that for air temperatures below 22.5°C, the draught should not exceed $0.15^{\mathrm{m}}/_{\mathrm{s}}$ if it is not under the occupants' direct control. In a comfort model that takes a user's preferences and the current state of the environment into account, a more precise estimate of the influence of draught can be made. As draught has a cooling effect that grows with the velocity of the air, it correlates with the thermal comfort temperature. That means that as air velocity (or draught) rises, the temperature felt by a user drops. Therefore, the optimal air temperature that maximises comfort needs to be higher (rises with the draught).

2.1.2 HVAC scheduler

The HVAC scheduler can be seen as the core component of a smart HVAC system. The scheduler coordinates where and when to turn on the heating or air conditioning in order to regulate the temperature in different rooms. Conventional thermostats have very limited scheduling abilities. Usually such thermostats define a comfort band and monitor the temperature. As soon as the temperature leaves the comfort band, the thermostat takes actions. Some thermostats might allow the configuration of different comfort bands for different times of the day. Smart thermostats offer more powerful scheduling options. The scheduler in a smart thermostat might for example consider the thermal environment, changing comfort bands, home occupancy and other factors when deciding when to turn on the heat. The Nest learning thermostat for example learns the heating properties of the house to be able to switch the heating on early enough and can use additional sensors to track occupancy and switch the heating off accordingly. HVAC schedulers of existing smart thermostats can be categorised into two types of schedulers: HVAC agents and manually configured schedulers. HVAC agents usually work autonomously. That means that schedules are created and changed autonomously by the agent without the need for manual inputs by a user. An HVAC agent does not neet to operate fully autonomously, it is also possible to have a collective between the HVAC agent and users, where users for example provide feedback about the temperatures or take over control every now and then. Manually configured schedulers require the user to configure an exact, complete schedule. This work aims to minimise user input to ensure unobtrusiveness (research requirement 4) and is therefore focusing on HVAC agents rather than manually configured schedulers. Of the existing smart heating and cooling approaches introduced in Section 1.2, all can be categorised as HVAC agents. In the following sections, the schedulers of those systems will be explained in more detail.

2.1.2.1 NEST scheduler

The Nest learning thermostat[2] can be used as both a manual HVAC scheduler and as an HVAC agent. The HVAC agent tries to maximise thermal comfort solely by learning a

[2] Nest learning thermostat (`https://nest.com/uk/`)

precise set point schedule from the user. When used as an agent, the creation of schedules is based on two input variables: a general set point schedule learned from user inputs and home occupancy. If the home is occupied, the scheduler will follow the general set point schedule. When unoccupied, the scheduler will turn the temperature down to an away set point defined by the user. The manufacturer only provides minimal information about the underlying mechanisms in the auto scheduler. The following explanations of the underlying scheduling mechanisms are therefore based on information given on the manufacturer's website and on a study by Yang and Newman (2012) evaluating the Nest thermostat.

To determine home occupancy, the Nest HVAC agent uses a combination of activity sensors and special purpose algorithms[3]. If multiple Nest products are installed in the home, those can communicate with each other to cover more areas with their activity sensors. After a week of learning, the thermostat creates a model of when the house is usually occupied. It is not stated how the model works or how the learning process functions.

While the exact algorithms for generating the set point schedule are not described, some descriptions of the general mechanics and influence of user inputs are given on the manufacturer's website[4]. The learning of schedules works as a two step process. The first step is a quick learning phase that starts immediately after installing or resetting the thermostat and will end when a basic schedule has been learned. During that time, the thermostat will remember every manual user input and repeat it during the next days. After a basic schedule has been learned, the thermostat switches to the second phase during which it is less sensitive to user inputs. During this phase, to make lasting changes to the schedule, a repeating pattern of at least two similar changes is needed. There are different possible patterns:

Weekdays If the same change was made during two weekdays (excluding weekends) in a row, the schedule will be adjusted to repeat this input for every weekday.

Weekend The weekend day change is similar to the weekdays change. However, instead of requiring the same changes during two weekdays in a row, the changes are required for the two weekend days. Likewise, the schedule will only apply the changes to weekend days.

Full Week The full week change combines the weekdays and weekend change: if the same adjustment happens on a weekday and a weekend day, the schedule will be altered for all seven days of the week.

[3]Nest Auto-Away (`http://support.nest.com/uk/article/What-is-Auto-Away`)
[4]Nest Auto-Scheduler (`http://support.nest.com/uk/article/How-does-Auto-Schedule-learn`)

Single Weekday If the same change occurred twice on the same weekday (e.g. Monday), the thermostat will adjust the schedule to repeat this change only on this particular weekday.

The learning capabilities described so far are only used to determine to what temperature the thermostat should heat or cool the home at a given point in time. To ensure that the correct temperature is reached in time, the Nest HVAC agent also learns how long it takes the heating system of the house to heat it up to a certain temperature. Again, the exact underlying mechanisms are not open for the public.

As the Nest learning thermostat does not utilize a detailed thermal comfort model to determine set points, it fails to meet the first research requirement of providing accurate estimates of users' comfort levels. Studies (Yang and Newman, 2012, 2013; Yang et al., 2014) show that frequent user input and slight variations in the daily schedule can lead to schedules with frequently varying set point temperatures that are hard to understand by the user and that prohibit energy savings. Varying inputs from multiple users with different preferences will likely increase this problem, making the Nest learning thermostat particularly unsuitable for environments with multiple occupants participating in the system.

2.1.2.2 Adaptive Heating Algorithm by Shann and Seuken

Shann and Seuken (2014) use a Markov decision process (MDP) (Bellman, 1957) to decide when to turn on the heater. MDPs are an extension of Markov chains, where not only the previous state, but also effects of decisions of participants in the system are taken into account. Shann and Seuken model states of the system as a tuple of internal temperature, external temperature and time. Transitions describe the change of internal and external temperatures between two time steps. The outcome of the transitions depends on predictions of external temperatures and whether the heating was on or off during the transition. States are evaluated using a reward function that calculates the user's utility based on the internal temperature and subtracts the cost of heating based on predicted prices. When the heating was not on during the transition, the cost is zero. The user's utility function represents the comfort model in this approach and will therefore be further described in Section 2.2.2. Shann and Seuken exploit known relationships between variables to reduce the computational complexity of the transition function. When creating a schedule, Shann and Seuken always calculate a schedule for the next 24 hours. Energy prices and external temperatures for this interval are predicted using a Gaussian process (GP) (Seeger, 2004).

Shann and Seuken compare their MDP-based approach to a simple rule-based heating policy and two versions of a mixed-integer program. The rule-based heating policy tries to keep the temperature within a defined comfort band at all times. The first mixed-integer program represents the quadratically constrained mixed-integer program described by Rogers et al. (2011). This approach minimizes heating costs with the constraint of the discomfort not falling under a pre-defined level. The second mixed-integer program is a modification of the first one, where the objective function is altered to maximise the user's utility. The algorithms are compared by simulating 30 different days with time intervals of 15 minutes and taking the average cumulative utility of the algorithms. Shann and Seuken show that the average utility of users can be improved by 20-26% using MDPs.

In addition to this evaluation, the model was also tested in a real deployment (Alan et al., 2016; Shann et al., 2017). For this, a smart thermostat was installed in 30 households for a period of 30 days. The main goal of this evaluation was to test different user interfaces for the smart thermostat and to evaluate by how much the thermostat can reduce peak energy consumption in a real-time pricing scenario. The results show that this approach is able to reduce peak energy consumption by up to 38%.

The main shortcomings of this approach are the lack of a detailed thermal comfort model and the direct use of energy prices. Shann and Seuken use a modification of the adaptive comfort model defined in ASHRAE 55 which will be further described in Section 2.2.2. While Shann and Seuken parameterise the model and learn the parameters based on user feedback, the model neglects important factors such as humidity or seasonal adaptations, meaning that predictions obtained from this model are likely to be inaccurate (violating research requirement 3). As we will discuss in Section 2.3.2, directly incorporating energy prices into the HVAC scheduler can make the system response unstable (for example by introducing rebound effects) and hard to predict (violating research requirement 9). The rebound effect describes the situation where a collective of agents reacts similarly to the same pricing signal (i.e. using more energy when prices are low) which might introduce new peaks during times of low energy prices. This effect will be discussed in more detail in Section 2.3.2.

2.1.2.3 Distributed DSM by Ramchurn et al.

Ramchurn et al. (2011) propose an agent-based control strategy for performing DSM with HVAC loads. In this approach, each individual household is equipped with an individual agent controlling the HVAC system. These agents use a MIQP to minimize both, energy cost of the HVAC system as well as discomfort experienced by the user. The energy cost is further scaled by a weighting factor κ that allows for more fine grained control of the pay-off between cost and comfort. Ramchurn et al. set this value to a very low value by default to ensure a comfortable environment for the user. User comfort is assessed as a weighted, quadratic deviation of the indoor temperatuer from a static, user-specified optimal comfort temperature. Ramchurn et al. further allow to specify different costs for staying below or above the comfort temperature. This can be used for distinguishing between summer and winter times. It is likely that for most users going above the static comfort temperature is more acceptable during summer than during winter.

To incorporate DSM capabilities in the agents, Ramchurn et al. use a second objective of minimising the cost of the HVAC system. This objective function calculates the energy consumed by the HVAC system at different time steps and multiplies it with the energy prices at that time. This way, dynamic energy prices as commonly found in DSM approaches are incorporated into the computation by the agent. During times of low energy cost, using the HVAC system is penalised less than during times of high energy cost, causing the agent to shift its energy usage towards times of low energy prices.

Without further corrections, this approach however can be prone to introducing new peaks during times of low energy prices. This is commonly referred to as the rebound effect (Palensky and Dietrich, 2011). To address this issue, Ramchurn et al. introduce a correction parameter α that denotes the probability with which agents should react to pricing signals. This parameter is adjusted over time based on resulting load profiles. If a rebound effect is happening, the value for α is reduced, causing fewer agents to follow the pricing signals. If the response to the pricing signals is too weak and not enough load is shifted, α is increased, causing more agents to react to the pricing signals. Ramchurn et al. show that using this mechanism, load peaks in the grid can be reduced by up to 17%.

2.1.2.4 HVAC based load shifting by Lu

Lu (2012) presents a direct load control strategy that intelligently schedules when to switch on the heating in different households to follow different DSM signals. Lu classify households into two separate groups: those where the heater is currently running and those where the heater is currently switched off. These two groups are ordered by different priorities. In the heater on list, houses where the indoor temperature is close to the upper bound of the thermostat's temperature band are prioritised. In the heater off list, houses where the indoor temperature is close to the lower bound of the thermostat's temperature band are prioritised. When receiving DSM signals, the central controller determines in which houses the heater can be switched on or off respectively. To do so, the central controller has to predict the indoor temperature of each separate household. User comfort is not modelled directly. Instead, a static temperature band is used. Lu test two different widths for this temperature band and evaluate the performance of the control strategy with both of them. Using simulations based on real load shifting signals, Lu shows that when using the wider (4°C wide) temperature band, the control strategy successfully shifts demand to follow the provided signal.

A major drawback of this approach is its reliance on a centralised controller. It is not clear how well the algorithm scales up with the number of households (research requirements 6 and 8). Further, the approach requires potentially sensitive information about live indoor temperatures in each household to be sent to the central controller. Lastly, this approach oversimplyfies thermal comfort by just using a static range (violating research requirement 3).

2.1.2.5 Occupancy based demand response by Erickson and Cerpa

Erickson and Cerpa (2010) create a model to predict room occupancy in a house. Knowing when which room is occupied, the HVAC agent is able to only heat or cool where necessary, leading to significant reductions in energy consumption. Erickson and Cerpa utilise a Moving Window Markov Chain to model room occupancy throughout the day. To account for multiple occupants, the markov chain does not only model whether a room is occupied or not, but also how many occupants are occupying a room. Since room occupancy is highly dependent on the time of day, Erickson and Cerpa add a moving window on top of the Markov Chain. The window can move by an arbitrary amount of minutes and calculates transition probabilities between different occupancy states based on the last hour.

Based on their occupancy model, Erickson and Cerpa develop control strategies for ventilation as well as the heating and cooling system. Through simulations, Erickson and Cerpa show that their control strategy is able to reduce annual energy consumption of the HVAC system by about 20%. During the winter months, the energy saving potential is even higher with up to 26.5% energy savings, suggesting that the proposed strategies are especially successful in colder climates.

While Erickson and Cerpa use a sophisticated, general thermal comfort model defined in ASHRAE 55 to decide set point temperatures, they do not account for preferences of individual users. Such general models only represent average preferences of large populations. When dealing with small numbers of occupants, individual users' prefences may deviate from these average preferences, meaning that the chosen set point temperatures might actually cause discomfort to some users (conflicting with research requirement 3). In addition, Erickson and Cerpa do not provide any means to incorporate DSM signals into their account (research requirement 7).

2.1.2.6 Optimisation friendly thermal comfort model by Zhou et al.

Zhou et al. (2016) define a novel, optimisation friendly thermal comfort model. The comfort model is built as a convex piecewise linear classifier. Based on different inputs (such as outdoor temperature and humidity), the classifier creates multiple hyperplanes denoting comfort boundaries. Only a point lying in between all hyperplanes is considered comfortable. The resulting space is convex and therefore allows for easy optimisation.

Zhou et al. build an HVAC agent based on this model and evaluate it using weather data from Singapore. In the optimisation, Zhou et al. however do not focus on just finding the best combination of set points to reduce energy consumption while maintaining user comfort. Instead, the HVAC agents also considers manipulating other factors such as humidity. Zhou et al. show that in hot climates, manipulating the humidity instead of using the air conditioning can reduce energy consumption by 12.81%.

Similar to other approaches, Zhou et al. do not account for individual users' preferences, potentially leading to inaccurate predictions for comfort temperatures (violating research requirement 3). In addition, no means of incorporating DSM signals are provided (research requirement 7).

2.1.2.7 Participatory approaches

There are a couple of approaches that rely on participatory sensing to control the HVAC system. Such approaches usually utilise a constant stream of feedback about the thermal environment from the occupants to adjust the set point of the HVAC system.

Shetty et al. (2015) for example propose a participatory HVAC control approach that uses the user feedback to choose appropriate set point temperatures but also to suggest space allocations based on user preferences. In this approach, a pop up opens up on the user's screen every 30 minutes, asking to provide feedback on a simplified version of the ASHRAE 7-point scale that we will introduce in the next section. Using this data, a range of of cold, neutral and warm temperatures is determined for each user. Shetty et al. propose two strategies for how this data can be used to reduce energy consumption and maximise user comfort. The first strategy is to choose a set point that is as close to the outdoor temperature (minimising the temperature gradient between indoors and outdoors) as possible while still being in the range of neutral temperatures of all/most occupants. The second strategy is to relocate occupants with lower temperature preferences to colder parts of the building and occupants with preferences for higher temperatures to warmer parts respectively.

Purdon et al. (2013) create a model and sensor free HVAC agent that relies on live user votes to determine occupancy and discomfort. Purdon et al. provide the users with a smartphone app allowing them to rate the thermal environment as hot, cold or OK. The general idea is to maximise the amount of OK votes and balance the amount of hot and cold votes. To be able to reduce energy consumption, Purdon et al. rely on occupancy prediction and a concept they call *drift*. Occupancy prediction is determined by a lack of user votes. If no feedback is provided for a certain amount of time, the HVAC is switched off, reducing energy consumption. Further reductions are achieved by letting the indoor temperature *drift* towards the outdoor temperature as much as possible. Once a user provides feedback going against the drift (for example a cold vote when the temperature is drifting down), the HVAC agent stops the drift. Using simulations, Purdon et al. show that energy savings of up to 65% can be achieved.

Ghahramani et al. (2014) go beyond the aforementioned approaches by modelling individual users. Ghahramani et al. allow users to rate their thermal environment on a scale from -5 ("prefer cooler") to 5 ("prefer warmer"). Using pattern recognition on the users' votes and the respective indoor temperature, an individual comfort range per user is computed. Using the comfort ranges of all users, each day a new set point minimising discomfort and reducing energy consumption is chosen. Ghahramani et al. show that using their approach, significant reductions in airflow rates can be achieved, improving the efficiency of the HVAC system.

While participatory approaches generally perform well with respect to achieving a comfortable environment for the users, this usually comes at the cost of being very intrusive, requiring users to provide feedback frequently (violating research requirement 4). Further, the reliance on live feedback of participatory approaches complicates tasks that require the HVAC system to plan ahead, as is the case in most DSM scenarios (violating research requirement 7). As set point temperatures or comfort ranges are generally determined just-in-time based on user feedback, it is difficult for the HVAC scheduler to for example pre-heat or cool a room in response to higher or lower energy prices during certain intervals. To be able to plan ahead, the HVAC scheduler needs to utilise a thermal comfort model.

2.2 Thermal Comfort Models

Thermal comfort models can generally be categorised into two different categories: static thermal comfort models and adaptive thermal comfort models. Static models are usually designed for HVAC equipped buildings and assume static activity levels and clothing of occupants. Adaptive models assume that occupants have more individual control over their thermal environment and take adaptive measures of occupants into account. Typical adaptive measures are for example opening of windows or turning on a fan. As many, especially larger HVAC equipped buildings prevent occupants from using such measures, adaptive models are mostly used for naturally ventilated (NV) buildings. Due to the increased interest of the research community in domestic heating, more advanced adaptive models have been created recently. Such models do not only consider adaptive measures of occupants, but also account for individual preferences of occupants. The solution proposed in Chapter 3 of this work combines findings from classic, established models with recent findings of individualised models to create a new, more robust and accurate, individualised thermal comfort model.

In the following, the two categories are explained in detail and examples for comfort models of those categories are given and discussed.

2.2.1 Static Thermal Comfort Models

The most widely used static thermal comfort model is the model by Fanger (1970). Fanger's model defines two main measures for thermal comfort: the predicted mean vote (PMV) and the predicted percentage dissatisfied (PPD). The PMV describes how a group of people would rate their perception of the current thermal environment. Ratings are based on the 7-point scale shown in Table 2.2. According to ASHRAE 55, the thermal environment can be seen as comfortable if the PMV lies within the limits of ($-0.5 < \text{PMV} < 0.5$).

Vote	Thermal sensation
3	Hot
2	Warm
1	Slightly Warm
0	Neutral
-1	Slightly Cool
-2	Cool
-3	Cold

Table 2.2: 7-point thermal sensation scale

The PPD predicts the percentage of people whose thermal sensation lies outside the comfort limits. When designing buildings and deciding temperatures, one generally aims to minimise the PPD. After ASHRAE 55, the PPD needs to be below 20% to achieve thermal comfort.

The calculations of PMV and PPD are based on human heat-balance. Based on empirical studies, Fanger (1970) defines very specific limits for skin temperature and sweat secretion. Within those limits, heat-balance is achieved, resulting in thermal comfort. Skin temperature and sweat secretion are calculated based on the following six factors: air temperature, mean radiant temperature, relative humidity, air speed, metabolic rate and clothing.

Static thermal comfort models have several limitations. The biggest limitation originates from the nature of the data that is needed to evaluate the model. Factors such as air speed, mean radiant temperature and relative humidity require specific sensors or detailed knowledge of the thermal environment. If present at all, such sensors or knowledge usually only exists for buildings that are equipped with modern HVAC systems. This limits the applicability of such models greatly as many buildings — especially in the domestic sector — are not equipped with such systems. The second limitation is due to the focus on metabolic rate and clothing. Static models assume that metabolic rate and clothing are not subject to changes. While this assumption may hold in some spaces (i.e. office spaces with strict regulation on clothing), there are many situations where individuals are able to influence their thermal environment through adaptive means such as changing their clothing or moving around. These limitations led to the development of adaptive thermal comfort models, described in the next section.

2.2.2 Adaptive Thermal Comfort Models

Adaptive thermal comfort models account for individual adaptations of occupants to the thermal environment. As described by Liu et al. (2012b), there are three main categories of adaptations that can take place: physiological adaptation, behavioural adaptation and psychological adaptation.

Physiological Adaptation describes adaptations to the environment by the human body. One has to distinguish between genetic adaptation which takes several generations and acclimatisation. In the field of thermal comfort special attention has to be paid to acclimatisation. Typical measures for acclimatisation is adjustments to the skin blood flow, shivering or sweating.

Behavioural Adaptation describes behaviours performed by an individual to adapt to the thermal environment. One distinguishes between three categories of behavioural adaptation: personal adaptation, technological adaptation and cultural responses. Personal adaptive measures are for example taking off an item of clothing. Typical technological measures are turning on a radiator or air conditioning. Cultural responses incorporate measures such as having breaks during the hottest hours of the day.

Psychological Adaptation describes altered perception of and reaction to temperatures as a reaction to the past and expectations of the thermal environment. de Dear and Brager (1998) and Paciuk (1990) for example analysed how the perception of having control influences people's thermal comfort. They found that occupants are more likely to accept certain disturbances to their thermal environment when they believe they are still in charge of them.

Adaptive thermal comfort models are usually aimed at NV buildings. This is due to the fact that in NV buildings occupants usually have more control over their thermal environment. Buildings equipped with sophisticated HVAC systems are often office buildings where occupants are not able or allowed to simply open windows or turn on fans. In addition, as shown by de Dear and Brager (1998), the effects of psychological adaptations are stronger in NV buildings. The relation between outdoor temperature and comfort temperatures is for example twice as strong in NV buildings. As a result, many adaptive thermal comfort models use the outdoor temperature as the main indicator for the comfort temperature. The best known adaptive comfort model is based on the findings of de Dear and Brager (1998) and has been included into ASHRAE-55 in 2004. The model however is based on average values from a large number of households and thus is mainly applicable to spaces with large groups of occupants. It fails to adapt to the preferences of individual users and is therefore likely to produce unsatisfying results in many cases. More recent approaches therefore shift the focus on learning relationships specific to individual users. In the following, we introduce a typical learning based approach.

Shann and Seuken (2013) model thermal comfort with respect to cost and comfort. The overall utility function subtracts the cost for heating to a given inside temperature from the comfort value of that temperature for the user.

Similar to de Dear and Brager (1998), Shann and Seuken focus on the relation between outside temperature and comfort temperature. Shann and Seuken model this as a linear relationship with a base temperature T^* and a slope m. The base temperature describes the user's optimal comfort temperature for an outside temperature of 0°C. This results in the following equation:

$$T_{pref}(T_{out}) = T^* + mT_{out} \tag{2.8}$$

where $T_{pref}(T_{out})$ is the preferred temperature T_{pref} for a given outside temperature T_{out}. To determine a user's utility for a given interior temperature, Shann and Seuken use the quadratic deviation of the interior temperature from the optimal comfort temperature of a user. To account for higher sensitivity to temperature deviations caused by outside temperatures below 0°C as described by Peeters et al. (2009), Shann and Seuken scale the utility with the term $b\,e^{-cT_{out}}$. The variable b represents the user's sensitivity at 0°C and c describes how much the user's sensitivity changes when the outside temperature changes. The complete utility function is described in the following equation:

$$v(T, T_{out}) = a - b\,e^{-cT_{out}}\big((T^* + mT_{out}) - T\big)^2 \tag{2.9}$$

where T is the actual interior temperature and a is the user's base utility for the optimal temperature.

The cost of heating is calculated using the following equation:

$$c(p, T, T_{out}) = p|T - T_{out}| \tag{2.10}$$

where p represents the current price for heating, T represents the desired interior temperature and T_{out} the outside temperature. Shann and Seuken base this function on the fact that heating costs are proportional to the temperature difference of inside and outside temperature.

Putting everything together, the user's utility can be described as follows:

$$u(p, T, T_{out}) = v(T, T_{out}) - c(p, T, T_{out}) \tag{2.11}$$

The optimal interior temperature is found by maximising the utility with respect to T.

To adapt to the user's preferences, Bayesian learning is used to learn the parameter vector (b, c, T^*, m), which describes the user's preferences in this model. Shann and Seuken define a Gaussian prior for this vector, which gets updated using Bayes' theorem whenever the user provides feedback.

While the adaptive thermal comfort model by Shann and Seuken tries to adapt to users' preferences, the underlying model is rather simple and neglects other factors defined by static thermal comfort models. As a result, certain limitations in the accuracy of predictions of user's comfort temperatures can be expected. Especially in varying weather conditions beyond outside temperature (for example in humidity) we expect some accuracy losses, violating the requirement of an accurate comfort model.

2.3 Demand Side Management

Demand side management (DSM) describes the control of loads on the electricity grid caused by end-users. Another common term in this field is demand response (DR) which describes a subcategory of DSM. DR focuses on designing incentives (for example financial savings through variable pricing) for customers to participate voluntarily (Gelazanskas and Gamage, 2014). There are a couple of existing problems and challenges which can be addressed by DSM.

One of them is to flatten peaks in demand. To ensure supply stability and avoid blackouts, the grid and generation capacities need to be large enough to cover the maximum expected demand. As a result, this infrastructure is working below its maximum capacity most of the time. As stated by Farhangi (2010), 20% of power generation capacity only exists to deal with peaks in demand. This capacity is only used 5% of the time, meaning that suppliers are investing in infrastructure that is rarely used.

Another factor is the optimal utilisation of renewable energy sources. Many renewable energy sources such as wind or solar energy are beyond human control and therefore not plannable. As a result, there may be overproduction at some times and not enough energy during other times. A straightforward approach to overcome this problem would be to store energy during times of overproduction and release it during times when the actual production does not suffice anymore. However, to this day such large storage capacities are not yet available. DSM can be utilised to compensate for fluctuations in supply by reshaping demand to use more energy while production is high and less when production is low.

In both applications the problems are solved through *load shaping*. As the name suggests, load shaping describes a process in which the load profile on the grid over a course of time is modified to fit certain requirements. There are six different types of load shaping (Gelazanskas and Gamage, 2014):

Peak Clipping describes a form of load shaping in which demand is limited to not go over a certain threshold during peak times. This results in a profile where the top of the peak is "cut off" at a certain level (Figure 2.1a).

Valley Filling means that during times of low demand (valleys), the demand is increased, evening out the dip in the load profile (Figure 2.1b).

Load Shifting refers to the process of moving parts of the load to a different time interval. Loads occurring during a peak can be for example shifted to time intervals before and after the peak (Figure 2.1c).

Strategic Conservation aims to reduce energy consumption in general, for example to bridge times of low production due to lack of wind or sun (Figure 2.1d).

Strategic Load Growth is the opposite of strategic conservation. Instead of reducing the overall demand, demand is increased to for example make best use of renewable energy sources at times of high production such as very windy and sunny days (Figure 2.1e).

Flexible Load Shaping is often a combination of all other types. Flexible load shaping tries to modify the load to fit any kind of predefined profile (for example filling valleys, clipping peaks and grow load during times of high supply at the same time) (Figure 2.1f).

In practice there are currently two main ways to approach load shaping: price-based and incentive-based approaches (Wang et al., 2013). Price-based approaches are based on the assumption that in order to save money, end-users are willing to adjust their energy consumption behaviour. By having variable prices, user behaviour is altered to change demand in a certain way. Incentive-based approaches usually do not rely on the assumption made in price-based DSM but rather use other means of controlling load and provide incentives such as general rebates or bonuses for the user to hand over control over their appliances.

In the following, we will list enabling technologies needed in DSM and briefly introduce and discuss price-based and incentive-based DSM approaches. This is followed by a rough categorisation of different load types and their possible utilisation in DSM. We show where HVAC fits in these categories and provide a short discussion of the current state of the art of HVAC based on approaches introduced in Section 2.1.2.

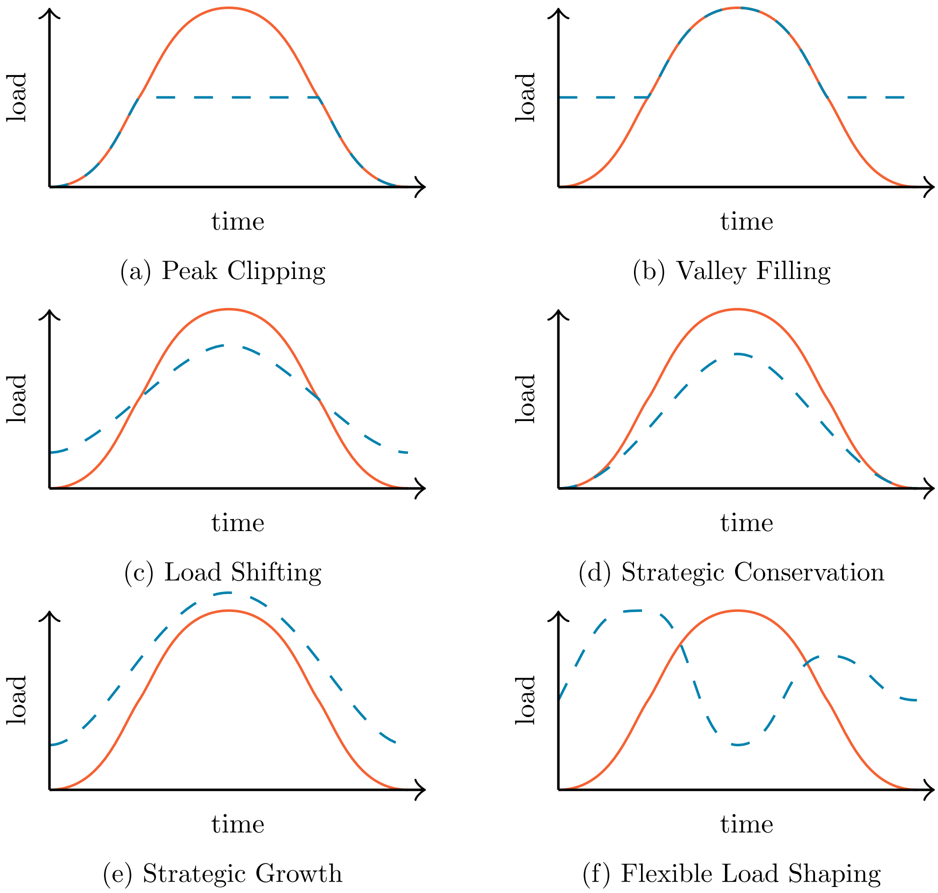

(a) Peak Clipping
(b) Valley Filling
(c) Load Shifting
(d) Strategic Conservation
(e) Strategic Growth
(f) Flexible Load Shaping

Figure 2.1: Different types of load shaping (red: original profile, blue: shaped profile) as defined by Wang et al. (2013)

2.3.1 Enabling Technologies

Most existing approaches to managing and influencing demand have common underlying requirements. To be able to evaluate how an approach will perform on a larger scale with several separate households, a model to calculate the aggregated load on the grid is required. This is referred to as *load modelling*. Another requirement is some mean of managing the energy in the household itself. Existing approaches cover a wide range from simply telling users to turn appliances on or off to managing appliances autonomously without any user interaction. This is referred to as an *energy management system*. In the following we will give a brief overview over existing technologies for those requirements.

2.3.1.1 Load Modelling

As mentioned before, to evaluate approaches before applying them in the real world simulations need to be carried out to see how the load is affected by an approach. There are two different ways of approaching load modelling: bottom-up and top-down (Swan and Ugursal, 2009).

Bottom up models usually create load profiles for individual parts of an area and aggregate those individual profiles into a full load profile for the whole area of interest. Depending on the model, the granularity for the individual parts can range from individual household appliances up to individual groups of houses. Within the field of bottom-up models, one can further distinguish between statistical methods relying on historical information on energy usage of different appliances or houses and engineering methods trying to accurately model consumption based on actual energy ratings and physical calculations (for example thermodynamics).

Top-down modelling is based on historical data and population statistics and does not model individual end-users. Typically, top-down models use information on macroeconomic indicators (gross domestic product, unemployment rates and others), climate and appliance ownership to estimate the aggregate load for a whole area. Within the field of top-down models one can further distinguish between econometric models that focus mainly on price and technological models that are mainly built around housing attributes such as typical appliance ownership in certain areas. In addition, there exist hybrids of econometric and technological models.

As top-down models are strongly relying on historical data, predictions given by such models are usually very precise. This makes top-down models especially suitable to predict changes in energy usage of a whole area if for example new houses are built or older houses are refurbished. However, the reliance on historical data also makes such models harder to apply in many modern DSM approaches for which not much historical data exists. A major drawback of bottom-up models is that extensive input data and complex models are required to accurately predict the load.

For this work, which is based on the fairly easily predictable process of heating and cooling, a combination of both approaches can be used. Based on historical data, a general load profile for individual households can be created. Using bottom-up calculations, individual profiles for heating and cooling systems can be created and removed from the overall household profile. To evaluate how a heating/cooling based DSM approach affects the overall load, new heating and cooling profiles resulting from said DSM approach can be created and factored back into the overall household profile.

2.3.1.2 Energy Management Systems

To be able to use DSM algorithms in practice, means to control (for example whether to turn them on or off) appliances in a household are required. For private homes this is referred to as a home energy management system (HEMS). The main task of a HEMS is to schedule when to turn on and off appliances in the household. There are two common ways of controlling the appliances: by getting a real person to control appliances or by controlling them automatically through specialised infrastructure (Wang et al., 2013).

While at first glance, getting a person to control the appliances seems like an easy task, existing research shows that getting the users to actually participate and keep them involved in the system can be very challenging (Murtagh et al., 2014b,a; Hargreaves et al., 2013; Snow et al., 2013). A typical implementation of a HEMS that relies on the user is to have in-home displays showing information about current energy prices, energy consumption and maybe even suggestions of which appliances to turn on or off. As shown by various researchers (Murtagh et al., 2014a; Hargreaves et al., 2013; Snow et al., 2013), users are often not very interested in this information and therefore quickly stop participating. For example, in the study conducted by Murtagh et al. (2014b), only 20% of participants had a positive attitude towards in home displays, 60% had no preference and 20% saw them as annoyance. Murtagh et al. (2014b) further found that especially HEMS that only provide the user with information about current energy prices and energy usage ask too much from the user. Users disliked the fact that they had to think and decide themselves when to turn which appliances on or off.

Automatically controlled appliances on the other hand often take away too much control from the user (Murtagh et al., 2014b) and require investments for upgrading existing or buying new appliances. Further, recent research suggests that users should be kept in the loop to some extent to avoid creating apathetic, disengaged users. As a result, hybrid approaches where some appliances are manually and some automatically controlled are often applied.

2.3.2 Price-based DSM

Price-based DSM refers to DSM techniques that use variable energy prices to influence demand. Such approaches are based on the assumption that in order to save money, users are willing to change their behaviour. In practice that means that users will reduce their energy consumption when prices are high and increase it when prices are low. This can be through manual control or through automated systems trying to optimise energy usage on the users behalf. Typically, prices would be higher during times of peak energy consumption or low energy production and low during times of excess production from renewable sources. Further, prices could reflect regional differences, taking into consideration weather constraints for renewables as well as cost of delivering the energy to a particular region. There are three main categories for how prices can be changed (Wang et al., 2013):

Time-of-use there are fixed prices for different time intervals. Typical setups define hourly prices for a 24 hour period.

Real-time pricing prices constantly change depending on different factors (wholesale price for electricity, demand and more). Changes are usually announced an hour or day ahead.

Critical peak pricing similar to time-of-use pricing but under certain circumstances (peak demand, low supply) much higher real-time rates can be applied.

In practice, price-based DSM is not well understood yet and yields mixed results. Thorsnes et al. (2012) for example show that during winter, time-of-use pricing can consistently reduce peak consumption on average by 10-15%. Over the course of the whole year however, time-of-use pricing had no significant effect on the overall daily consumption or peak consumption. Further, while off-peak consumption correlates with off-peak prices, higher peak prices had no significant effect on peak consumption. There is also a risk of introducing a so called *rebound effect* (Palensky and Dietrich, 2011). As the number of participating households increases, there is a growing risk that a large number of households reacts in a similar fashion to pricing signals. This mass behaviour could lead to new peaks, for example during times when energy prices are very low. Lastly, as mentioned in Section 2.3.1.2, users are not always keen to participate (Murtagh et al., 2014a,b; Hargreaves et al., 2013; Snow et al., 2013).

2.3.3 Incentive-based DSM

Incentive-based DSM describes approaches that set different incentives to motivate users to participate or hand over control to a DSM system. Similar to price-based approaches, a typical incentive is saving money. However, instead of letting the system regulate itself through variable prices, incentive-based approaches usually utilise rewards, rebates and possibly penalties. Users could be granted rebates for handing over some control to their DSM system. Handing over control can be in the form of the system restricting usage for certain appliances during certain times (for example using the washing machine would only be allowed during non-peak hours). Some approaches even go as far as modelling homes as suppliers that autonomously store and sell back energy to the grid (for example using their electric vehicles as storage) (Kahrobaee et al., 2013).

A major downside of many incentive-based approaches is that they often are fairly intrusive, demanding control over appliances. In direct load control (DLC) for example, appliances can be switched off remotely if required. As shown by Murtagh et al. (2014b), most users dislike the fact of not being in control as it can interfere with their personal schedules. In reaction to this, incentive-based approaches usually focus on loads that are not interfering with the user's schedules, such as thermal loads, which will be explained further in the next section.

2.3.4 Load types

As mentioned earlier, there are different kinds of loads that can be utilised differently in DSM. Ramchurn et al. (2011) define three categories of loads in a household:

Shiftable static loads include appliances that are not linked to heating or cooling but could still be controlled by a DSM system. Washing machines, tumble dryers and dishwashers belong to this category.

Thermal loads include appliances that are linked to heating or cooling. Such loads are usually considered shiftable and include appliances such as water and space heating but also fridges or freezers.

Non-shiftable loads include appliances that cannot or should not be controlled by the DSM system. This includes appliances such as lights, oven, kettle, microwave and many others.

It can be seen that the only categories of loads that can be utilised in DSM are shiftable static loads and thermal loads. In theory, shiftable static loads can be shifted without disrupting a user's daily life. Whether the dishes are cleaned in the evening or during the night will not make a big difference for the user. However, there are certain limitations. As Holyhead et al. (2015) suggest, static shiftable loads can usually be only shifted backwards as they have to be loaded manually by the user first. To overcome this issue, users could be prompted to prepare appliances by the DSM system or have to notify the system in advance about planned usages. This however would mean several disruptions and changes to the user's daily routines. There are other constraints like the amount of laundry or dirty dishes to fill a machine. Usually one would wait until the machines are filled before turning them on. A study by Costanza et al. (2014) on automating the laundry using agents confirmed several of these points and also highlighted that trying to control such loads can in fact increase load rather than shifting it. Some users for example noticed that in reaction to low prices they were doing the laundry more often with only partly filled machines, resulting in overall higher energy consumption of the household.

As opposed to shiftable static loads, thermal loads can be shifted fairly easily with only minor interference with the user's schedules. This is due to the fact that heating and cooling is an inert process, meaning that for example the water in a boiler will stay warm for some time after turning off the heating module. Further, these appliances mostly work in the background already without the need of manual control (Du and Lu, 2011). So as long as the water is always warm enough, rooms stay at a comfortable temperature and the fridge stays cold enough, users are generally not concerned with details of when those appliances are actually turned on and off.

2.3.5 HVAC and DSM

As mentioned earlier, in theory thermal loads offer great potential for DSM (Ramchurn et al., 2011). In practice there are a couple of challenges. In case of controlling a fridge for example, hardware to remotely control cooling cycles of the fridge needs to be developed and added to the existing infrastructure. For water and space heating or cooling this particular challenge has been addressed already. In most households, water and space heaters are already controlled by a thermostat. By replacing existing thermostats with smart thermostats, full control over these appliances can be obtained. Another issue is that in many countries, only the minority of households use electric heaters at this point (according to Ramchurn et al. (2011) only 7% of UK households use electric heaters). However, this is expected to change as production of electricity becomes cleaner due to the growing share of renewables. Heating is expected to move away from fossil fuels to use this clean energy (Department of Energy and Climate Change, 2013). In combination with already electrified air conditioning units, significant electric loads are accumulated.

Another major challenge of DSM on HVAC loads is user comfort. Shifting HVAC loads usually means to switch off or lower the heating or cooling during certain timese, resulting in higher or lower temperatures. These temperatures need to be in accordance with users' preferences. To predict users' preferences, one can utilise a thermal comfort model. However, as discussed in Section 2.2, amongst other shortcomings, existing thermal comfort models usually only provide single comfort temperatures or very narrow bounds. As the temperature in a room will decay towards the outside temperature as soon as the heating or AC is turned off, these loads can only be shifted in fairly short intervals. As a result, many approaches focusing on DSM have so far avoided space heating or cooling or accepted comfort trade-offs to achieve good results in DSM. Du and Lu (2011) for example only focus on controlling the boiler and avoid space heating and cooling completely. Even though Ramchurn et al. (2011) emphasize heating in their approach, they do not define how the user's optimal comfort temperature used in their model can be obtained.

2.4 Belief Networks

Belief networks (BNs) (sometimes also called Bayesian Networks or graphic models) are probabilistic models first introduced by Pearl (1986) that allow easy integration of prior knowledge. In case of the model presented in this work this means that findings from established comfort models can easily be incorporated to speed up the learning process. Probabilistic models are used to model complex relationships between numerous random variables. BNs are directed acyclic graphs in which nodes represent latent variables, unknown parameters or hypotheses. Edges represent the relationships between nodes. Unconnected nodes are conditionally independent. BNs are based on Bayesian statistics that in contrast to frequentist statistics model degrees of beliefs in certain states of the system. Bayesian statistics allow constant updating of beliefs as additional data is acquired. Beliefs are updated using Bayes' theorem:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{2.12}$$

where $P(A)$ is the *prior*, representing the initial belief in A and $P(A|B)$ is the *posterior* belief in A given B. $P(B|A)/P(B)$ is the *support* or *likelihood* for A given by B.

In systems with many latent variables, calculating exact conditional probabilities requires the evaluation of integrals of high dimensions. Computational complexity can be greatly reduced by considering dependencies of variables. Assume for example a system with four latent variables $\mathbf{x} = \{x_1, x_2, x_3, x_4\}$. The normal chain rule would result in the following equation:

$$P(x_1, x_2, x_3, x_4) = P(x_4|x_1, x_2, x_3)P(x_3|x_1, x_2)P(x_2|x_1)P(x_1) \tag{2.13}$$

Evaluating probabilities conditioned on many variables is costly. If some variables are known to be independent, this knowledge can be incorporated. Assume for example that x_4 is known to be independent from x_1 and x_2 and consider x_1 and x_2 to be completely independent. Using that knowledge, Equation 2.13 simplifies to:

$$P(x_1, x_2, x_3, x_4) = P(x_4|x_3)P(x_3|x_1, x_2)P(x_2)P(x_1) \tag{2.14}$$

In the simplified equation, x_4 for example only depends on one additional variable (x_3) which greatly simplifies the calculation.

However, this notation is very unintuitive and becomes harder to read as more variables are introduced. Graphical representations such as BNs can be used to express such equations in a more readable format. The resulting BN for Equation 2.14 is shown in Figure 2.2.

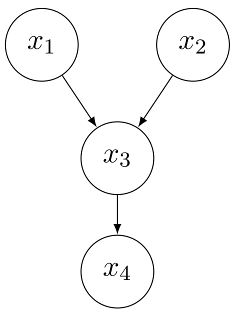

Figure 2.2: Simple example of a Belief Network

2.4.1 Inference

One of the main operations performed in BNs is to infer unobserved variables. In a thermal comfort model, such a variable could be the comfort temperature or the expected vote by the user on the temperature. Even with the reduced computational complexity achieved by incorporating dependencies between variables, the task of calculating exact probabilities is very complex (NP-hard). Therefore, numerous algorithms to perform inference more efficiently have been developed. There are two different kinds of algorithms: stochastic and deterministic algorithms. Stochastic algorithms use sampling methods to get good approximations of variables to be inferred. Deterministic algorithms use specific update schemes to infer variables. This work focuses on the expectation propagation (EP) algorithm by Minka (2001) due to its fast convergence, good support for continuous variables and accurate results.

EP is a message passing algorithm that extends the max-sum algorithm to support discrete and continuous variables at the same time. Such algorithms iteratively update their beliefs in variables by sending messages to other nodes with their current belief in other nodes. To make the computation tractable, the messages often only represent approximations of the actual distribution that minimize some divergence measure. As all message passing algorithms, EP is a subclass of power expectation propagation (power-EP). It differs from power-EP in the choice of divergence measure (Minka, 2005). While power-EP uses general α-divergence, EP minimizes the Kullback-Leibler divergence (KL divergence) — a special case of α-divergence where $\alpha = 1$ — between the actual probability function p and its approximation q.

EP is usually performed on factor graphs. Factor graphs are bipartite graphs where nodes can either be variables or factors. In a factor graph, factor nodes can only be connected to variable nodes and vice versa. Variable nodes represent random variables while factors model the relationship between different random variables. Factor graphs are used for message passing as they result in an intuitive interpretation of passing messages between factors or, in case of a fully factorised graph, messages between nodes and factors. The basic algorithm describes the joint probability distribution p as the product of all factors f:

$$p(\mathbf{x}) = \prod_a f_a(\mathbf{x}) \tag{2.15}$$

The structure of the network is defined by the set of variables each single factor depends on. With mixed distributions for different factors, calculating this product would result in hard to calculate integrals of high dimensionality. To overcome this problem, Minka proposes to approximate the factors with distributions of the same exponential family. The reason for constraining the approximations to members of an exponential family is closure under multiplication. The product of two distributions of the same exponential family results in another distribution of the same exponential family, simplifying the multiplication. Using approximations, the approximate joint probability q is described as follows:

$$q(\mathbf{x}) = \prod_a \tilde{f}_a(\mathbf{x}) \tag{2.16}$$

Finding a good approximation for each factor is a challenging task by itself as the approximation should be chosen in respect to the joint probability. In EP this happens under the premise that all other factors already have good approximations and therefore it is sufficient to use those as a reference for the full joint distribution. As a result, an approximation for a factor f_a can be found by minimising the divergence between $f_a(\mathbf{x})q^{\setminus a}(\mathbf{x})$ and $\tilde{f}_a(\mathbf{x})q^{\setminus a}(\mathbf{x})$

As a further simplification, the graph can be approximated by a fully factorised graph. In that case, approximations of factors are calculated as the product of all messages the factor sends to its neighbouring nodes:

$$\tilde{f}_a(\mathbf{x}) = \prod_i m_{a \to i}(x_i) \tag{2.17}$$

Those messages represent the conditional probability of the distribution of the factor conditioned on the expectations for all other neighbouring nodes. This conditional is again projected to an exponential family to allow easy multiplication. Expectations of nodes are calculated as the product of all incoming messages from neighbouring factors except for the targeted factor. A message to a factor f_a is calculated as follows:

$$m_{i \to a}(x_i) = \prod_{b \neq a} m_{b \to i}(x_i) \tag{2.18}$$

One can see that for calculating one type of message, its counterpart is needed, leading to a chicken-and-egg problem. This is overcome by initialising $m_{a \to i}(x_i)$ to some initial distributions. The algorithm then iterates over the different factors and recalculates its messages based on the updated messages of the other factors until all messages eventually converge. Once the algorithm converged, variables can easily be inferred by multiplying all incoming messages from its neighbouring factors.

2.4.2 Parameter Learning

Another big task performed on BNs is to learn their parameters. A common way to learn parameters is to model them as additional variables in the network and infer them based on observations of other variables in the network. Since parameters are treated as normal variables of the network, they can be inferred using the same methodologies explained in the previous section. Those additional nodes however introduce additional complexity, making the network harder to understand and generally slowing down operations on the network. It is therefore advised to only use this approach on smaller networks. For more complex networks, there are numerous other approaches that for example use expectation maximization (EM) (Bauer et al., 1997) or evolutionary algorithms (Myers et al., 1999). Such approaches do not require the model parameters to be modelled as additional variables.

In this work, model parameters are learned by modeling them as additional variables in the network and inferring them using the same inference algorithm also used for inferring normal variables of the network. Evaluations of the model showed that the network presented in this work is small enough to make this approach feasible. Even with the additional complexity, the scalability requirement is met.

2.5 Summary

The introduction of smart thermostats into smart homes is one of the next logical steps for the future to meet agreed energy saving goals. Existing approaches however fail to meet all requirements for a smart HVAC system. Most approaches neglect the user's preferences by either relying too strongly on manual configuration or over-simplifying their models. Some other approaches on the other hand focus too much on user comfort that they either become a burden to the user by requiring constant feedback or neglect other parts of a smart HVAC system such as load shifting capabilities.

One reason for that is that existing comfort models are often not applicable to domestic spaces due to their focus on the average satisfaction, or fail to incorporate important factors that determine thermal comfort found by previous studies. Comfort models that focus on individuals' preferences usually base their calculations of comfort temperatures solely on the current outside temperature. This model of adaptive thermal comfort is based on the assumption that not much about other variables such as humidity is known. However, as static thermal comfort models show, those variables have a major influence on thermal comfort as well and should therefore be also considered in adaptive models. In addition, only behavioural adaption of individuals is covered by this approach. Acclimatisation and psychological adaption are ignored completely.

While research in DSM is very active and thermal loads offer great potential, only few approaches deal with HVAC specifically. Of these approaches, even less are concerned with space heating and cooling. This is mainly due to the fact that maintaining user comfort can be a very challenging task using existing thermal comfort models due to the shortcomings mentioned before. Further, a number of existing approaches rely on a constant stream of feedback provided by the users which in practice is hard to maintain. Lastly, some approaches operate in a centralised manner, raising concerns about privacy of users and scalability of such approaches.

The next chapter introduces a novel approach to a personalised thermal comfort model that combines findings from static thermal comfort models with adaptive thermal comfort models and adds effective learning capabilities. As a result of using belief networks, generalised parameters for the model can be learned. The model is designed to work with a minimal set of sensors required. The outputs of the model are chosen in a way to integrate well with existing approaches to smart thermostats.

Chapter 3

The Bayesian Comfort Model

In this chapter we introduce a BCM and evaluate it against real world data. We start by introducing the model itself as a Bayesian network. We split the model into two parts: the general, core comfort model and the part of the model containing adaptive parts. We show how optimal comfort temperature, user votes and the user's comfort ranges can be calculated from the model and how inference and learning is performed. This is followed by an evaluation of the model. Apart from using real world data available from the ASHRAE RP-884 database, we also include data obtained from a deployment conducted by us. We explain the experimental setup of the deployment and how the data was processed. We evaluate the model's accuracy in comparison to the standard models defined in ASHRAE 55. Other models based on the user's optimal comfort temperature rather than the user's optimal vote were not evaluated since this temperature is not provided with the evaluation data set and therefore no training data for these models can be generated.

3.1 The 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 at any given time. The model can be trained to use any numeric scale for the comfort vote (for example the ASHRAE 7-point scale) as long as a zero on the scale expresses maximum comfort or total absence of discomfort. As opposed to existing models, our model combines the human-body centered approach of static models with the outdoor environment based approach of adaptive models. 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

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" that has been scaled with its γ_{var} counterpart.

Table 3.1: Notions in 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 sensitive a user is to deviations from the 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 influencing factors are eliminated.

For simplicity, the model has been broken down into two parts: the general comfort model (Figure 3.1) and its adaptive parts (Figure 3.4). 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 3.1.1. The adaptive parts of the model show the detailed calculation of the influence of adaptive measures and are explained in Section 3.1.2. Table 3.1 lists the different types of nodes and variables in the figures and explains their meanings.

3.1.1 The General Comfort Model

The general comfort model, shown as a factor graph in Figure 3.1, 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. 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}.

3.1.1.1 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 3.1.2. The unscaled humidity h describes the measured relative humidity inside the room.

According to existing thermal comfort models, there are other factors influencing a user's thermal comfort. In the static comfort model, metabolic rate, clothing level, operative

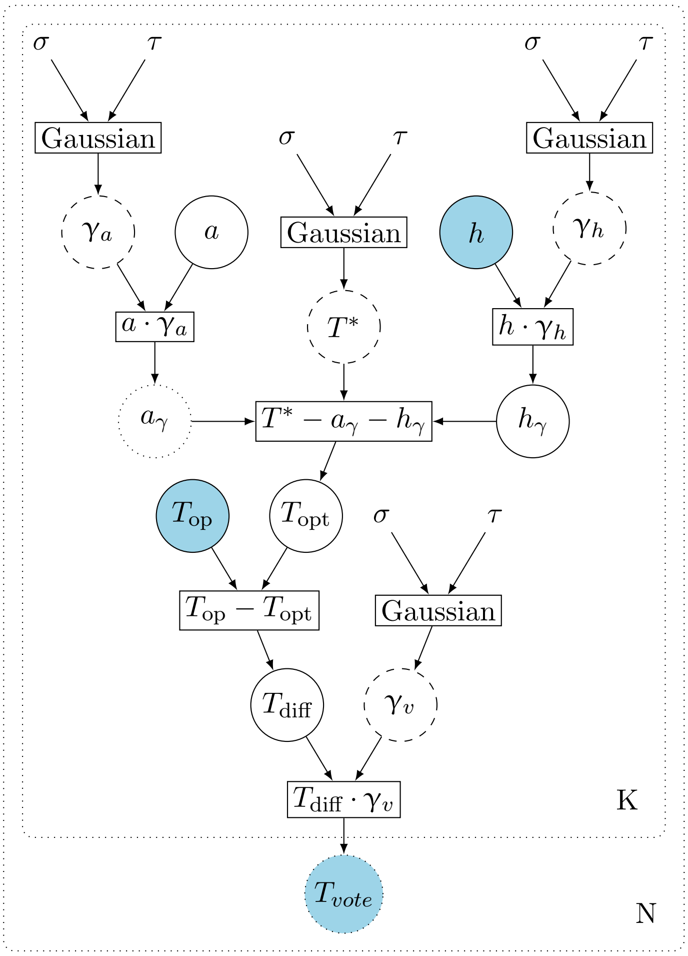

Figure 3.1: The general comfort model

temperature, humidity and air velocity (draught) are identified as influences (Fanger, 1970). Adaptive models name the outside temperature as the main influence (de Dear and Brager, 1998). In our model, metabolic rate, clothing level and air velocity are omitted when calculating the user's comfort temperature. Farhan et al. (2015) further states that age and gender play an important role in determining a person's thermal comfort level. However, since these variables are unlikely to change or change very slowly, they were not included directly into the model. These variables are therefore better suited for determining more accurate and individualised priors for the hyperparameters of the model.

For thermal comfort modelling, metabolic rate describes how much heat is produced through a person's metabolism. The clothing level defines the amount of insulation provided by the clothes a person is wearing. Together, those variables can be used to calculate the current heat loss of a person's body. However, these variables may be subject to frequent changes. Further, these variables are hard to measure or estimate. While modern portable fitness trackers measuring a user's heart frequency could be used to create a rough estimate of a user's metabolic rate, the clothing level is very hard to measure. Because of that, the clothing level was omitted in our model. For the metabolic

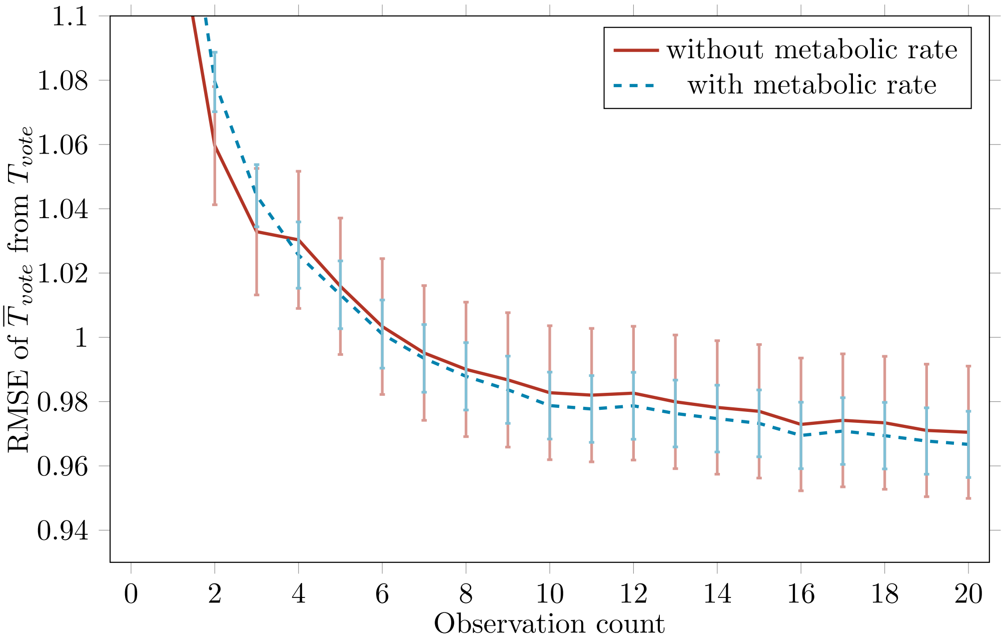

Figure 3.2: Accuracy of our model with and without accounting for the metabolic rate (with 2σ confidence intervals)

rate, we ran additional evaluations comparing how including the metabolic rate affects the accuracy of predictions of user's votes obtained from the model. For the evaluation, we used data about user's metabolic rates provided with the ASHRAE RP-884 data base. To include the metabolic rate into the model, we added a variable denoting a user's metabolic rate. The metabolic rate was scaled with a user-specific scaling factor learned from user feedback (similar to the γ factors in the model). The resulting user specific metabolic rate variable denoted by how many °C the users metabolism lowers the optimal comfort temperature. The calculation of the optimal comfort temperature was adjusted to also subtract he user specific metabolic rate from the base temperature T^*, meaning that higher metabolic rates would lower the user's optimal comfort temperature T_{opt}. Figure 3.2 shows the accuracy of the model's predictions of users' votes with and without metabolic rates. One can see that including the metabolic rate has no significant effect on the accuracy of the predictions. As a result, we omitted the metabolic rate in our model.

Moving air can have a cooling effect on the perceived temperature and can therefore have great influence on a person's thermal comfort. This cooling effect grows with the air velocity, which is why this variable is usually included into thermal comfort models. While anemometers can be used to measure air velocity, in practice this task turns out to be quite challenging as air velocities can vary greatly within a single room. To get a good idea of air movement in a room, several anemometers would have to be placed in a room, which would likely interfere with the occupant's normal life. Apart from measuring air velocities directly, one could try to calculate them. There are several possible sources of draught, such as open windows or doors or the stack effect described

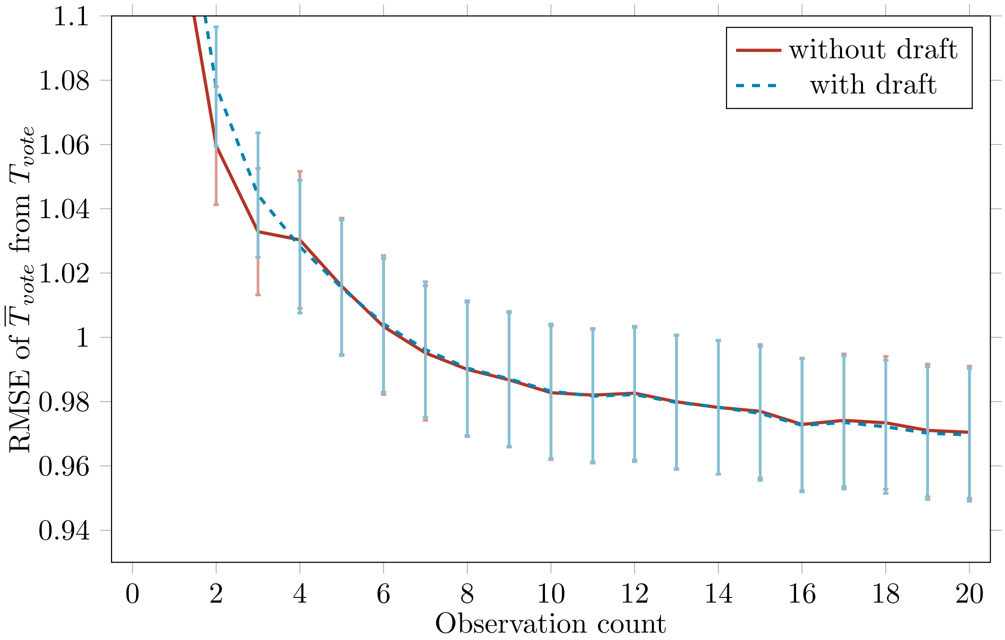

Figure 3.3: Accuracy of our model with and without accounting for draught (with 2σ confidence intervals)

in Section 2.1.1.4. To calculate the air flow caused by these sources, detailed knowledge of room geometry and other variables is required. Estimating these parameters or letting user's input them seems impractical as most of these parameters would require precise measurement and basic knowledge and understanding of thermodynamical terminology. In addition, as people move around in a room, previous assessments of thermal comfort based on air velocity are obsolete. Due to these problems, air velocity was excluded from our model. Experiments using data from the ASHRAE RP-884 data base, which contains data on air velocity, further showed that including these measurements had no significant effect on the predictive quality of our model. Figures 3.3 and 3.2 show that as soon as the model converges, the differences between a model accounting for draught or metabolic rate and a model not accounting for these factors are insignificant.

3.1.1.2 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, the range of comfortable temperatures is defined as all temperatures where the predicted mean vote (equivalent to T_{vote}) lies between -0.5 and 0.5. We can rearrange the calculations shown in Figure 3.1, to calculate the operative temperature T_{op} depending on the desired user's vote:

$$T_{\text{op}} = \frac{T_{\text{vote}}}{\gamma_v} + T_{\text{opt}} \tag{3.1}$$

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 $T_{\text{vote}} = 0.5$.

3.1.2 Adaptive Components

To cover a variety of adaptations by the user, the model includes a detailed section for adaptations (see Figure 3.4). 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., 2012b). 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). Further, as Auliciems (1981) suggest, repeated exposure to certain thermal environments, such as changing weather during different seasons, reduces the thermal sensitivity to these conditions. This is modelled by equation (3.2), which takes the current day of the year t_y as an argument:

$$a_s = \cos \frac{2\pi t_y}{365} \tag{3.2}$$

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

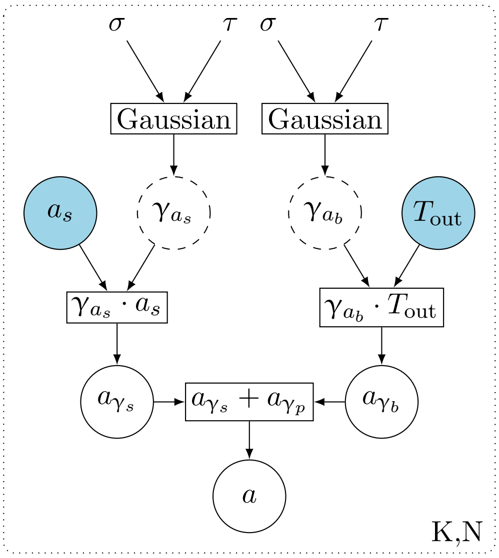

Figure 3.4: Adaptive part of the comfort model

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. This equation only represents a simple approximation to represent different seasons. In practice, more accurate equations (for example based on the hottest or coldest day of the year; deviation of the average temperature; etc.) could be used. However, since the evaluation data set does not span multiple seasons for single users, we did not focus on optimising this equation.

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}.

3.2 Evaluation

In this section, we evaluate the accuracy of predictions made by the Bayesian Comfort Model (BCM). 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. The thermal properties of a typical house and HVAC system are chosen based on data presented by Rogers et al. (2013).

3.2.1 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 HVAC 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. In the studies we chose, users were asked to provide consecutive 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 HVAC system for the user and get more accurate feedback, we further introduce and 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, our *desired change scale* asks 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 records for the locations of the studies were downloaded from Weather Underground[1]. 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 3.2. 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.

3.2.1.1 The ASHRAE data sets

Since many studies in the ASHRAE RP-884 data set were not concerned with consecutive feedback of occupants, we can only use a small subset of the overall data. Overall, there are three different studies that contain multiple data points per individual: a study conducted in different cities in Pakistan, a study conducted in Athens, Greece and a

[1]Weather Underground - `http://wunderground.com`

	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 3.2: Description of the different data sets. The Pakistan, Athens and Bay Area data sets were taken from the ASHRAE RP-884 database, the Southampton dataset was collected as part of this research and will be introduced in Section 3.2.1.2. *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).

study conducted in the Bay Area, California, USA. 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 when using an automated or semi-automated HVAC system. 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.

3.2.1.2 Thermal Comfort study in Southampton

To gather more data and be able to test our alternative feedback scale, in collaboration with Stephen Snow, we conducted a deployment at the University of Southampton, UK, to measure people's thermal comfort levels. This study was approved by the University's ethics department (ERGO number 20302). For this study, we developed posters inviting people to log how they feel about their thermal environment. On each poster, we attached a temperature logger measuring the temperature periodically, 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 is the temperature?", a large QR code and a URL address unique to each different poster (see Figure 3.5a). Scanning the QR code or entering the URL into a browser linked the user to a simple web-interface, shown in Figure 3.5b where they could log their thermal comfort.

(a) The poster encouraging to provide feedback.

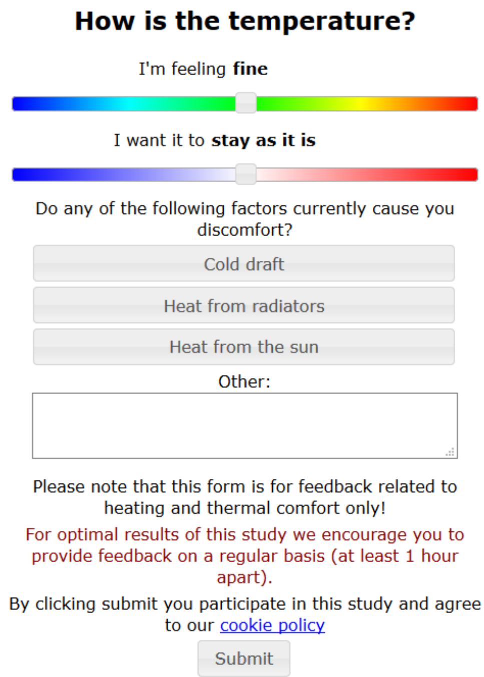

(b) The web interface used to provide feedback

Figure 3.5: The user interface of the study

The user interface of the feedback page featured two different sliders showing both feedback scales simultaneously. The first slider represented the ASHRAE 7-point scale, the second our own *desired change scale*. As the users moved the sliders, the label above the slider would adjust to the respective value on the scale. In case of the ASHRAE scale, these labels were: I'm feeling [very cold, cold, slightly cool, neutral/fine, slightly warm, warm, hot]. Under the assumption that some people might like it to be slightly warm or cold, we added the desired 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]. Similar to the ASHRAE scale, labels correspond to values between -3 and 3. Note that compared to the ASHRAE scale, the desired change scale is inverted, meaning that negative values roughly correspond to positive values on the ASHRAE scale and vice versa. In addition to the temperature feedback, we allowed users to list other things affecting their thermal comfort. Namely, those things were *cold draught*, *heat from radiators* and *heat from the sun*. In addition to that, users were also given a comment box where they could write what else caused them discomfort. Feedback was linked to individual users by assigning each user an anonymous, unique ID with their first feedback that was stored in a cookie in the browser used to provide feedback.

Leaving the decision when to report feedback to the participants likely biases the data towards only containing users that were unhappy with the thermal environment. As a result, user's feedback tends to be on the more extreme end (values lower than -0.5 or higher than 0.5). This was especially the case in the library deployment as we will discuss in the next two paragraphs. Due to this bias, the data collected in these studies should not be used to assess the general satisfaction of occupants with the thermal environment. However, since this bias is likely to occur in a real deployment of a smart thermostat (where people only report feedback or manually adjust the temperature when they are uncomfortable), this data provides a realistic scenario for the evaluation of our comfort model.

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. Since the interviews were mainly conducted by our collaborator Stephen Snow and focused mainly on the user experience rather than the actual comfort model the results are not included in this work. A more detailed description of the interview process and anlaysis of the results can be found in the original publication (Snow et al., 2017).

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. The deployment lasted for 5 weeks between May and June 2016. 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 ethical restrictions and 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

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

Table 3.3: Statistics of votes and users in the deployments

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 (desired change scale). Such situations should not occur in a system controlled using our HVAC algorithm and comfort model, since the model would likely 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 3.3 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.

3.2.2 Evaluation Results

We benchmark our model's prediction accuracy against the existing, standardised ASHRAE comfort models described in Section 2.2:

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

The approaches were compared based on the root mean square error (RMSE) of their predictions for T_{vote}. 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 3.3).

$$T_{\text{vote}} = 0.29(T_{\text{op}} - T_{\text{opt}}) \tag{3.3}$$

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.

Figure 3.7 shows the accuracy for predictions of T_{vote} achieved by our model, the PMV and the adaptive model. One can see that after 3-5 observations, our model starts outperforming existing approaches. 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 3.5 for the values chosen for the priors). Table 3.4 compares shows absolute and relative accuracy gains obtained by our model (after converging) compared to the PMV and adaptive model. One can see that apart from the Pakistan data set (see Figure 3.7a), our model achieves significant accuracy gains (7-55% smaller prediction error for T_{vote}) in comparison to the other approaches (Figures 3.7b and 3.7c). 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 environment, possibly due to the participants misunderstanding the trial protocol, which hinders the learning process. A closer look at individual results further revealed that the Pakistan data set contains single individuals for which the performance is extremely poor

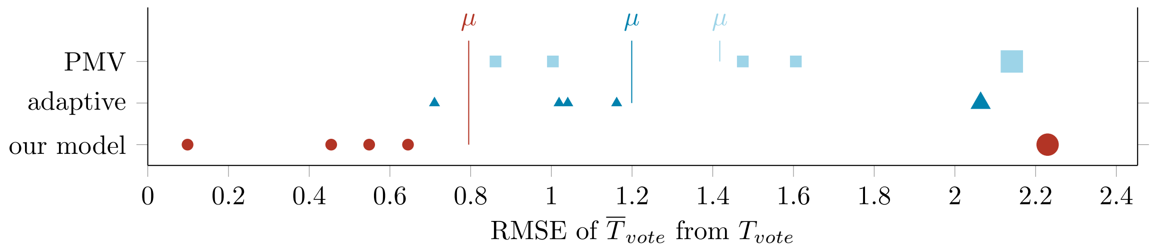

Figure 3.6: Prediction error for single individuals in Quetta (winter) dataset.

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 3.4: 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

(RMSE of $\overline{T}_{\text{vote}}$ from $T_{\text{vote}} > 2$), while for the remaining individuals the performance is good. Figure 3.6 shows single individuals of the Quetta data set. One can see that due to the generally lower numbers in our model, the outlier has a much greater effect on the average than with the other models. Looking at the actual data for such outliers revealed that in similar conditions (often only a couple of minutes apart), these individuals often only provided either extreme votes (-3 or 3) or zero votes. This indicates that there might have been a misunderstanding of the trial protocol or some techincal errors in the recording of users' votes.

Further, our model seems to benefit from the continuous scale used in the Bay Area data set. Despite having very limited data per participant (only 4-5 observations for the majority of participants), the model reaches a similar solution quality (see Table 3.4) to the other data sets which required 10 or more observations to converge to similar solution qualities (see Figure 3.7). In general, our model typically converges after 10 observations (see Figure 3.7c), but starts outperforming the other models after 2 - 4 observations.

Figure 3.8 shows the different histograms of the final model parameters for different subjects. Table 3.5 shows the means and variances of the parameters and the priors we chose for the variables. As convergence cannot be observed on the Bay Area and Southampton data sets and therefore no final parameters can be determined, data from these data sets has been excluded from the histograms.

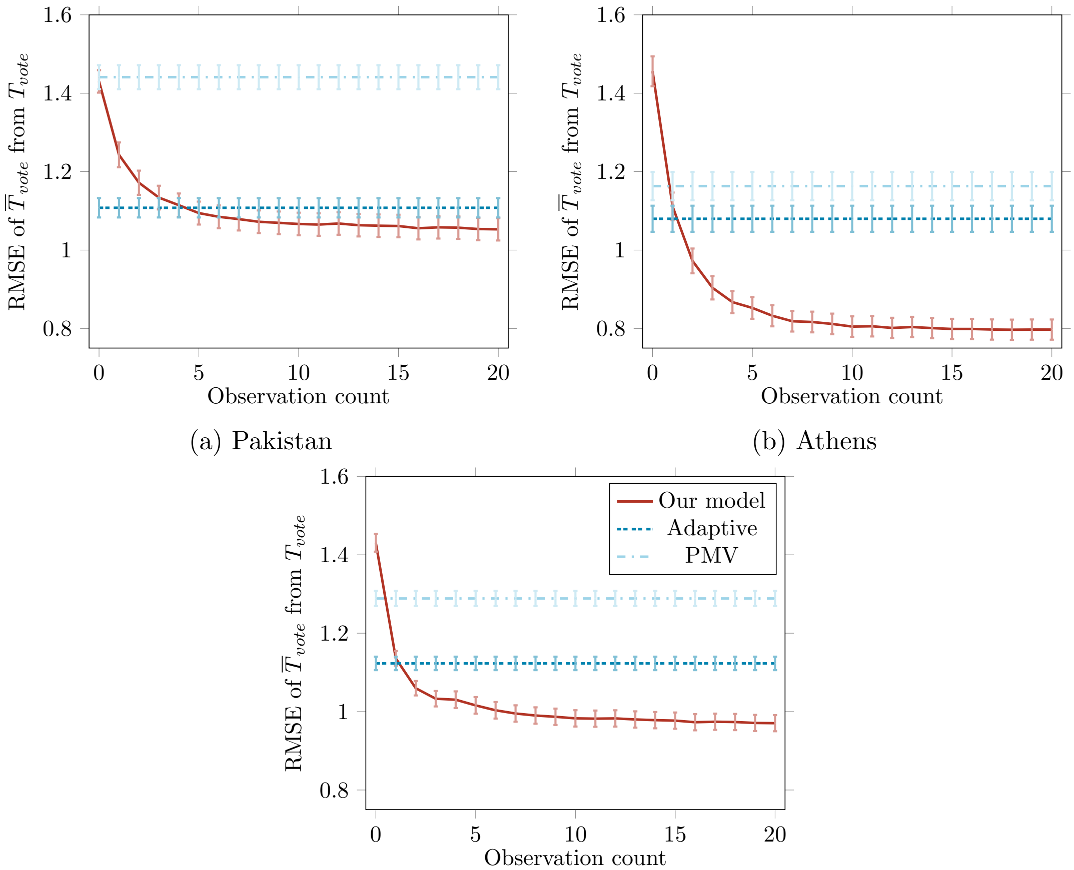

(a) Pakistan (b) Athens

(c) All data

Figure 3.7: RMSE of the predicted vote depending on observation count (with 2σ confidence interval)

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

Figure 3.8a shows that the base temperatures T^* ranges between values from 19°C to 25°C. The distribution of the learned parameters for T^* roughly resembles a Gaussian distribution with a mean around 22°C. Those results indicate that while there seems to be a general trend towards a neutral temperature of 22°C, many users still deviate from this, making it worth learning this parameter.

In contrast to this, the influence of adaptive measures γ_a does not seem to follow a general trend but is distributed more evenly with values ranging between 0 and 1 with a tendency to values of $\gamma_a > 0.5$ (see Figure 3.8b). This underlines the importance

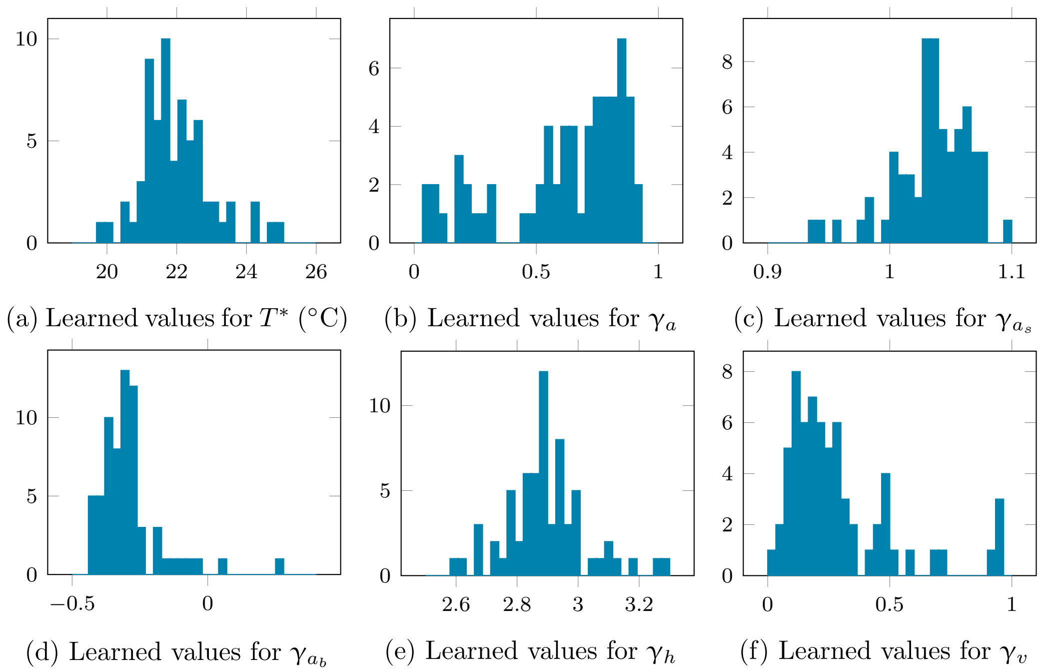

(a) Learned values for T^* (°C) (b) Learned values for γ_a (c) Learned values for γ_{a_s}

(d) Learned values for γ_{a_b} (e) Learned values for γ_h (f) Learned values for γ_v

Figure 3.8: Histograms of learned model parameters

of learning user specific values for this parameter as different users can be more or less effective in adapting to their thermal environment. The influence of seasonal adaptations γ_{a_s} does not vary a lot between different users in the data sets used. The learned parameters' values all lie in the narrow range between 0.95 and 1.1 (see Figure 3.8c). This may be due to the strict separation of seasons in the used data sets where for each individual only data for a single season is available. Behavioural adaptations γ_{a_b}, which were modelled with respect to the outside temperature T_{out}, seemed to play a more significant role with values mainly ranging between -0.4 and 0 (see Figure 3.8d). The mean, which is around -0.3, complies with values used in existing adaptive models, where the factor is often chosen to be between 0.3 and 0.33. The negative values in our model originate from the subtraction of the value from adaptive measures from the base temperature, which inverts the factor.

With values ranging from 2.6 to 3.3 (see Figure 3.8e), the influence of humidity γ_h also underlies significant variations among different individuals. The same holds for the users' thermal sensitivities γ_{v}. Overall the values range between values of 0 and 1 with an average of 0.29 (see Figure 3.8f). Using the equation $T_{\text{vote}} = T_{\text{diff}} \cdot \gamma_v$ with the average value, one can calculate that one step on the ASHRAE 7 point scale is on average equal to a deviation T_{diff} of about 3.5°C of the actual temperature T_{op} from the optimal comfort temperature T_{opt}. This is similar to the average of comfort bands for NV and HVAC buildings defined by de Dear and Brager (1998). Most thermal sensitivities however lie in the interval $[0.1, 0.4]$, which means that the temperature range for a single step on the 7 point mostly varies between 2.5°C and 10°C. This big range underlines the importance of learning individual thermal sensitivities.

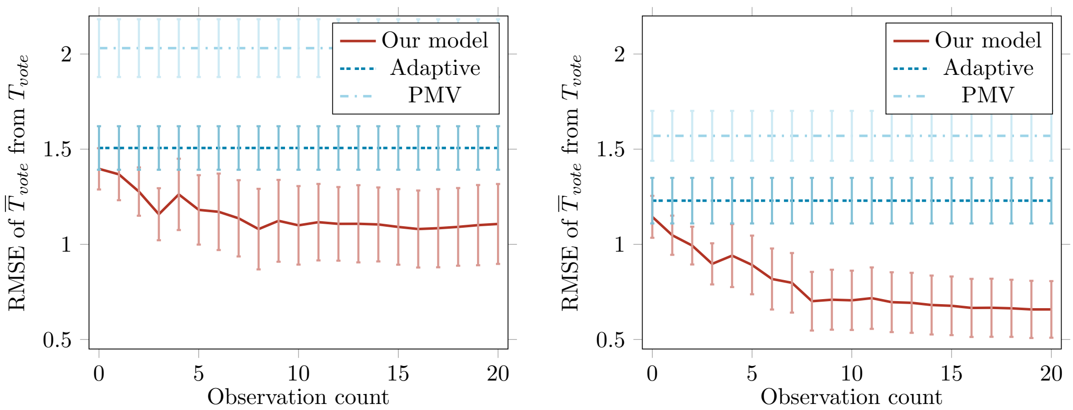

(a) ASHRAE scale

(b) Desired change scale

Figure 3.9: Comparison of prediction error on Southampton data set by feedback scales (with 2σ confidence interval)

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 3.9 shows the model's prediction accuracy using the standard ASHRAE 7-point scale (Figure 3.9a) as well as using the desired change scale described in Section 3.2.1.2 (Figure 3.9b). Despite the low number of participants and resulting large error, one can see that using our desired 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 3.2.1.2, 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.

3.3 Summary

This chapter introduced the Bayesian comfort model (BCM), a novel approach to a personalised thermal comfort model that combines static comfort models with adaptive thermal comfort models and extends them to model the users more accurately. The model was realised as a belief network, which allows to incorporate prior knowledge and enables the model to adapt efficiently to the users' preferences. The model presented consists of a core network with components directly related to the optimal comfort temperature, additional standardised components and additional individual-centric components. The standardised components represent factors defined by static thermal comfort models. Those factors are the operative temperature and humidity.

The individual-centric parts represent adaptive thermal comfort but instead of focusing just on the outside temperature, they also take seasonal adaptations and acclimatisation into account.

Through empirical evaluations we showed that the BCM is able to adapt quickly and efficiently to a user's preferences. Predictions of a user's votes on their thermal comfort are overall 13.2% - 25.8% more accurate than those of established thermal comfort models. We further proposed the *desired change scale*, an alternative scale for assessing users' thermal comfort levels to the established ASHRAE 7 point scale. Based on a novel study of individual thermal comfort in work environments, we showed that the desired change scale greatly improves the accuracy of the BCM.

The BCM introduced in this chapter has been published at the International Joint Conference on Artifcial Intelligence (IJCAI 2015) (Auffenberg et al., 2015a). A deeper discussion of the BCM, the *desired change scale* and our own study on users' thermal comfort has been accepted for publication in the next special issue on Urban Intelligence in ACM Transactions on Intelligent Systems and Technology (TIST) (Auffenberg et al., 2017, in press).

In the next chapter, we introduce an HVAC agent that utilises the BCM to minimise energy consumption of the HVAC system while keeping a comfortable environment for the occupants. To do so, we introduce a simple mechanism to aggregate the thermal comfort preferences of different users. Using simulations partially based on real data, we show the theoretical energy saving potential of the HVAC agent.

Chapter 4

A semi-autonomous HVAC agent

In this chapter we introduce a semi-autonomous HVAC agent that utilises the personalised thermal comfort model introduced in Chapter 3 to ensure a comfortable environment for the user. We show how this agent can utilise the comfort model to minimise energy consumption of the HVAC system without trading off user comfort. We first introduce a novel HVAC scheduling algorithm used by the agent to schedule the HVAC based on energy consumption and users' comfort ranges in Section 4.1. We model the problem as a linear program, allowing us to compute globally optimal HVAC schedules. In a realistic scenario, the algorithm will usually have to satisfy several occupants simultaneously. Thus, in Section 4.2, we introduce and compare two simple mechanisms to aggregate the comfort ranges of multiple users. In Section 4.3, using simulations based on real data, we evaluate how much energy the HVAC agent is able to save.

4.1 An Optimal HVAC Control Algorithm

We now introduce our optimal HVAC control algorithm that uses the Bayesian Comfort Model (BCM) introduced in Chapter 3 to minimise energy consumption of the HVAC system while maximising user comfort. 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.

4.1.1 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 affects the room temperature. We built upon the simple model by Rogers et al. (2013) introduced in Section 2.1.1.1. In this model, the HVAC system and room are described by a leakage rate $\Phi \in R^+$ (in 1/hr) and a heater output R (in °C/hr). The leakage rate Φ describes the rate at which the indoor temperature adjusts to the outdoor temperature T_{out}. The heater output describes by how many °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 as well, we added a third variable, the cooling rate C (in °C/hr), which describes by how many °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 computationally 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 0.3 in a time interval would mean that the heater/AC runs for 30% of the time and is switched off for 70% of the time, effectively meaning that in practice the heater AC is running at approximately 30% of its full capacity. As a result, the indoor temperature T_{air}^{t+1} in the house at time $t+1$ with a heater output R and cooler output C is calculated as follows:

$$T_{\text{air}}^{t+1} = T_{\text{air}}^{t} + \Big[\rho_r^t\ R - \rho_c^t\ C - \Phi\big(T_{\text{air}}^{t} - T_{\text{out}}^{t}\big)\Big]\Delta t \tag{4.1}$$

4.1.2 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 to ensure a comfortable environment. We use linear programming for this optimisation since it provides globally optimal results in polynomial time, at the cost of having to express the problem as a linear combination.

For the formulation as a linear program, we add constraint 5.1 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.

$$\text{lb} \leq T_{\text{op}} \leq \text{ub} \qquad \text{(constraint 4.2)}$$

By including user comfort in the form of a constraint rather than adding it as an objective on its own, the overall objective of the algorithm reduces to minimising only the energy consumption. Modelling user comfort as a constraint rather than a second optimisation objective prevents the algorithm from compromising comfort regardless of how much it could improve the objective. In addition, it greatly simplifies the optimisation problem itself. To actually calculate the user's discomfort, we would have to assess the indoor temperature at any point t in the optimisation. The indoor temperature at point t however depends on every previous indoor temperature, introducing recursive calculations into the objective function. This is avoided when only minimising energy consumption as the actual indoor temperature does not need to be evaluated at each time step.

To be able to minimise the energy consumption, the algorithm needs 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:

$$\gamma^t = (\rho_r^t\ \xi^r + \rho_c^t\ \xi^c)\ \Delta t \tag{4.3}$$

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 simultaneously and minimise the aggregated energy consumption. The resulting objective is shown in equation 4.4.

$$\min(\sum_{t=0}^{n} \gamma_t) \tag{4.4}$$

While the algorithm currently only factors in user comfort and energy savings, it can be easily extended to incorporate other measures such as space occupancy or variable energy prices. Space occupancy can easily by incorporated by removing the comfort constraints for the indoor temperature at times when the space is predicted to be unoccupied. To allow the algorithm 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 minimising cost is shown in equation 4.5.

$$\min(\sum_{t=0}^{n} \gamma_t \, p_t) \tag{4.5}$$

By design, the algorithm only considers a single comfort range, regardless of the number of occupants. In the next section, we introduce different mechanisms for aggregating multiple users' preferences into a single comfort range for the algorithm.

The full linear program looks as follows:

$$\min(\sum_{t=0}^{n} (\rho_r^t \; \xi^r + \rho_c^t \; \xi^c) \; \Delta t) \tag{4.6}$$

$$\forall t, \mathrm{lb} \leq T_{\mathrm{op}}^t \leq \mathrm{ub} \qquad \text{(constraint 4.7)}$$

where $T_{\mathrm{op}}^t = T_{\mathrm{air}}^t$ and T_{air}^t is calculated as shown in equation 4.1.

4.2 Aggregating preferences of 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 optimises for a single comfort range, we need to merge all occupants' comfort ranges. We evaluate two simple *comfort compromisers* that aggregate comfort ranges: the overlap comfort compromiser (OCC) and the average model parameter compromiser (AMPC).

The main aim of the OCC is to maximise user comfort, regardless of energy usage. The OCC uses the minimum of the upper bounds and the maximum of the lower bounds to get the area in which all comfort ranges overlap. If there is no clear overlap (lower bound > upper bound), the next closest value to the respective bounds are taken until the lower bound is less than the upper bound. If the overlap is smaller than 0.5°C, we expand the comfort range evenly on both sides to a width of 0.5°C. This is to aid the algorithm with the optimisation, as wider comfort ranges allow for more pre-heating or pre-cooling. Algorithm 1 shows how the lower bound lb and upper bound ub of the compromise comfort range are calculated with the OCC.

Algorithm 1 Pseudocode for calculating OCC

```
LB ← users.lowerBounds()
UB ← users.upperBounds()
lb = max LB
ub = min UB
while lb > ub do
    LB ← LB \ {lb}
    UB ← UB \ {ub}
    lb = max LB
    ub = min UB
end while
if ub − lb < 0.5 then
    lb = mean(lb, ub) − 0.25
    ub = mean(lb, ub) + 0.25
end if
```

The AMPC aggregates user parameters by creating a single comfort model for all users. In our simulations, this is done by creating a new instance of the BCM and setting the model parameters to the average of all occupants' model parameters. However, in practice, a similar result could be obtained by simply having all occupants provide feedback to a shared comfort model instead of a single comfort model per occupant. The AMPC further gives some degree of control over the energy usage vs comfort trade-off. When the thermal sensitivity/user vote factor γ_{v} is increased, the resulting comfort range gets narrower and user comfort increases. Decreasing the factor leads to greater energy savings at the cost of users' comfort levels.

4.3 Evaluation

To assess the theoretical energy savings achievable by our algorithm in combination with the comfort model and different comfort compromisers, we simulate households and their HVAC systems. We evaluate two main metrics: energy consumption and user discomfort. Energy consumption is calculated as the product of running times of the heater and air conditioning, multiplied by the maximum energy consumption of the heater and air conditioning (similar to equation 4.3). We chose values of $\xi^r = 8$kW and $\xi^c = 12.5$kW for the maximum energy consumption of the heater and air conditioning respectively. These values were chosen to match real energy consumption values observed in the Pecan Street data set, a data set that provides ousehold energy data disaggregated by various appliances. The algorithm was implemented using the CPLEX solver.

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 calculated as the aggregated discomfort of each user over each time step of the simulation. More specifically, user discomfort for a time interval t is calculated by multiplying the deviation of the indoor temperature from the user's comfort range with the length of the time interval Δt. We simulate 1000 households. We vary the amount of occupants per household between values of 1 occupant to 6 occupants per household. We add an extra value of 2.4 occupants per household, which is the average number of occupants per household in OECD countries[1]. The fractional value of 2.4 occupants per household is achieved by simulating 60% of the households with 2 occupants and the remaining 40% of households with 3 occupants.

Users are simulated using random configurations of our comfort model. In more detail, the parameters are drawn from the distributions shown in Table 3.5. To cover a variety of climates, houses were placed in different parts of the world and weather of that particular location was used. Specifically, households were simulated in the following cities: Austin, Texas, USA; Brussels, Belgium; Cape Town, South Africa; Moscow, Russia; San Francisco, CA, USA; Seattle, WA, USA; Shanghai, China; Sydney, Australia. The simulation covers the time period between 1st of January 2014 to the 31st of December 2015, covering all seasons.

We define two main benchmarks for our model: a fixed set point (FSP) HVAC system 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°C from the set point temperature. This is realised by defining this as the lower and upper bounds (lb and ub) of the comfort range (see Section 4.1.2). The FCR is mainly used to demonstrate the benefits of using a (wider) range of temperatures (such as comfort ranges) over a single set point temperature. Similar to the FSP, the FCR is implemented by setting the lower and upper temperature bounds of the algorithm to static values. For the FCR, we allow a 0.5°C deviation from the base temperature.

The set point temperature for an individual household is the daily average of the temperature maximising user comfort. That means that for each day, we calculate a single set point maximising the comfort of the household's users. This is done using a linear program similar to the one introduced earlier in this chapter. This set point is calculated for each day of the simulation. The household's set point for the actual evaluation is set to the average of the daily optimal set points. We use the average of each day since it is computationally not tractable to optimise the set point for the whole year.

4.3.1 Results

Figure 4.1 shows the evaluation of the HVAC agent's energy saving potential and comfort achievements with respect to the number of occupants in the building. One can see that

[1] Numbers taken from the OECD family database `http://www.oecd.org/els/family/database.htm`

the energy saving potential of the OCC highly depends on the number of occupants (see Figures 4.1a and 4.1b). The more occupants, the lower the energy savings. In case of cooling, with 4 or more occupants the OCC even uses more energy than the FCR. This can be explained with the shrinking size of the resulting comfort range. With more occupants, overlaps between all their comfort ranges are likely to be smaller, limiting the algorithm's ability to let the indoor temperature follow the outdoor temperature. The AMPC is more stable with regards to the number of occupants. This is because the average thermal sensitivity/vote factor, which directly determines the width of the comfort range is unlikely to vary significantly as more occupants are added. For two or more occupants, the AMPC therefore saves significantly more energy than the other approaches.

When comparing the discomfort with different comfort compromisers (see Figure 4.4b, one can see that compromisers utilising the comfort model (AMPC, OCC) can lead to significant reductions of discomfort when only a single occupant is present. Discomfort in this case is defined as the average hourly deviation of the indoor temperature from a user's comfort range. When only dealing with a single occupant, the hourly discomfort is reduced to almost zero. One can see that discomfort is slightly higher ($0.0161°C/h$) for a single occupant with the OCC (Figure 4.1c) as compared to the AMPC. This is most likely due to the artificial expansion of the comfort range to be at least 0.5°C wide. Starting from 2 or more occupants, the OCC however outperforms the AMPC. In a real setting this difference is likely too small to be noticeable by users. As the number of occupants grows, the AMPC even causes more discomfort than the static FCR and FSP. As we show in Section 4.3.2, this can be controlled by modifying the thermal sensitivity parameter.

4.3.2 Controlling comfort vs energy savings with the AMPC

One advantage of the AMPC is that it offers simple control over the prioritisation of comfort vs energy savings. This can be done by scaling the thermal sensitivity/vote factor γ_v of the comfort model. Reducing the factor results in wider comfort ranges which increase discomfort but reduce energy consumption. Increasing the factor results in narrower comfort ranges, decreasing discomfort at the cost of an increase in energy consumption. Figure 4.2 shows how the energy consumption changes with different scaling factors. We set the occupant count to a fixed value of 2.4, which represents the average number of occupants in OECD countries. One can see that especially for heating, the increase in energy consumption appears to flatten out for larger values.

Figure 4.3 compares the running times of heater and cooler of the AMPC compared to the other compromisers. One can see that the AMPC has consistently lower running times than the FSP and FCR. Compared to the OCC, heating and cooling running times are lower up to a scaling factor of 1.3.

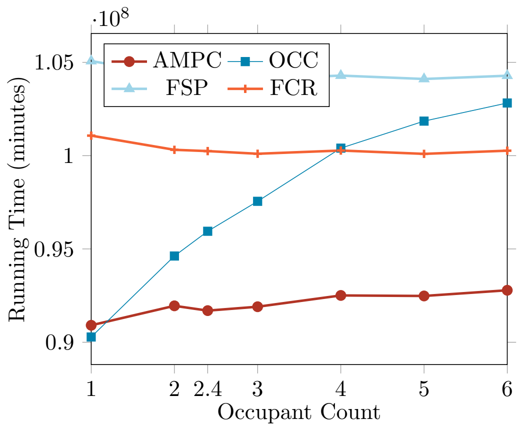

(a) Running time of the heater

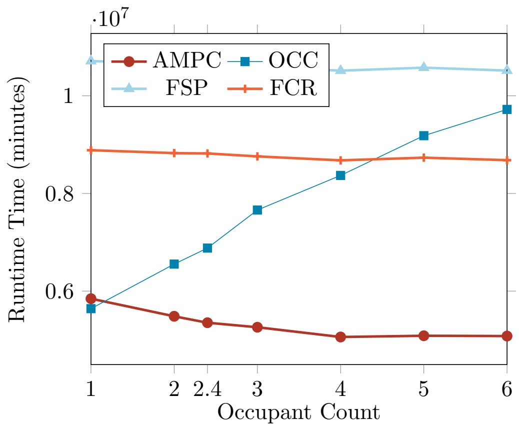

(b) Running time of the cooler

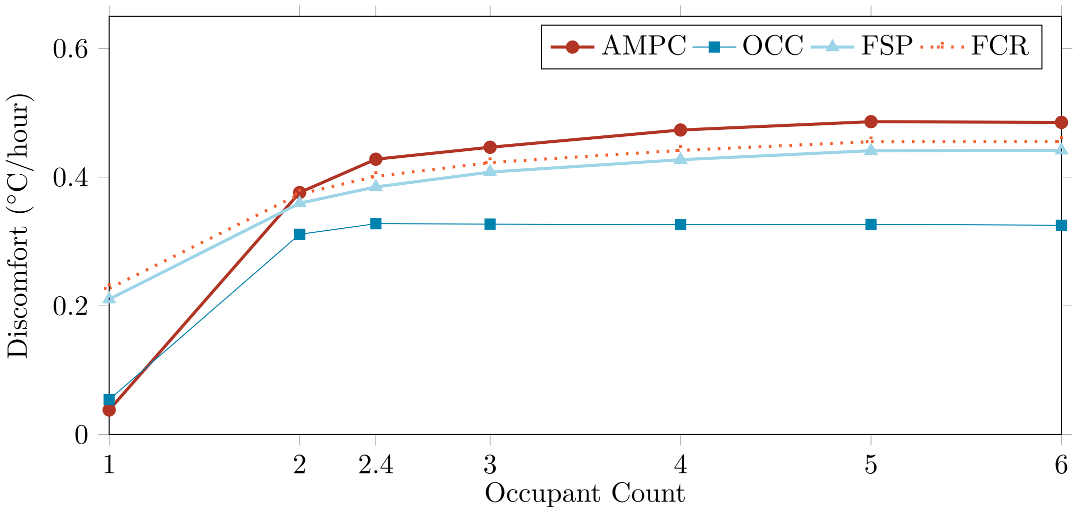

(c) Hourly discomfort per user per household

Figure 4.1: Evaluation of the HVAC agent's performance using different comfort compromisers

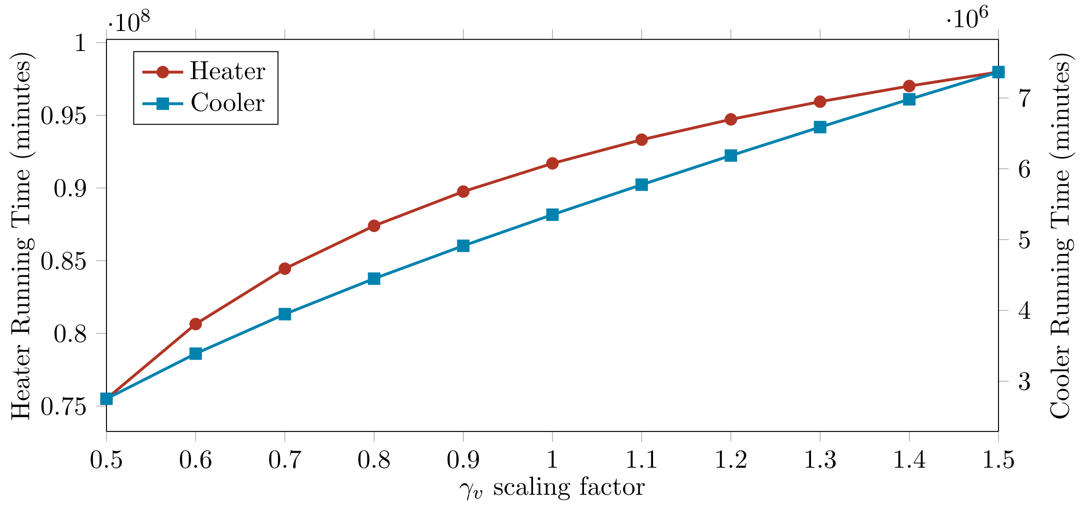

Figure 4.2: HVAC running times depending on the vote factor for γ_v (with 2.4 occupants)

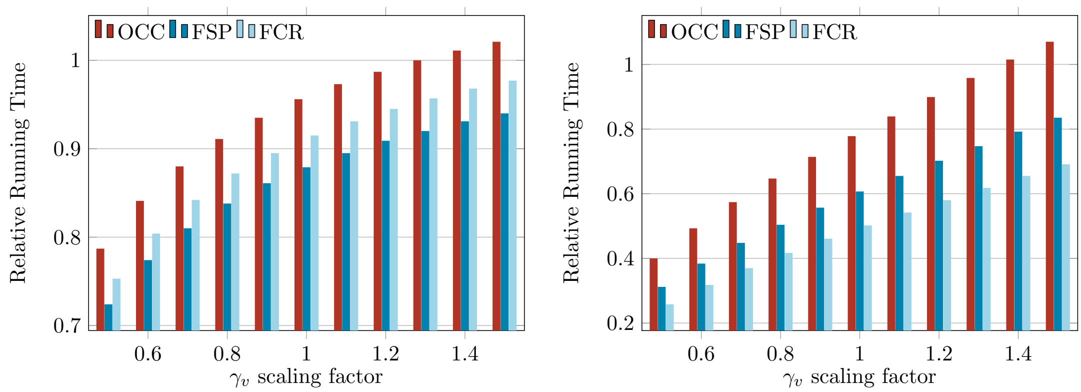

(a) Heater (b) Cooler

Figure 4.3: Heater and cooler running times of the AMPC relative to the other approaches

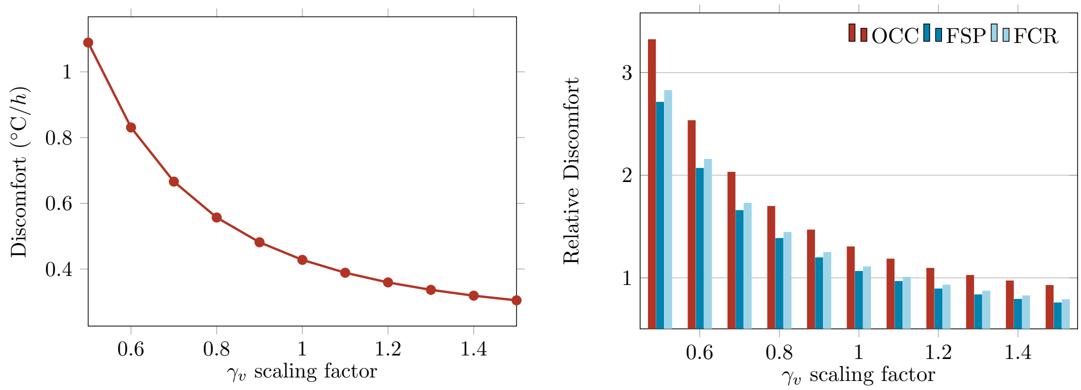

(a) Discomfort vs vote scaling (b) Discomfort relative to AMPC

Figure 4.4: Evaluation of AMPC with varying γ_v scaling factors

Figure 4.4a shows how the vote scaling influences the discomfort levels of the occupants. One can see that, as expected, higher vote scaling factors reduce discomfort. However, the discomfort seems to flat out with higher values. Figure 4.4b shows how the discomfort achieved with the AMPC compares with the other compromisers. One can see that for FSP and FCR, a factor of about 1 is needed to achieve similar discomfort levels. For the OCC, parity is achieved with a factor of about 1.3. With scaling factors greater than 1.3, the AMPC is able to outperform even the OCC but at the cost of higher heater and cooler running times (see Figure 4.3). Figure 4.5 shows how the scaling factor influences the discomfort with respect to the occupant count. One can see that lower scaling factors have a significant impact with three occupants and less. With three or more occupants, the impact of the scaling factor stays nearly constant.

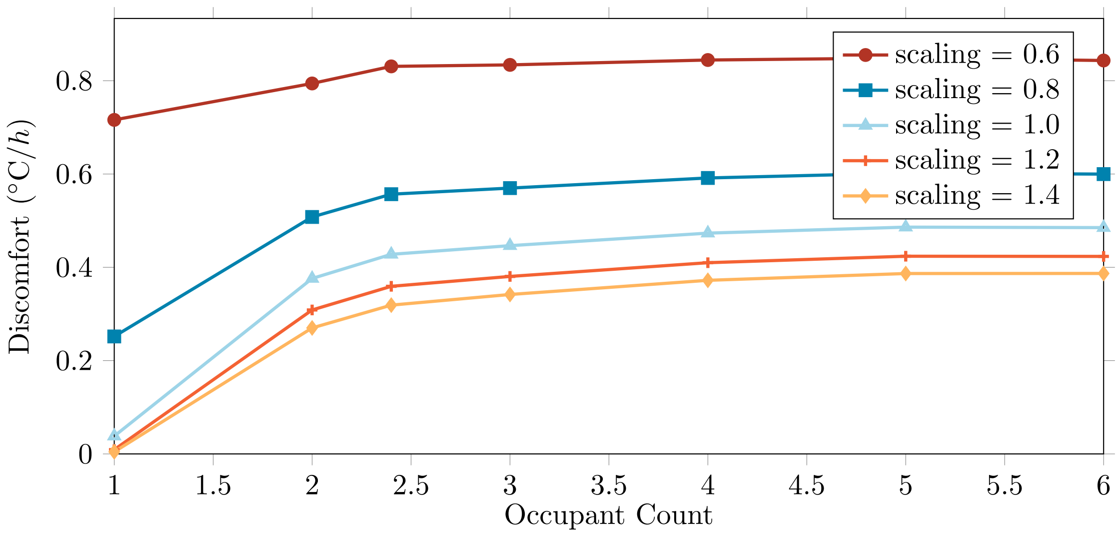

Figure 4.5: Vote scalings with respect to occupant counts

4.4 Summary

In this chapter we introduced an HVAC agent that utilises the BCM introduced in Chapter 3 and an optimal HVAC control algorithm to minimise energy consumption while retaining a comfortable environment for the occupants. The control algorithm models the problem of scheduling heating and cooling times as a linear problem. This allows the solver to find a globally optimal solution to the problem, meaning that the heating schedules obtained from the algorithm are optimal with respect to our objective function. To allow for efficient optimisation and the representation as a linear problem, we limit the objective to simply reducing energy consumption. Instead of adding a second objective for user comfort, we add a constraint limiting the indoor temperature to stay within the comfort range.

To be able to handle multiple occupants simultaneously, we introduce two comfort compromisers that aggregate different occupants' comfort preferences. The overlap comfort compromiser (OCC) aggregates different occupants' comfort ranges by finding the biggest overlap between different comfort ranges. The average model parameter compromiser (AMPC) utilises the BCM to aggregate preferences. Instead of only aggregating comfort ranges, it aggregates different occupants' comfort models directly into a single comfort model adjusted to the needs of all occupants. This is done by taking the average of all occupants' model parameters and feeding those into a new instance of the BCM. In a real-world setting where it might be difficult to collect and store each user's model parameters, a similar effect could be achieved by letting all occupants submit feedback to the same model instead of a single model for each occupant.

Based on simulations, we show the energy saving potential of the algorithm in combination with the BCM and the aforementioned comfort compromisers. We show that both comfort compromisers outperform traditional strategies such as the fixed set point (FSP) and fixed comfort range (FCR) with respect to both, energy saving and user comfort in most cases. In addition, we show how that when using the AMPC, scaling the vote factor γ_v of the BCM can be used to control the user comfort vs energy savings payoff.

The contents of this chapter have been accepted for publication in the next special issue on Urban Intelligence in ACM Transactions on Intelligent Systems and Technology (TIST) (Auffenberg et al., 2017, in press).

In the next chapter, we extend the HVAC control algorithm to incorporate DSM measures into its optimisation to shift heating and cooling times to approximate arbitrary load profiles.

Chapter 5

Comfort-Based Load Shifting

In this chapter, we extend the HVAC agent introduced in the previous Chapter 4 to incorporate DSM signals into its computation and perform load shifting based on these signals. We first introduce a centralised version of the algorithm that is able to calculate globally optimal results. Due to scalability issues with the centralised algorithm, we introduce a distributed version of the algorithm. Based on simulations using real data, we evaluate the load shifting potential of both versions of the algorithm. We do so with respect to energy usage, number of occupants and peak reduction capabilities.

5.1 Comfort-based load shifting

Our comfort-based load shifting approach consists of three separate steps, which are carried out on a daily basis:

1. **Collecting households' estimated profiles:** Using the BCM, households assess their comfort ranges, estimate their HVAC usage for the next day and send the resulting expected load profiles to the grid operator.

2. **Creating a target profile:** Using the profiles from step 1, the grid operator creates a target profile denoting a desirable profile meeting desirable characteristics like maximising renewable energy use or reducing peaks.

3. **Optimising HVAC schedules:** Using the target profile, either the grid operator (centralised version) or the households (decentralised version) run the algorithm presented in this work to calculate adjusted HVAC schedules. The schedules are optimised to approximate the target profile while still ensuring indoor temperatures to stay within comfort ranges calculated in step 1.

This work mainly covers steps 1 and 3. Since in practice the creation of target profiles will depend on a multitude of grid operator specific objectives, we only provide a simple approach for step 2 (see section 5.2). This approach calculates target profiles that aim to achieve a flat load profile which allows maximum utilisation of existing grid infrastructure. In the remainder of this section we introduce our comfort-based load shifting algorithm for HVAC loads. This algorithm can be used to estimate a globally optimal baseline HVAC load profile minimising energy consumption for each household (step 1) as well as to optimise the HVAC schedule to fit a given target profile (step 3).

We model the demand response problem as an MIQP, optimising the predicted interior temperature of each house to keep energy usage below a certain level and to minimise the quadratic deviation from the target profile provided by the energy supplier. Using the quadratic deviation ensures that going over the target is penalised equally to staying below. This allows target profiles to be tuned to, for example, maximise utilisation of renewable energy sources as opposed to only reducing peak energy consumption. Interior temperatures are modelled using the same, modified version of the model by Rogers et al. (2013) introduced in Section 4.1.1. Similar to the algorithm introduced in Section 4.1.2, the optimiser is constrained to keep the interior temperature within occupants' comfort ranges. Further, the maximum temperature change per time interval is limited by the maximum output power of the heating and AC systems used in a house.

In the following, we first describe the role of target profiles, which are central to our algorithm. We then formalise the algorithm as an MIQP and show how to run the algorithm in a distributed manner, locally for each household.

5.1.1 Target Profiles

Our algorithm uses target profiles instead of penalty profiles. As opposed to a penalty profile defining different penalties (such as higher energy prices) for each time slot, a target profile ω defines a specific target ω^t for each time slot t. These targets represent the most desirable actual load during a slot t, accounting for factors such as renewable energy production or grid and production capacities.

As opposed to penalty profiles, target profiles give more direct control over the resulting load profile. As pointed out by Ramchurn et al. (2011), penalty profiles have the risk of not scaling well. For example, if all households react similarly to the same pricing incentives, a rebound effect can develop, creating new peaks at times when energy is cheap (Palensky and Dietrich, 2011). Finding penalty profiles/pricing rates that prevent rebound effects can be challenging. Target profiles can prevent rebound effects since deviations from the target in either direction (going over or staying under) are penalised equally in the utility function. As a result, target profiles are easier to control since one can always expect the algorithm to get as close as possible to the given target.

Var	Description
γ_h^t	energy consumption of house h at time t
Γ^t	aggregated energy consumption at time t
G	aggregated energy consumption
$G_{\min}$	lowest possible aggregated energy usage
ω	target profile
ω^t	target load at time t
$\Delta\omega^t$	aggregated quadratic deviation at time slot t
Ω	aggregated quadratic deviation from target
l_t	aggregated load of target profile ω
l_h	aggregated load of house h
b	extra energy usage scaling factor
T_h^t	indoor temperature in house h at time t
lb_h^t	lower bound of comfort range in house h at time t
ub_h^t	upper bound of comfort range in house h at time t
T_{out}^t	outdoor temperature at time t
ρ_r^t	heating ratio at time t
ρ_c^t	cooling ratio at time t
Φ_h	leakage rate of house h
R_h	heater output of house h
C_h	cooler output of house h
ξ_h^r	heater power consumption
ξ_h^c	AC power consumption

Table 5.1: Nomenclature

Further, target profiles are more flexible with respect to the load. With target profiles, different types of loads or even appliances can be provided with different target profiles. This means that the profile can take limitations of load type or appliances into account. While target profiles for HVAC loads are relatively unrestricted due to the possibility of highly automating the process, target profiles for manually used appliances such as dishwashers, washing machines and tumble dryers can be tailored towards accommodating the users (for example by not asking for an increase during the night), increasing the likelihood of users actually following the profile.

The main downside of target profiles is that they usually require forward-planning (for example for the next 24 hours). This planning requires accurate forecasts of energy produced (especially by renewable sources) and expected demand. Inaccuracies in these forecasts can lead to undesireable load profiles when using target profiles. Other approaches, such as variable price rates with real-time pricing genreally allow for more rapid adjustments to current conditions.

5.1.2 Modelling HVAC-based DSM as an MIQP

In this section we introduce our DSM algorithm. We first describe how heating dynamics in the house are modelled. This is followed by a formalisation of the DSM problem as an MIQP. The formalisation shows a centralised version that provides an optimal solution to the DSM problem. We then show a mechanism to distribute this algorithm.

5.1.2.1 Centralised MIQP formulation

The main task of the DSM algorithm is to compute heating and cooling ratios for each house at every time step t so that the deviation from the target profile is minimised. To guarantee a comfortable environment for the user, the indoor temperature T is constrained to stay within the comfort range. The comfort range of a household h at time t is defined by a lower bound lb_h^t and upper bound ub_h^t, together describing the temperature band that is considered to be comfortable for the user. The indoor temperature T_h^t of a house h is constrained as follows:

$$\mathrm{lb}_h^t \leq T_h^t \leq \mathrm{ub}_h^t \qquad \text{(constraint 5.1)}$$

In addition to keeping the indoor temperature within the comfort range, the DSM algorithm should try to choose heating and cooling ratios that minimise the overall energy consumption. To calculate the actual energy usage of a house h for a time interval Δt ending at time t, the heating and cooling ratios get multiplied with the respective maximum energy usage of the heater or air conditioning unit. The maximum energy usage of the heater is described by ξ_h^r (in kW), the maximum energy usage of the air conditioning by ξ_h^c (in kW). The energy consumption between two time steps for house h is therefore calculated as follows:

$$\gamma_h^t = (\rho_r^t\ \xi_h^r + \rho_c^t\ \xi_h^c)\ \Delta t \tag{5.2}$$

The overall energy consumption Γ^t (in kWh) across all households for an interval ending at time step t is defined as the sum of all households' consumption-values:

$$\Gamma^t = \sum_h \gamma_h^t \tag{5.3}$$

To assess how close the resulting profile Γ is to the target ω, we calculate the deviation $\Delta\omega^t$ (in kWh) between the two profiles. The deviation at a time step t can be calculated as follows:

$$\Delta\omega^t = \left|\omega^t - \Gamma^t\right| \tag{5.4}$$

Using these calculations, we can formulate two objectives for the optimiser, the overall energy consumption and the overall deviation from the target profile. The overall energy consumption G (in kWh), shown in Equation 5.5, is the sum of the energy consumption of all time steps. We define the overall deviation from the target profile Ω (in $(\text{kWh})^2$) as the sum of the quadratic deviation from all time steps (see Equation 5.6). The quadratic deviation is used to penalise larger deviations from the target more and to simplify the optimisation. The resulting objective for the optimiser is shown in equation 5.7.

$$G = \sum_t \Gamma^t \tag{5.5}$$

$$\Omega = \sum_t \Delta\omega^{t2} \tag{5.6}$$

$$\min\bigl(G + \Omega\bigr) \tag{5.7}$$

There are a couple of issues with this objective function. The first issue is that these two objectives tend to be conflicting. Approximating a target profile will require some pre-heating or cooling which usually goes against the outdoor temperature. This however increases the temperature gradient between indoor and outdoor environment. Larger gradients mean increased heat transfer between outdoor and indoor environment, increasing the amount of energy required to maintain the indoor temperature.

Since the base profile minimizes energy consumption, adding the objective of minimising the target deviation will likely increase energy consumption. A simple way to control extra energy used for load shifting is to take a weighted sum of the objectives. This however can be troublesome due to the two objectives scaling differently. While the energy consumption (in kWh) scales linearly, the target deviation scales quadratically (in $(\text{kWh})^2$). This means that different scales (e.g. different number of households involved) require different weighting factors. Another issue is that the algorithm's ability to approximate the target profile is influenced by other factors, such as the outside temperature. This means that the same weighting of the objectives might yield different results depending on these factors. This bears the risk of unexpected, large increases in energy consumption during some days.

To address these issues, we modify the objective function to only minimise the deviation from the target profile. Energy consumption is controlled by adding a constraint restricting it to stay below a fixed threshold. The threshold is computed relative to the energy consumption of the profile with the lowest possible energy consumption G_{min}. This gives full control over how much extra energy can be used for load shifting, regardless of the conditions. The minimum energy profile G_{min} can be calculated by running a separate optimisation where the target deviation is removed from the objective function:

$$G_{\text{min}} = \min\bigl(G\bigr) \tag{5.8}$$

Using this, the load shifting optimisation problem can be simplified to reducing just the target deviation ($\min(\Omega)$) while energy consumption is controlled using constraint 5.9, which ensures that energy consumption stays below the specified bounds. The bound is specified by variable b (≥ 1), which defines how much energy can be used relative to $G_{\min}$. A value of $b = 1$ for example implies an equal amount of energy while $b = 1.05$ implies a 5% increase in the maximum consumption.

$$\sum_t \Gamma^t \leq b\ G_{\min} \qquad \text{(constraint 5.9)}$$

5.1.2.2 Distributed MIQP formulation

While the algorithm presented so far gives optimal results, its runtime and memory requirements scale exponentially with the number of households involved. To have a real impact on grid load, demand side management however has to be applied to a large number in the range or several thousands to hundreds of thousands of households, rendering the centralised algorithm infeasible. Along with scalability issues, a centralised algorithm also requires sensitive user information like individual comfort preferences to be collected. These issues can be addressed by distributing the algorithm to individual households.

A simple way of distributing the algorithm is to let households optimise their own load profiles individually. While only giving locally optimal results for each single household, the aggregated results of the distributed algorithm get close to the optimal solution as we will show in Section 5.2. To run the algorithm for a single household, the target profile needs to be scaled down to match the household's consumption. The target is scaled by the ratio between the aggregated predicted load of a household l_h and the aggregated load of the target l_t (see equation 5.10). The aggregated predicted load is obtained by multiplying the aggregated load of a house's minimum energy profile (similar to $G_{\min}$) with the extra energy usage factor b (see equation 5.11). The optimisation objective for a household h can be expressed as shown in equation 5.12) subject to (constraint 5.13).

$$l_t = \sum_t \omega^t \qquad (5.10)$$

$$l_h = b\ \min\Big(\sum_t \gamma_h^t\Big) \qquad (5.11)$$

$$\min\Bigg(\sum_t \Big(\frac{l_h}{l_t}\omega^t - \gamma_h^t\Big)^2\Bigg) \qquad (5.12)$$

$$\sum_t \gamma_h^t \leq b\ \min\Big(\sum_t \gamma_h^t\Big) \qquad \text{(constraint 5.13)}$$

5.2 Empirical Evaluation

To evaluate how our DSM algorithm performs in different areas, we empirically test it using simulations based on real data. We use the CPLEX optimiser for solving the

MIQP. Household energy usage data is taken from the Pecan Street Dataport. This data set contains usage data for several hundred households in the United States over a period of several years. The data features minutely intervals with energy usages being broken down by different appliances. Data is aggregated into intervals of 15 minutes length. We chose 15 minute intervals to allow for fine-grained control while taking into account slow room response to changes in HVAC control. In addition to aggregated intervals, the overall usage excluding the usages for heating and cooling is calculated.

Users' comfort ranges are modelled using the thermal comfort model described in Chapter 3. Different users are generated by drawing model parameters randomly from the distributions mentioned in Table 3.5. For most of our evaluations, we use the AMPC (see Section 4.2) with a vote scaling factor of 1.3. The scaling factor was chosen to produce comfort levels similar to the OCC.

5.2.1 Benchmark

We compare our DSM algorithm to two benchmarks. The first benchmark is the HVAC agent introduced in Chapter 4. Load profiles obtained from this agent minimise energy consumption without performing any load shifting. This benchmark is used to obtain the lower bound of energy consumption to evaluate the relationship between load shifting capabilities of our algorithm and extra energy usage. The second benchmark compares state-of-the-art approaches using a FSP (Hubert and Grijalva, 2011; Lu, 2012; Ramchurn et al., 2011). The set point is set to 23°C. We allow the temperature to deviate by 0.2°C in either direction. We assess how using the BCM changes the energy consumption and its effect on the peak reduction capabilities of our algorithm.

5.2.2 Simulation Setup

Our algorithm and the effects of different comfort models, occupant counts and energy consumptions are evaluated by simulating individual households and their HVAC systems. A household h is defined by the parameters listed in Table 5.2, with values randomly drawn from the uniform distributions shown in Table 5.2. These values were chosen so that simulated HVAC energy consumptions match up with real HVAC energy consumption data found in the Pecan Street data set.

The simulation covers the time span from January 2014 until the end of December 2014 to include all seasons. The time resolution is 15 minutes. The algorithm runs on a daily basis, each time planning ahead the next 24 hours. This means that each day, per household 96 different set point temperatures are calculated. Since the majority of households in the Pecan Street data set are located in Austin, TX, we base our simulations on this area. To get realistic consumption data, we only consider households

Variable	meaning	distribution
R_h	heater output	$U([1.75, 2.25])$
C_h	cooler output	$U([1.25, 1.75])$
Φ	leakage rate	$U([0.02, 0.03])$
ξ_h^r	heater power (kW)	$U([6, 8])$
ξ_h^c	AC power (kw)	$U([9.5, 15.5])$

Table 5.2: Thermal properties of a house

with electric HVAC systems, which leaves a total of 225 different households. Weather data for this area was downloaded from Weather Underground[1].

To create realistic scenarios for the simulation, target profiles are created based on real usage data of the simulated households. The target profile for HVAC loads is chosen to complement the profile of non-HVAC loads so that the combined profile approximates a flat line. This is done by creating a flat profile where each value is equal to the maximum energy usage of the real usage data. We then subtract the real usage data from this flat profile. The resulting profile is then scaled so that the integral of this profile matches the expected overall energy consumption of the HVAC systems in the simulation.

5.2.3 Evaluation Results

We evaluate the performance of the algorithm in a number of ways. First, we investigate how the extra energy usage factor b influences the algorithm's ability to approximate the target profile. We then evaluate the influence of the number of occupants per household on the algorithm's ability to reduce the target deviation. This is followed by an evaluation of the algorithm's ability to reduce peaks in demand. We then assess seasonal effects on the algorithm's load shifting potential, followed by a brief example showing how to use target profiles to improve the utilisation of renewable energy sources.

5.2.3.1 Extra Energy Usage

The two objectives of minimising energy consumption and minimising the target deviation are conflicting as pre-heating or cooling the house usually increases the temperature gradient between inside and outside, resulting in higher losses to the environment. Consequentially, to be able to perform load shifting, more energy needs to be used. It is likely that at one point the drawbacks of increasing the energy consumption will outweigh the benefits of shifting demand. To be able to decide on a suitable upper bound for the amount of extra energy used to shift loads, knowledge about the relationship between the two objectives is required.

[1]Weather Underground `http://www.wunderground.com/`

Figure 5.1 shows the algorithm's load shifting capability with respect to the extra energy usage factor b. The minimum energy baseline ($b = 1$) is equal for the distributed and centralised case as the energy consumption objective is additive with respect to individual households, resulting in identical profiles between the centralised and distributed version. This means that results with respect to the extra energy usage factor b are directly comparable between the centralised and distributed version.

The results show that allowing the algorithm to consume more energy greatly increases its ability to reduce the target deviation. The improvement rate however reduces with growing values for b. The algorithm further benefits from using the more sophisticated BCM as opposed to the simple FSP. Using the AMPC, on average the distributed version leads to 39.9% (up to 53% in some cases) closer approximations of the target profile, the OCC to 29.75% (up to 40% in some cases) respectively. The AMPC in general performs slightly better than the OCC. We believe that this might be due the generally smoother changes of comfort ranges. In some cases, for example when comfort ranges stop to overlap, the comfort ranges between two consecutive time steps can change significantly. Such jumps in the comfort range can limit the algorithm's ability to pre-heat or cool the house.

One can see that the distributed version in general achieves slightly worse results than the centralised version. The centralised version using the AMPC gets 12.8% (up to 15.4% in some cases) closer to the target than the distributed version. For the FSP, the average gap between centralised and distributed version is a bit smaller with the centralised version getting 8.9% closer to the target than the distributed version. We believe that this is because the central algorithm cannot exploit different comfort ranges. On a hot day, for example, the centralised algorithm can exploit the fact that some houses can be pre-cooled further than others. With the FSP, the main difference between households are the thermal properties of the house. However, one can see that using the FSP, the centralised version benefits more from larger values for b. With a value of $b = 1.2$, the centralised version gets 19.75% closer to the target than the distributed version.

5.2.3.2 Occupant Count

As opposed to algorithms using a FSP, algorithms using personalised thermal comfort models are affected by the number of occupants in a household. The more occupants, the harder it becomes to find a compromise between keeping users comfortable, saving energy and shifting demand. We evaluate what impact the amount of occupants per household has on the algorithm's load shifting ability depending on the comfort model and compromiser used. For the evaluation, we vary the occupant count between 1 and 5 in steps of 0.5. Half counts are achieved by letting half of the households have x occupants and the other half $x + 1$. For example, to obtain an occupant count of 2.5, half the households will be assigned 2 occupants whereas the other half will be

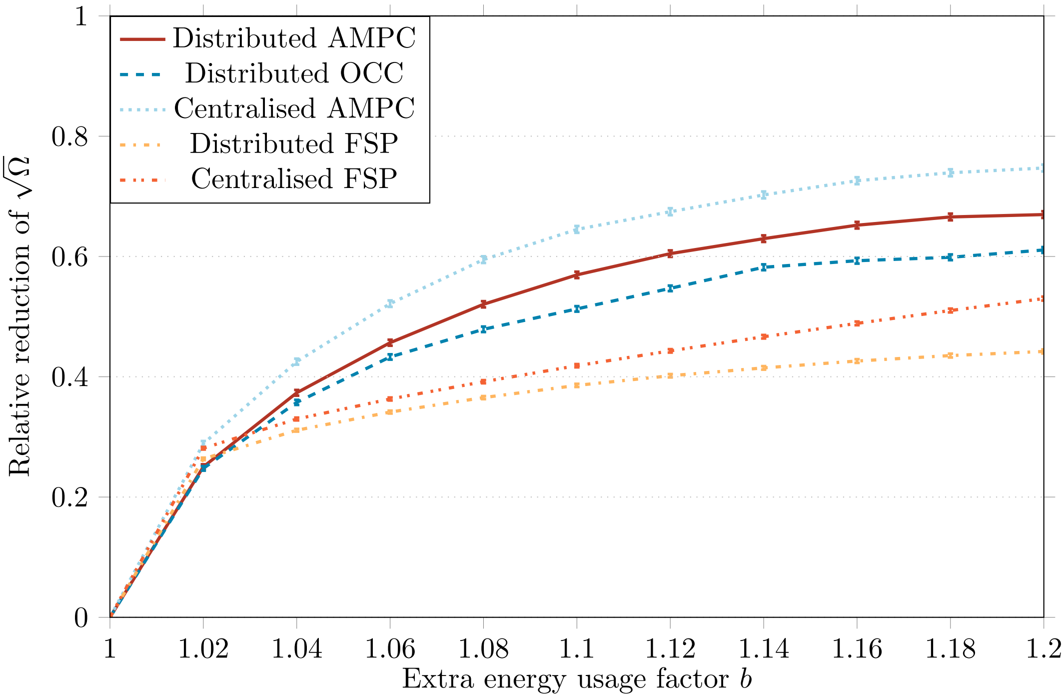

Figure 5.1: Reduction of target deviation compared to energy optimal profile (with 2σ confidence intervals)

assigned 3. As the energy consumption is based on the minimum energy consumption, energy consumption scales similarly as for the HVAC agent introduced in Chapter 4. See Section 4.3.1 for an evaluation of energy consumption with respect to the comfort compromiser used and occupant count.

As shown in Figure 5.2, the number of occupants has varying effects the algorithm's ability to reduce the target deviation depending on which comfort model or comfort compromiser is used. In case of the AMPC, the amount of occupants seems to have no significant effect on the algorithm's load shifting ability. In contrast, when using the OCC the algorithm's performance suffers the more occupants are present. We believe that this is due to the same effect described in the previous section. As more occupants need to be satisfied at the same time, jumps in the comfort ranges resulting from the OCC become more likely.

5.2.3.3 Peak Reduction

To assess the algorithm's peak reduction capability, we measure by how much the time the load goes over a certain threshold can be reduced. We define peaks as intervals with unusually high loads compared to other intervals during the same day. This definitions allows to class peaks by percentiles. For example, 99th percentile peaks describe the 1% of intervals with the highest load. A 10% reduction in the 99th percentile would mean that the aggregated load of intervals in the 99th percentile was reduced by 10%.

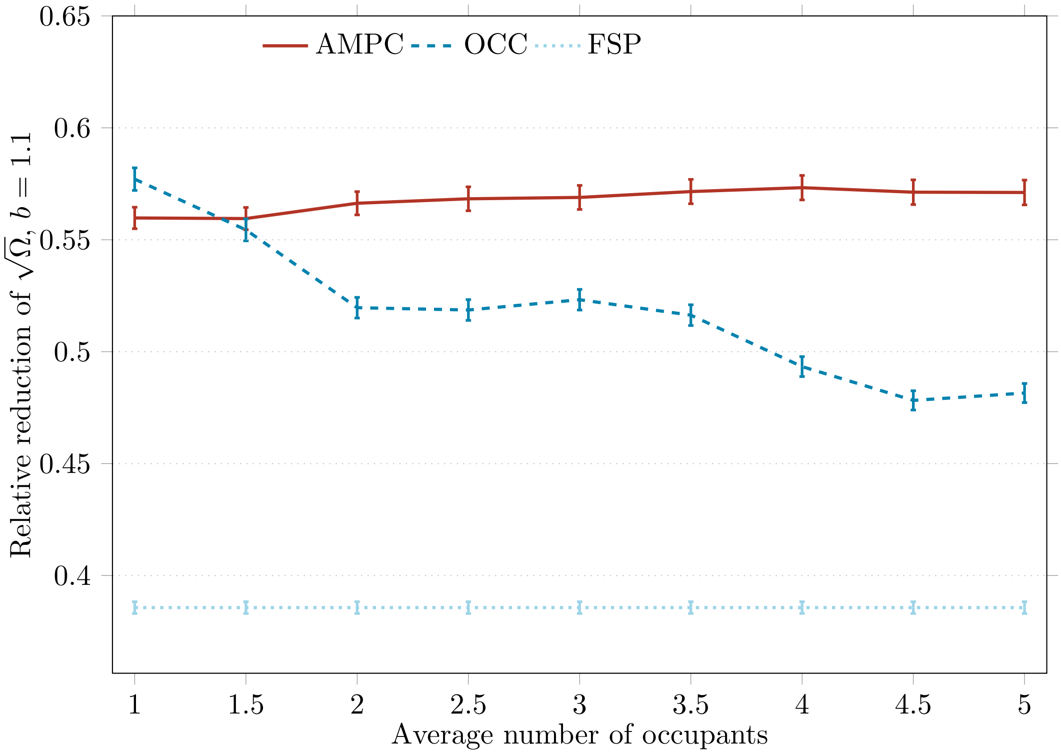

Figure 5.2: Influence of occupant counts on the load shifting potential for $b = 1.1$ (with 2σ confidence intervals)

We set the threshold to different percentiles of the actual load values. Actual load values are non-HVAC loads plus HVAC loads resulting from the minimum energy profile $G_{\min}$ without load shifting. We compare the peak reduction capabilities of the distributed algorithm using AMPC with a vote scaling of 1.3 and a FSP. Figures 5.3a and 5.3b show that using the more sophisticated comfort model with the AMPC allows for much greater peak reductions. The increased peak reduction capability is mainly due to the larger comfort ranges and the resulting larger load shifting capabilities. One can see that while greatly reducing the 99th and 95th percentiles, reductions decrease in lower percentiles. When using the AMPC, usage in the 70th percentile actually increases, indicating that the algorithm's capability to reduce peaks is exhausted at this point as loads from higher percentiles are shifted into this area.

When comparing the results of the AMPC and FSP one can see that the extra energy usage factor b has a greater influence on the algorithm's peak reduction performance when using the AMPC. This dependence however scales differently in different percentiles. With higher percentiles the effect is limited only to the lower value range of b ($1 \leq b \leq 1.05$). For lower percentiles, this value range tends to be larger. This is because for the higher percentiles, only relatively small loads during short time periods need to be shifted.

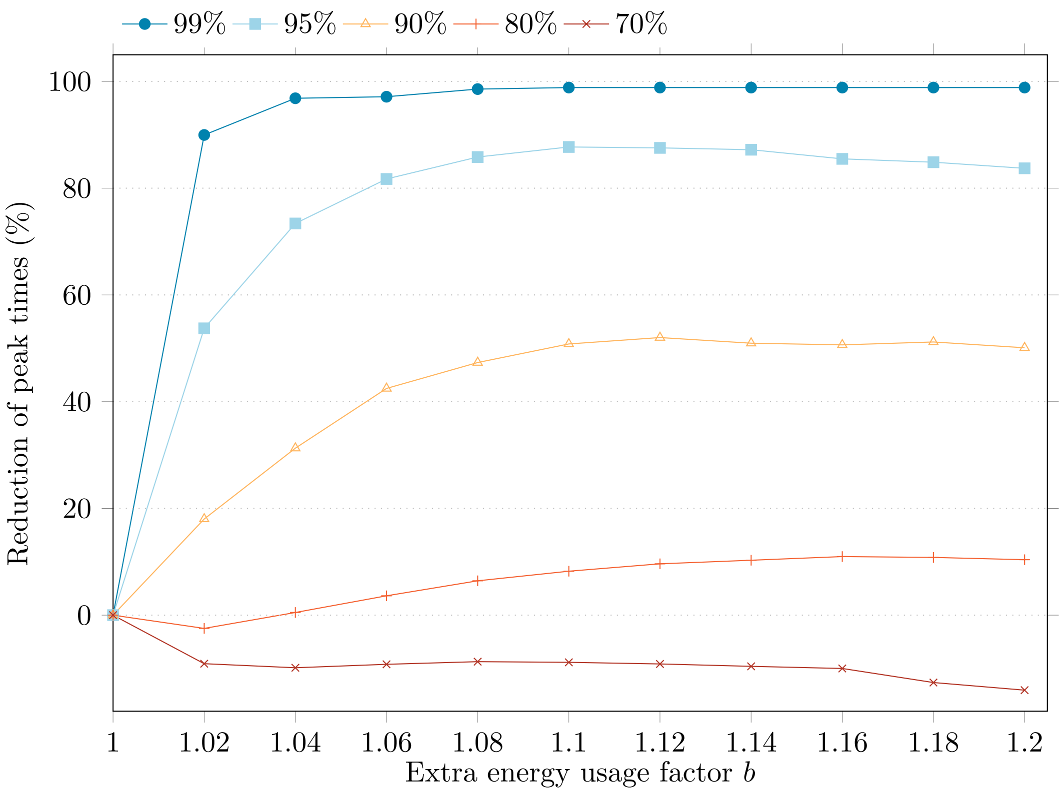

(a) Comfort Model

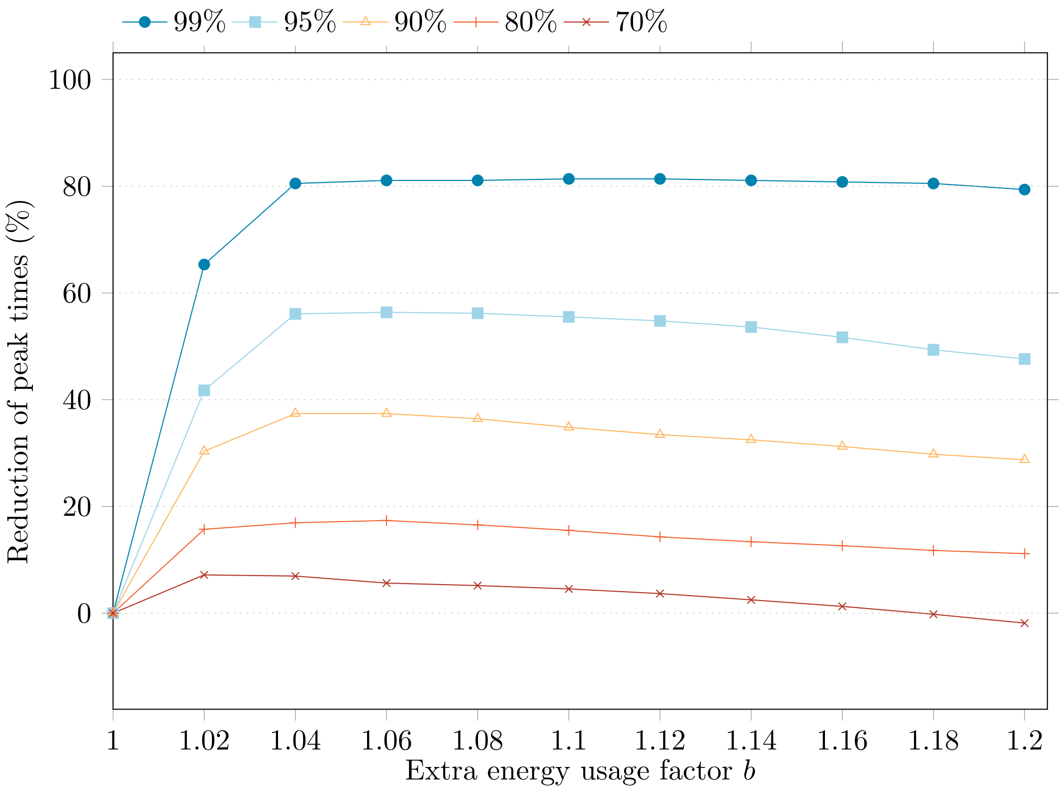

(b) FSP

Figure 5.3: Relative peak reduction for different perecentiles

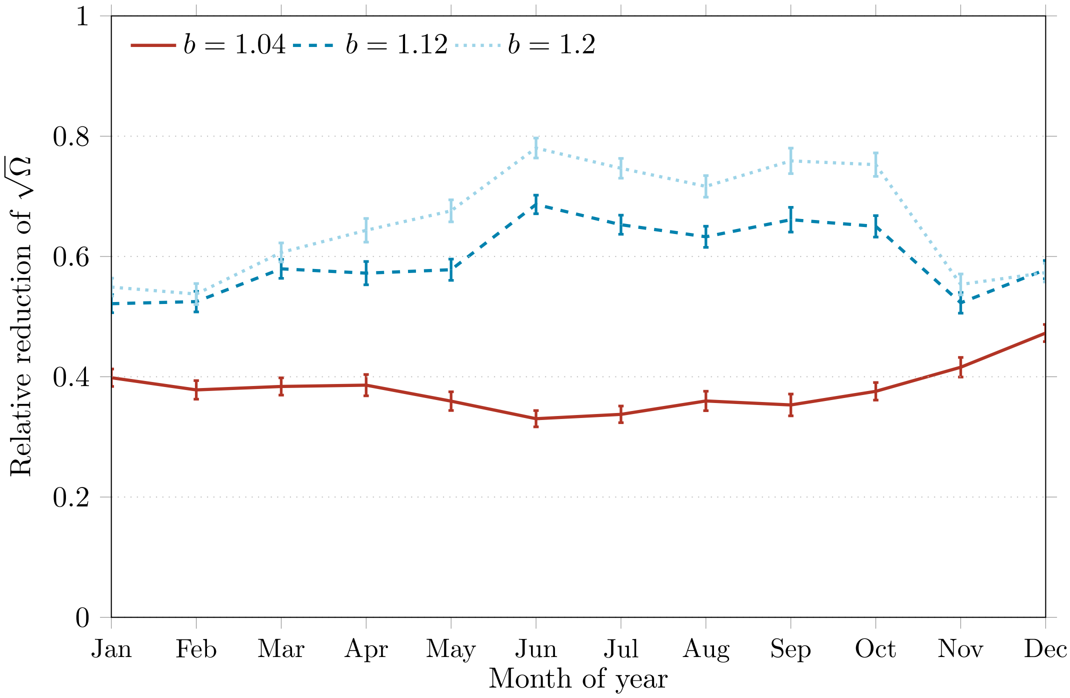

Figure 5.4: Reduction of target deviation by month using the AMPC (with 2σ confidence intervals)

5.2.3.4 Seasonal effects on load shifting potential

Since space heating and cooling is highly dependent on the outside temperature, we expect the algorithm's performance to vary throughout the year. To assess these variations, we measure the algorithm's performance separately for each month. In addition, we compare three different upper energy bounds. We compare a rather small value of 1.04, a medium value of 1.12 and a large value of 1.2. The algorithm was configured to use the AMPC with a vote scaling of 1.3. The average occupant count was set to 2.5.

The results are shown in Figure 5.4. With a small upper energy bound, the algorithm's performance is limited during the summer months. Increasing the upper energy bound however reverses this effect and the algorithm performs much better during summer than during winter. This stark difference between summer and winter can be explained by the location of our simulations. In the generally warm climate in Austin, Texas, the HVAC system plays a much more important role during the hot summer months where the AC is running most of the time than during winter when outside temperatures are fairly moderate. The larger temperature gradient between inside and outside during summer increases the energy required to pre-heat and cool the house. This extra energy however, can be spread out more easily, improving load shifting capabilities. This suggests that in colder climates, the algorithm will benefit more from higher energy bounds during winter.

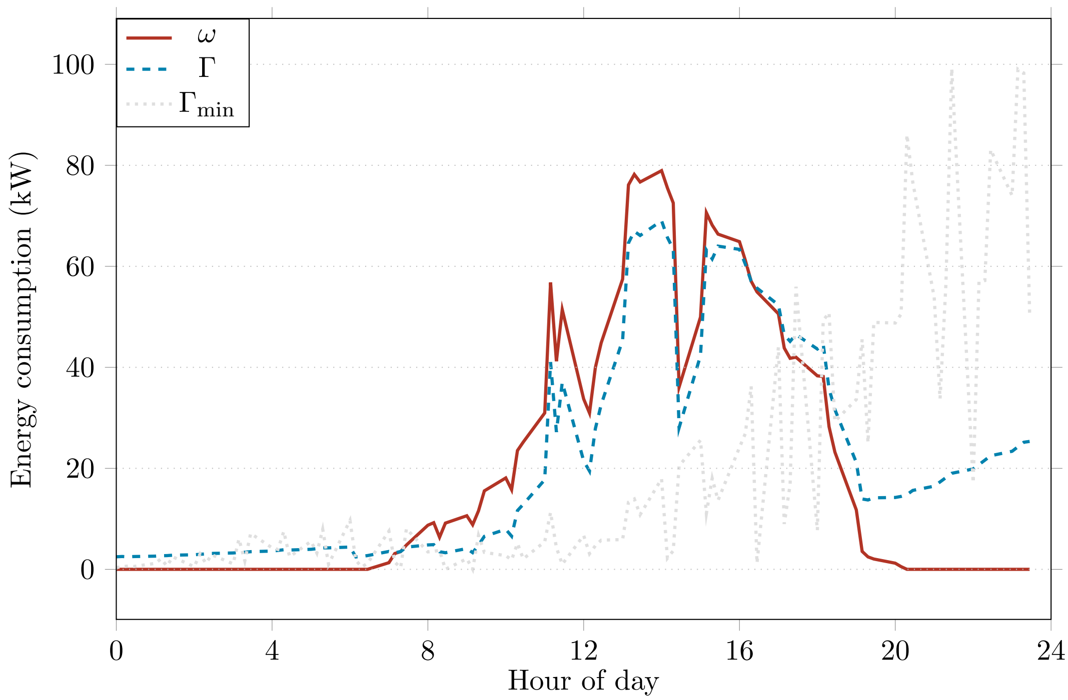

Figure 5.5: Example profile with solar radiation target

5.2.3.5 Incorporating Renewable Energy Sources

In this section we illustrate how to maximise utilisation of renewable energy sources by using target profiles and our DSM algorithm. To do so, we create a target profile that incorporates the expected production from renewable energy sources for the next day. In the example shown in Figure 5.5 we use solar radiation measurements obtained from Weather Underground for the 1st of August 2014 in Austin, Texas to create the target profile. As this is an illustrative example, for simplicity's sake, we assume that solar radiation and production from solar panels are directly correlated. The target profile therefore represents the solar radiation profile, scaled up to match the integral of $b\,G_{\text{min}}$ scaled by an upper energy bound of $b = 1.1$.

One can see that the algorithm successfully approximates the target profile, maximising the utilisation of energy from renewable energy sources. Most of the loads from the evening are shifted into the period between 12pm and 7pm. The algorithm however successfully preserves the dips around 12pm and 3pm. While most of the loads after 8pm are shifted to earlier periods, a base load remains. This is because houses cannot be pre-cooled indefinitely as this would cause discomfort to the occupants.

5.2.4 Discussion

The results presented in the previous sections lead to several implications for how to deploy our algorithm in practice. If peak reduction is the main target, not much extra

Household Count	100	200	300	400
Execution time (s)	8	30	60	230
Mem. Consumption (GB)	0.8	2	3.5	5.5

Table 5.3: Execution times and memory consumption of the centralised algorithm for a single day

energy needs to be used. A value of $b \leq 1.06$ for the extra energy usage factor is likely to suffice. When approximating the exact shape of a target profile, for example when trying to maximise the utilisation of energy produced from renewable sources, larger values between 1.06 and 1.14 for b might be more suitable. Values larger than 1.14 only provide minor improvements on peak reduction and target approximation. Such values could be advisable during times of very high HVAC utilisation, such as the hot summers in Texas. During times of usual HVAC utilisation, the extra energy used does not lead to significant improvements in the algorithm's load shifting capabilities.

While a centralised execution of the algorithm generally offers better results, the distributed version provides a good balance between scalability and performance. Table 5.3 shows execution times and memory consumption of the centralised algorithm on our test system (Intel Core i5-3570, 8GB RAM). Since solving a MIQP scales exponentially with the number of variables, and each household adds a number of new variables to the problem, the centralised version scales exponentially with the number of households. As DSM tends to happen on larger scales from several thousand to hundreds of thousands of households, the centralised version therefore quickly becomes infeasible. Solving the MIQP for a single household however only took between 1-2ms on our test system, suggesting that it can also be solved in acceptable time (less than a minute) on typical smart thermostat hardware (e.g. ARM Cortex A8 in the Nest thermostat). Apart from potential bandwidth and network latency bottlenecks, for example when collecting estimates of baseline profiles from households to calculate suitable target profiles, the distributed algorithm scales independently from the number of households.

5.3 Summary

In this chapter we introduced an extension to the HVAC agent to incorporate DSM signals to shift HVAC demand. The DSM algorithm is based on target profiles instead of commonly used dynamic energy price profiles. This helps to prevent rebound effects from happening and allows for more direct control over the resulting load profiles. We model the problem of optimising HVAC schedules with respect to user comfort, energy usage as well as deviation from the target profile as a MIQP, allowing to compute optimal solutions to the problem. Since solving the MIQP scales exponentially with the number of households, we provide a simple mechanism to run the algorithm in a distributed manner. This way, the algorithm scales independently from the number of households involved.

Using simulations based on real energy usage data, we evaluated the load shifting potential of the DSM algorithm. We showed that using the BCM, depending on the comfort compromiser used, the algorithm is able to achieve load profiles 29.75% to 39.9% closer to the target as compared to using a simple FSP, as it is common in most regular thermostats. We further investigated how the number of occupants influences the performance of the algorithm depending on the comfort compromiser used. We showed that when using AMPC or FSP, the algorithm's performance is mostly unaffected by the number of occupants. Using the OCC, the algorithm's ability to shift loads however suffers the more occupants are present. Regarding peak reduction capabilities, using the AMPC, the algorithm was able to almost completely flatten out peaks in the 99th and 95th percentiles. For lower percentiles, the algorithm's performance strongly depended on how much extra energy can be used to shift loads. Evaluations of the algorithm's performance throughout the year revealed that especially during times of high HVAC utilisation, the algorithm benefits from the ability to use more energy. Lastly, we gave an outlook how target profiles can be used to maximise utilisation of energy produced from renewable sources.

One of the shortcomings of the agent presented in this chapter is that the load shifting actually increases energy usage. In order to motivate users to participate and accept an increase in energy consumption, certain incentives must be set. There are a number of ways to address this. Energy suppliers could for example offer financial incentives in the form of a general rebate on the final bill for participants that opt in to load shifting. Opting in in this case means that they agree to hand over control to the HVAC agent, allowing it to increase their energy consumption if necessary. The rebate would be financed by savings through better utilisation of renewable energy sources as well as savings in grid infrastructure. Another approach could be to only sell the agent with DSM capabilities activated. The main incentive for consumers would still be to save energy and therefore money since the agent most likely would still reduce energy consumption as compared to a regular thermostat. More research needs to be done on how to incentivise consumers. However, this would go beyond the scope of this dissertation.

Chapter 6

Conclusions and Future Work

This work introduced the bayesian comfort model (BCM), a novel approach to a personalised thermal comfort model that unifies static and adaptive thermal comfort models, extends them and adds learning capabilities to adapt to users' preferences. Such a model can be used to increase user acceptance of autonomous smart heating systems by accurately predicting the optimal comfort temperature or range of acceptable temperatures for its users. This work further introduced an HVAC agent that uses information about users' ranges of acceptable temperatures to optimise the usage of the HVAC system. The agent can optimise solely for minimising the energy consumption of the HVAC system, but is also able to incorporate DSM signals to optimise the load profile generated by the HVAC system.

6.1 Discussion

In the following section, we discuss the outcomes of this research. We start by discussing the BCM. We then include a short discussion of the thermal comfort study and lessons learned from the deployments. This is followed by a discussion of both HVAC agents: the energy-centric agend and the DSM-centric agent.

6.1.1 The Bayesian Comfort Model

In comparison to existing thermal comfort models, the BCM includes a detailed user model to account for a wide variety of circumstances such as changing weather and different seasons and provides several possible outputs. The model was implemented as a belief network, which allows easy incorporation of prior knowledge about the problem and adds learning capabilities to constantly refine the model and adapt to the user's preferences. As a result, the BCM is more accurate than existing thermal comfort models. However, this accuracy comes at the cost of complexity, meaning that calculating users' comfort levels with the BCM is in general more computationally expensive than evaluating user comfort with existing, more simplistic models.

Through an empirical evaluation, we showed that using a belief network to learn user's preferences enables our model to give 13.2% to 25.8% more accurate results than existing models. Next to user comfort data from the ASHRAE RP-884 database, we conducted our own deployment in which we assessed thermal comfort of occupants in shared work spaces. In this deployment, we further tested an alternative to the ASHRAE 7-point thermal comfort feedback scale. Instead of asking for the user's current perception of the thermal environment, the desired change scale asks how the user wants the temperature to change. Using this scale, our model gives up to 55.4% more accurate predictions of user comfort as compared to existing models. While actual deviation from the optimal comfort temperature could not be assessed due to the lack of such data, the accuracy gains for predictions of user votes indicate that our model outperforms existing models on the research requirement of providing an accurate comfort model. These accuracy gains are achieved after about 8 observations. In a scenario where a user provides feedback every one or two days, this would equal an initial learning phase of roughly one to two weeks. This lies within the bounds defined for the unobtrusiveness and quick adaptation requirements.

One benefit of the constant learning in belief networks is that each new observation will most likely improve the quality of the predictions immediately. This means that even if the results are not optimal during the initial learning phase, a user should be able to see constant improvements, motivating the user to keep interacting with the system. Using the expectation propagation (EP) algorithm, we showed that the frequent task of performing inference on the model for predicting the user's current optimal comfort temperature is simple enough to be executed on small, inexpensive computers with low computational power or on remote servers. The scalability requirement can therefore be considered to be met.

6.1.2 Thermal Comfort study in Southampton

While the thermal comfort study in Southampton resulted in a sizeable data set that proved useful for the evaluation of the BCM, some lessons can be learned. A major issue was that the majority of feedback had to be discarded. The main reasons for this were that the feedback was either too one-sided or that single participants had not provided enough data points for a per-individual evaluation.

The large amount of one-sided, extreme feedback has several causes. In both deployments (office and library), interior temperatures were rather high, resulting in votes being skewed towards values indicating *too hot*. In case of the office deployment, the building was known to be warmer than other buildings. As a mitigation, the study was carried out early in the year when outside temperatures were still low. However, despite low outside temperatures, the interior temperature often went above 25°C. For future deployments, we therefore would recommend a more thorough assessment of what temperatures to expect during the duration of the deployment. This would have also helped with the library deployment. The high temperatures in the library deployment can be accounted to two different factors: high occupancy due to students studying for exams and a heatwave coinciding with the time of the deployment. The time of the study was deliberately chosen to be during the exam period in order to maximise the number of feedback received by students. To mitigate the effect of unusual occurences such as heat waves, we chose a longer time frame for the library deployment. However, by the time the heat wave was over, the exam period was close to its end as well and participation dropped significantly.

Apart from amplifying the effect of the heat wave in the data set, the drop in participation also resulted in a large amount of unusable data. By the time particpation dropped (roughly 10 days after the start of the study), the majority of participants had only provided feedback between 1-4 times, strongly limiting the usability of these participants' feedback for the evaluation of the model's learning capabilities. This effect was less pronounced in the office deployment, where despite having more than 10 times less participants, more useable feedback was collected. We believe that this is due to the office participants being more involved in the study. In the office deployment, occupants of affected offices were made aware of the study and its aims beforehand via email. Further, most occupants were present when the posters were deployed. Many occupants took this opportunity to ask questions about the study and the study design.

The qualitative analysis of the thermal comfort study is not in the scope of this dissertation. See Snow et al. (2017) for a more in-depth discussion of the studies and a quantitative analysis.

6.1.3 Reducing energy consumption of the HVAC system

Based on the BCM, we further created an HVAC agent that minimises energy consumption of the HVAC system while keeping a comfortable environment for the occupants at the same time. This agent utilised information about occupants' comfort ranges obtained from the BCM to autonomously decide and change set point temperatures. In comparison to most existing HVAC agents, our agent uses a detailed thermal comfort model with learning capabilities and puts a hard constraint on the set point temperature to stay within the users' acceptable bounds. The problem of computing such HVAC schedules was modelled as a linear program, enabling the agent to compute optimal HVAC schedules with respect to its optimisation goals.

In a realistic scenario, most of the time there will be multiple occupants to satisfy. This means that the agent either needs to incorporate multiple users' preferences at the same time or aggregate their preferences first. We proposed two simple approaches to a comfort compromiser that try to aggregate different users' preferences. The first compromiser is the overlap comfort compromiser (OCC) that tries to find the biggest overlap between different users' comfort ranges. The second compromiser is the average model parameter compromiser (AMPC) that creates a new instance of the BCM configured with the parameters set to the average of all users' model parameters.

Using simulations, we empirically evaluated the energy saving potential of our HVAC agent. We compared the BCM using both comfort compromisers against a simple fixed set point (FSP) and fixed comfort range (FCR). The evaluation showed that using the BCM can lead to energy savings of around 8% for heating and 24.3% for cooling, satisfying the research requirement of the agent being energy aware, while significantly reducing discomfort (by about 17.5%) at the same time, showing that agent is comfort constrained. We further investigated the effects the amount of occupants has on the algorithm's performance depending on the different compromisers and showed that while with the OCC the performance suffers as more occupants need to be satisfied, the AMPC seems mostly unaffected by it.

6.1.4 Performing HVAC-based load shifting

We extended the HVAC agent presented in the previous section to incorporate DSM signals into its scheduling algorithm and perform load shifting based on these signals. By modelling the HVAC scheduling problem with respect to DSM signals as a Mixed Integer Quadratic Program (MIQP), we were able to obtain optimal schedules. The scheduler uses target profiles describing a desirable load profile instead of dynamic energy pricing profiles, making it more stable and plannable while minimising the chance of rebound effects happening. We further increased the plannability of the algorithm by only optimising the deviation of the resulting load profile from the target profile. User comfort and energy usage were included as constraints rather than optimisation objectives. This ensures that the algorithm keeps the indoor temperature within the comfort range at all times. We constrained the energy consumption based on a scalar of the minimum possible energy consumption. In doing so, we provided full control over how much extra energy can be used by the algorithm to shift HVAC loads. To allow the algorithm to scale independently from the number of households, we modified the algorithm to run in a decentralised manner, independently in each household.

We empirically evaluated the algorithm's ability to shift HVAC loads. We used real world household energy consumption data to simulate multiple households and their HVAC consumption. Based on these simulations, we evaluated how increasing energy consumption aids the load shifting process. We showed that the distributed version generally yields results close to the optimal results obtained from the centralised version of the algorithm, with the centralised version on average achieving load profiles 12.8% closer to the target than the distributed version. In addition, we showed that using the BCM increases the algorithm's load shifting capabilities by 29.7% to 39.9%. Evaluations with varying numbers of occupants showed that the load shifting capabilities of the OCC decrease with a growing number of occupants, the AMPC scales independently from the number of occupants. In the scope of peak reduction, our evaluations showed that using the BCM, our algorithm is able to almost completely remove peaks in the 99th and 95th percentiles. We further investigated how the algorithm's load shifting abilities vary within the year and gave an outlook how target profiles can be used to maximise utilisation of renewable energy sources.

6.2 Limitations and Future Work

While we have presented two different comfort compromisers in this work, further work needs to be done on this topic. Currently, both comfort compromisers are easily gameable by the users. For instance, if someone generally prefers a warmer environment, this person could always provide feedback asking for a warmer environment/stating that it is too hot instead of reporting truthful feedback. This would bias this user's model towards learning more extreme parameters, leading it to suggest generally warmer comfort ranges for this user. For the AMPC this would mean that the average of the parameters would be skewed towards the manipulative user's parameters due to their more extreme values. In case of the OCC, the warmer comfort range of the manipulative user would cause the overlap with other user's comfort ranges to be minimised and shifted towards warmer temperatures.

For the AMPC, a simple way to address this issue would be to use the median of all parameters instead of the average. This way, parameters of manipulative users are likely to take more extreme values at the edge of the general value range would not have any influence on the final parameters. However, this approach only works well with larger number of occupants. For example, with just two occupants finding the median becomes problematic. Just choosing the parameters of one of the two users would likely cause a uncomfortable environment for the other user. Taking the average of these two user's parameters would open the system back up to manipulations by the users. There is promising work by Gupta et al. (2015) on finding strategyproof mechanisms to aggregate different users' comfort preferences, in which each user's agent is given a fixed budget that can be used to bias the temperature towards its user's comfort temperature. In future, a similar mechanism based on comfort ranges could be incorporated into our HVAC agent.

More work needs to be done towards evaluating the accuracy of the BCM and the agent's abilit to save energy and shift HVAC loads. While the accuracy of the BCM was already evaluated using real data, we believe further evaluations based on real deployments of the comfort model are required to provide deeper insights into the model's ability to keep a comfortable environment for its users at all times. Such deployments could also be used to collect data on actual comfort range profiles, providing useful data for more realistic simulations of the DSM algorithm. In addition, such deployments would provide useful insights into learning rates and how the learning and resulting increased accuracy of the model changes the user's interaction with the thermostat.

Similar to the BCM, we believe that a more thorough evaluation of the HVAC agent is required. At this point, while driven by real data, evaluations of possible energy savings and load shifting are purely based on simulations. To evaluate the actual energy savings and load shifting capabilities, long term deployments of the HVAC agent in actual households need to be conducted. Ideally, such deployments should take place in buildings where HVAC energy usage is already monitored (for example houses from the Pecan Street data port). This would allow for an easy assessment of energy savings by our HVAC agent. Alternatively, additional data on energy usage needs to be collected, for example by letting the HVAC agent act like a conventional thermostat. The same deployment could be used for experiments to assess the agent's load shifting capabilities. As a first step, the reaction of single households to different target profiles should be assessed. If households react similarly to how the simulations suggest, the number of households to be sent target profiles can be increased, with an increasing number of households receiving similar target profiles.

Another limitation of the HVAC agent presented in this work is its focus on only using the user's comfort ranges for optimising the set-point temperature. While this optimisation alone already yields good results, further improvements could be achieved. As a next step, the HVAC agent could for example take room occupancy into account to only adjust the HVAC in rooms that are actually occupied. in addition to that, the HVAC agent should also be able to distinguish between different types of rooms. In its current form, it is mainly tailored towards spaces of long, sedentary occupancy such as living rooms or office spaces. However, especially in a domestic setting, there are a number of other room types with very different needs. A kitchen for example is usually only for relatively short periods of time, while being fairly active. Such rooms should generally be kept colder by the HVAC agent. In addition, such rooms have more regular occupancy patterns, meaning that occupancy prediction could achieve significant savings in such rooms Another type of room are bedrooms which are usually unoccupied during the day and follow very predictable usage patterns during the night. For such rooms, it would be useful to take the time of the day into consideration. While the time of day is partially reflected in the BCM through outside temperatures (being lower during the night), for rooms like bedrooms very clear times could be specified during which the room is very likely to be inactive.

Bibliography

A. T. Alan, M. Shann, E. Costanza, S. D. Ramchurn, and S. Seuken. It is too hot: An in-situ study of three designs for heating. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 5262–5273, New York, NY, USA, 2016. ACM.

ASHRAE 55. Thermal Environmental Conditions for Human Occupancy (ANSI Approved), 2010.

ASHRAE, FUNIP. *Fundamentals Handbook*. IP Edition, 2009.

F. Auffenberg, S. Snow, S. Stein, and A. Rogers. A comfort-based approach to smart heating and air conditioning. *Special Issue on Urban Intelligence in ACM Transactions on Intelligent Systems and Technology (TIST)*, 2017, in press.

F. Auffenberg, S. Stein, and A. Rogers. A Personalised Thermal Comfort Model using a Bayesian Network. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 2547–2553. AAAI Press, 2015a.

F. Auffenberg, S. Stein, and A. Rogers. A heating agent using a personalised thermal comfort model to save energy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1799–1800, Richland, SC, 2015b. International Foundation for Autonomous Agents and Multiagent Systems.

Auliciems. Towards a psycho-physiological model of thermal perception. *International Journal of Biometeorology*, 25(2):109–122, 1981.

E. Bauer, D. 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.

R. Bellman. A Markovian Decision Process. Technical report, DTIC Document, 1957.

H. Chappells and E. Shove. Debating the future of comfort: environmental sustainability, energy consumption and the indoor environment. *Building Research & Information*, 33(1):32–40, 2005.

A. K. Clear, J. Morley, M. Hazas, A. Friday, and O. Bates. Understanding Adaptive Thermal Comfort: New Directions for UbiComp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, pages 113–122, New York, NY, USA, 2013. ACM.

Commitee on Climate Change. Meeting Carbon Budgets - 2017 Progress Report to Parliament Meeting. Technical report, 2017.

E. Costanza, J. Fischer, J. Colley, T. Rodden, S. Ramchurn, and N. Jennings. Doing the laundry with agents: a field trial of a future smart energy system in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 813—-822, Toronto, Ontario, Canada, 2014. ACM.

D. Daum, F. Haldi, and N. Morel. A personalized measure of thermal comfort for building controls. *Building and Environment*, 46(1):3–11, 2011.

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

R. de Dear and G. S. Brager. Developing an adaptive model of thermal comfort and preference. In *ASHRAE Transactions*, volume 104, pages 145–167. ASHRAE, 1998.

Department of Energy and Climate Change. Climate Change Act. Technical report, UK Government, 2008.

Department of Energy and Climate Change. The Future of Heating : Meeting the challenge. Technical Report March, UK Government, 2013.

Department of Energy and Climate Change. Energy consumption in the UK. Technical report, UK Government, 2016.

P. Du and N. Lu. Appliance commitment for household load scheduling. *IEEE Transactions on Smart Grid*, 2(2):411–419, 2011.

V. L. Erickson and A. E. Cerpa. Occupancy based demand response HVAC control strategy. *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building - BuildSys '10*, page 7, 2010.

P. O. Fanger. *Thermal comfort. Analysis and applications in environmental engineering.* R.E. Krieger Pub. Co, 1970.

A. A. Farhan, K. Pattipati, B. Wang, and P. Luh. Predicting individual thermal comfort using machine learning algorithms. In *2015 IEEE International Conference on Automation Science and Engineering (CASE)*, volume 2015-Octob, pages 708–713, 2015.

H. Farhangi. The path of the smart grid. *IEEE Power and Energy Magazine*, 8(1): 18–28, 2010.

P. Gadonneix, Y. D. Kim, K. Meyers, G. Ward, and C. Frei. World Energy Resources 2013. Technical report, World Energy Council, 2013.

M. D. Galus and G. Andersson. Demand Management of Grid Connected Plug-In Hybrid Electric Vehicles (PHEV). In *Energy 2030 Conference, 2008. ENERGY 2008. IEEE*, pages 1–8, November 2008.

L. Gelazanskas and K. Gamage. Demand side management in smart grid: A review and proposals for future direction. *Sustainable Cities and Society*, 11:22–30, 2014.

A. Ghahramani, F. Jazizadeh, and B. Becerik-Gerber. A knowledge based approach for selecting energy-aware and comfort-driven HVAC temperature set points. *Energy and Buildings*, 85:536–548, 2014.

S. K. Gupta, K. Kar, S. Mishra, and J. T. Wen. Incentive Compatible Mechanism for Coordinated Temperature Control in Multi-Occupant Buildings. In *2015 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 438–443, 2015.

T. Hargreaves, M. Nye, and J. Burgess. Keeping energy visible? Exploring how householders interact with feedback from smart energy monitors in the longer term. *Energy Policy*, 52:126–134, 2013.

J. C. Holyhead, S. D. Ramchurn, and A. Rogers. Consumer Targeting in Residential Demand Response Programmes. In *Proceedings of the ACM International Conference on Future Energy Systems*, pages 7–15, Bangalore, IN, 2015.

T. Hubert and S. Grijalva. Realizing smart grid benefits requires energy optimization algorithms at residential level. *Isgt 2011*, pages 1–8, 2011.

S. Kahrobaee, R. a. Rajabzadeh, L. K. Soh, and S. Asgarpoor. A multiagent modeling and investigation of smart homes with power generation, storage, and trading features. *IEEE Transactions on Smart Grid*, 4(2):659–668, 2013.

J. H. Klote. A General Routing for Analysis of Stack Effect. Technical report, NISTIR, 1991.

J. Liu, Y. Xiao, S. Li, W. Liang, and C. L. P. Chen. Cyber security and privacy issues in smart grids. *IEEE Communications Surveys and Tutorials*, 14(4):981–997, 2012a.

J. Liu, R. Yao, and R. McCloy. A method to weight three categories of adaptive thermal comfort. *Energy and Buildings*, 47(0):312–320, 2012b.

N. Lu. An evaluation of the HVAC load potential for providing load balancing service. *IEEE Transactions on Smart Grid*, 3(3):1263–1270, 2012.

T. Minka. Divergence measures and message passing. Technical report, Microsoft Research, 2005.

T. Minka. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 362–369, 2001.

N. Murtagh, B. Gatersleben, and D. Uzzell. 20:60:20 - Differences in Energy Behaviour and Conservation between and within Households with Electricity Monitors. *PloS one*, 9(3):e92019, 2014a.

N. Murtagh, B. Gatersleben, and D. Uzzell. A qualitative study of perspectives on household and societal impacts of demand response. *Technology Analysis & Strategic Management*, 26(10):1131–1143, 2014b.

J. W. Myers, K. B. Laskey, and K. A. DeJong. Learning Bayesian networks from incomplete data using evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 1, pages 458–465, 1999.

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.

M. Paciuk. The Role of Personal control of the Environment in Thermal Comfort and Satisfaction at the Workplace. In *21th Annual Conference of the Environmental Design Research Association*, pages 303–312, 1990.

P. Palensky and D. Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3): 381–388, 2011.

J. Palmer and I. Cooper. Housing energy fact file. Technical report, Department of Energy and Climate Change, UK Government, 2013.

J. Pearl. Fusion, propagation, and structuring in belief networks. *Artificial intelligence*, 29(3):241–288, 1986.

L. Peeters, R. De Dear, J. Hensen, and W. D'haeseleer. Thermal comfort in residential buildings: Comfort values and scales for building energy simulation. *Applied Energy*, 86(5):772–780, 2009.

L. Pérez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394–398, 2008.

S. Purdon, B. Kusy, R. Jurdak, and G. Challen. Model-free HVAC control using occupant feedback. In *Proceedings - Conference on Local Computer Networks, LCN*, pages 84–92, 2013.

S. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-Based Control for Decentralised Demand Side Management in the Smart Grid. In *AAMAS 11, Taipei,*, pages 5–12, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

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.

A. Rogers, S. Maleki, S. Ghosh, and N. 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.

A. Rogers, R. Wilcock, S. Ghosh, and N. R. Jennings. A Scalable Low-cost Solution to Provide Personalized Home Heating Advice to Households. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '12, pages 211–212, New York, NY, USA, 2012. ACM.

C. Sarkar, A. U. N. S.N., and V. Prasad. iLTC: Achieving Individual Comfort in Shared Spaces. In *Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks*, pages 65–76, 2016.

M. Seeger. Gaussian processes for machine learning. *International journal of neural systems*, 14(2):69–106, April 2004.

M. Shann, A. Alan, S. Seuken, E. Costanza, and S. D. Ramchurn. Save money or feel cozy?: A field experiment evaluation of a smart thermostat that learns heating preferences. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, pages 1008–1016, Richland, SC, 2017. International Foundation for Autonomous Agents and Multiagent Systems.

M. Shann and S. 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.

M. Shann and S. 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.

S. S. Shetty, H. D. Chinh, and S. K. Panda. 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*, number September, pages 23–27, 2015.

S. Snow, F. Auffenberg, et al. Log it while its hot: designing human interaction with smart thermostats for shared work environments. In *Proceedings of the 35th annual ACM conference on Human factors in computing systems - CHI '17*. ACM, 2017.

S. Snow, L. Buys, P. Roe, and M. Brereton. 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*, pages 245–254, 2013.

G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12): 4419–4426, December 2008.

L. G. Swan and V. I. Ugursal. Modeling of end-use energy consumption in the residential sector: A review of modeling techniques. *Renewable and Sustainable Energy Reviews*, 13(8):1819–1835, 2009.

The Engineering Toolbox. Air flow and velocities due to natural draft `http://www.engineeringtoolbox.com/natural-draught-ventilation-d_122.html`. retrieved at June 18 2014.

P. Thorsnes, J. Williams, and R. Lawson. Consumer responses to time varying prices for electricity. *Energy Policy*, 49:552–561, October 2012.

K. Valogianni, W. Ketter, and J. Collins. A Multiagent Approach to Variable-Rate Electric Vehicle Charging Coordination. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1131–1139, 2015.

Z. Wang, S. Member, R. Paranjape, and Z. Chen. Residential Demand Response : an Overview of Recent Simulation and Modeling Applications. *26th IEEE Canadian Conference Of Electrical And Computer Engineering (CCECE)*, pages 1–6, 2013.

F. Wernstedt, P. Davidsson, and C. Johansson. Demand side management in district heating systems. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems - AAMAS '07*, volume 5, page 1, 2007.

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.

R. Yang and M. 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.

R. Yang, M. W. Newman, and J. Forlizzi. 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*, pages 823–832, 2014.

Y. Zhou, D. Li, and C. J. Spanos. Learning Optimization Friendly Comfort Model for HVAC Model Predictive Control. *Proceedings - 15th IEEE International Conference on Data Mining Workshop, ICDMW 2015*, pages 430–439, 2016.