Development and evaluation of occupancy-aware HVAC control for residential building energy efficiency and occupant comfort

by

Christina Anne Turley

B.S., California Polytechnic State University San Luis Obispo, 2011

M.S., University of Oregon, 2012

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Master of Science

Department of Civil, Environmental, and Architectural Engineering

2020

This thesis entitled:
Development and evaluation of occupancy-aware HVAC control for residential building energy efficiency and occupant comfort
written by Christina Anne Turley
has been approved for the Department of Civil, Environmental, and Architectural Engineering

Prof. Gregor Henze, Ph.D., P.E.

Prof. Moncef Krarti, Ph.D.

Prof. Kyri Baker, Ph.D.

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Turley, Christina Anne (M.S, Architectural Engineering)

Development and evaluation of occupancy-aware HVAC control for residential building energy efficiency and occupant comfort

Thesis directed by Prof. Gregor Henze, Ph.D., P.E.

Occupancy-aware HVAC control offers the opportunity to reduce energy use without sacrificing thermal comfort. Residential HVAC systems often use manually-adjusted or constant setpoint temperatures which heat and cool the house regardless of whether it is needed. By incorporating occupancy-awareness into HVAC control, heating and cooling can be used for only the hours it is needed.

Bringing this technology to fruition is dependent on accurately predicting occupancy. Non-probabilistic prediction models offer an opportunity to use collected occupancy data to predict future occupancy profiles. Smart devices such as a connected thermostat, which already include occupancy sensors, can be used to provide a continually growing collection of data that can then be harnessed for short-term occupancy prediction by compiling and creating a binary occupancy prediction. Real occupancy data from six homes located in Colorado is analyzed and investigated using this occupancy prediction model.

Results show that non-probabilistic occupancy models in combination with occupancy sensors can be combined to provide a hybrid HVAC control with savings on average of 5.0% and without degradation of thermal comfort. Model predictive control provides further opportunities, with the ability to adjust the relative importance between thermal comfort and energy savings to achieve savings between 1% and 13.3% depending on the relative weighting between thermal comfort and energy savings. In all cases, occupancy prediction allows the opportunity for a more intelligent and optimized strategy to residential HVAC control.

Dedication

To my family and friends for their continuous and all encompassing support.

Acknowledgements

Occupancy data was collected and provided by Dr. Gregor Henze and Margarite Jacoby. The MPC optimization code was provided by Dr. Gregor Henze and Dr. Gregory Pavlak. This thesis would not have been possible without them.

Contents

Chapter

Figures

Figure

May 12, 2020

Chapter 1

Introduction

A finite quantity of fossil fuels and the mounting concern of climate change makes reducing energy use a global necessity. According to the IEA, buildings are major consumers of energy worldwide, using around 3,060 million tons of oil equivalent (Mtoe) in 2018 [7]. In the United States, heating, ventilation, and air conditioning (HVAC) systems account for 50% of building energy consumption [8]. In 2015, U.S. homes used 4.676 quadrillion British thermal units (Btu) for space heating and air conditioning [9]. Reducing U.S. residential heating and cooling consumption by 1% would save 46.76 trillion Btu, the equivalent of 3.5 Hoover Dams [10].

Conventionally, heating and cooling in residential buildings is controlled by a thermostat with a single setpoint temperature, keeping the indoor temperature constant. Over time technologies have been added to improve temperature control and energy use. Occupancy-based HVAC control is one. This method controls the indoor temperature to provide thermal comfort only when the building is occupied, turning the HVAC system on when the building occupied and off when vacant. This results in reduced energy use during unoccupied hours. Previous studies estimate savings between 5-23% depending on climate, building vintage, and occupant behavior [11].

While occupancy-based HVAC control has potential benefits, questions still exist on how best to detect and implement occupancy-based HVAC control. Because thermal comfort standards should be met during all occupied hours, a good control strategy needs to accurately predict occupancy and condition the space just in time for the occupants' arrival. The purpose of this thesis is to optimize residential building control based upon an occupancy prediction model and then compare the estimated energy use of two control strategies to each other and conventional controls.

1.1 Background and Motivation

Energy use has been growing worldwide. The built environment accounted for 40% of total global energy use and 33% of total greenhouse gas (GHG) emissions in 2016 [12]. In the United States, residential buildings use 22% of the total annual energy, while commercial buildings use 18% [7]. Occupants now spend more than 90%, on average, of their time indoors [5]. With an increasing population and increasing amount of time spent indoors, building energy use is predicted to continue to rise. Annual energy use in the U.S. residential sector has increased by 261% since 1950, see Figure 1.1. The Energy Information Administration (EIA) estimates that energy use in the built environment will increase by 34% between 2010 and 2030 [7].

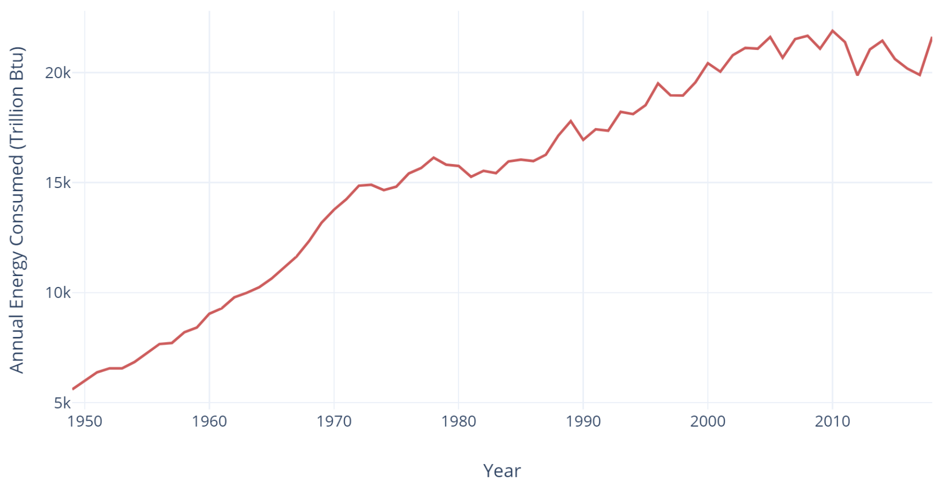

Figure 1.1: Annual energy consumption for U.S. residential sector [TBTU]

In parallel, the Paris Agreement, signed in 2016 as part of the United Nations Framework Convention on Climate Change (UNFCCC), sets forth the goal of "holding the increase in the global average temperature to well below 2°C above pre-industrial levels" [13]. To meet this goal, countries worldwide will need to reduce energy use. Energy reduction and efficiency is a critical component of buildings moving forward.

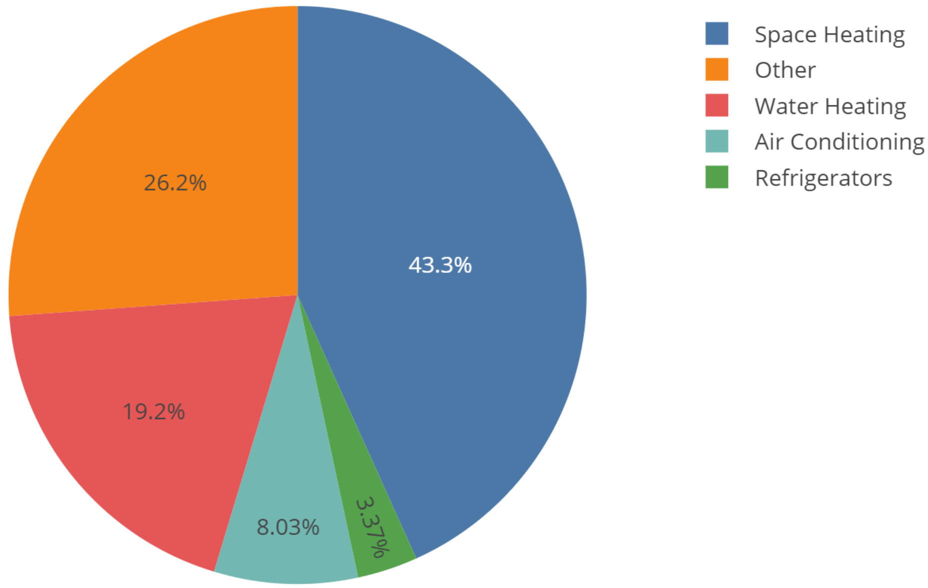

Figure 1.2: Residential building energy end-use consumption in the United States in 2015 [%]

Economic growth, expanded housing stock, and expanded building services, particularly HVAC systems, have all attributed to the growth in total residential energy consumption. While climate, architectural design, energy systems, and economic status of the occupants affect individual building consumption, an average of 50% of energy consumption goes to HVAC systems, see Figure 1.2 [7]. Thermal comfort, once considered a luxury, is now expected. This has created a growing reliance on HVAC systems, increasing energy use as air-conditioning becomes more widespread, see Figure 1.3 [9, 14].

Incorporating occupancy-based HVAC control offers reduced energy consumption while maintaining thermal comfort. 60% of American adults are employed in full or part-time jobs, leading to vacant homes for an average of 21% of the day [15]. In lower population households, vacancy times were measured up to 46% [3]. At a 21% vacancy, over 1800 hours of the year would not require space conditioning.

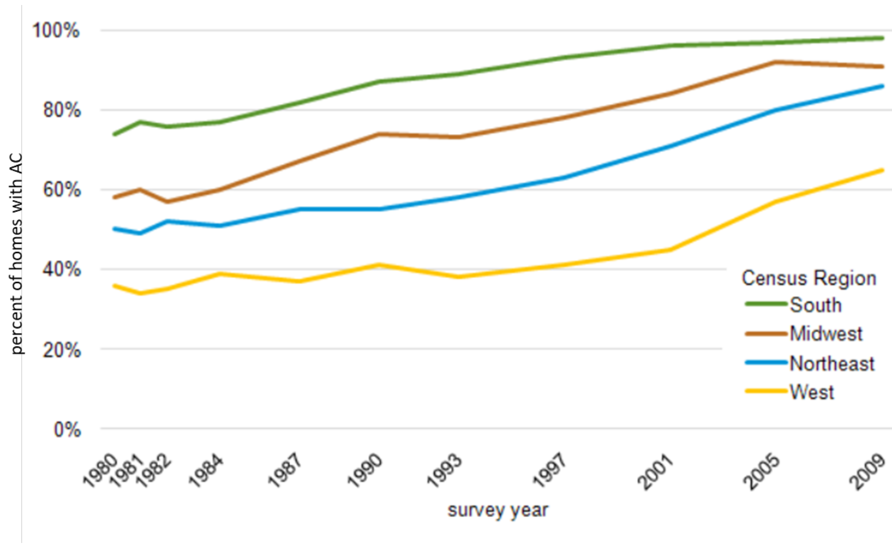

Figure 1.3: United States residences with AC by year [%]

1.2 Objectives

Realizing energy savings from occupancy-based HVAC control is contingent upon accurately predicting occupancy. Errors can result in two failure modes: 1) reduced energy savings when a vacant space is conditioned or 2) increased unmet thermal comfort hours when an occupied space is not properly conditioned. Both errors fail the objective of providing thermal comfort only when occupied.

The overarching goal of this project is to determine the optimal occupancy prediction model for the measured occupancy data, using different training durations and temporal resolution, and then simulate the estimated energy usage for typical single-family residences. The energy usage will be modeled for three HVAC control scenarios: 1) conventional setpoint operation, 2) occupancy-based HVAC control, and 3) occupancy-based model predictive control (MPC). The energy savings potential of the latter two will be analyzed and compared to conventional operation.

The aim of this thesis is to:

(1) Provide insights into how HVAC systems in residential buildings are controlled

(2) Provide a framework for how occupancy can be used in HVAC control

(3) Identify the best model for occupancy prediction

(4) Estimate the energy used in standard, MPC, and occupancy-based HVAC control

(5) Identify the best control method to reduce energy use

(6) Evaluate how climate, setpoint temperature, and house design affect energy usage

1.3 Thesis Organization

The thesis is organized into five chapters: the introduction, literature review, occupancy model generation and discussion, building simulation evaluation and discussion, and summary and conclusions. In the literature review section relevant studies on housing, HVAC control, and occupancy modeling will be examined. The occupancy model generation and discussion section will cover how the occupancy data was collected and an optimized occupancy prediction model was generated. The building simulation evaluation and discussion section will review how the building energy models were set up and the results of each control strategy will be analyzed. Summary and conclusions will present the effectiveness of occupancy-based HVAC control and provide possible future work based on the results from this thesis.

Chapter 2

Literature Review

Understanding of how buildings are changing is critical to reducing energy used in the built environment. Looking at how building control systems have performed historically can indicate opportunities for improvement and how current trends may shape the future. This chapter discusses general trends in residential buildings, comfort requirements, temperature control in buildings, and how occupancy has been incorporated in homes to date.

2.1 Current Landscape in Residential Heating and Cooling

2.1.1 Trends in U.S. Housing

The 2015 Residential Energy Consumption Survey found that the United States residential sector was comprised of 118.2 million homes, totaling 223 billion square feet of floor space [9]. Residential buildings use 22% of U.S. annual energy. In the three decades from 1980 to 2009, energy consumption increased by 0.9 quads or 8.9% [1].

Main contributors to increasing energy use were increasing size, number of homes, and appliance use. From 1980 to 2009, the number of households increased by 33% while the floor space increased by 52% [1]. The average single-family detached home increased from 2100 square feet to 2688 square feet, see Figure 2.1. In parallel, appliance electricity consumption increased by 30.6%, with the largest increase from microwave ovens, personal computers, air-conditioning, and clothes dryers [1].

Contributors to energy reduction were population shift, changes in weather, and decline in

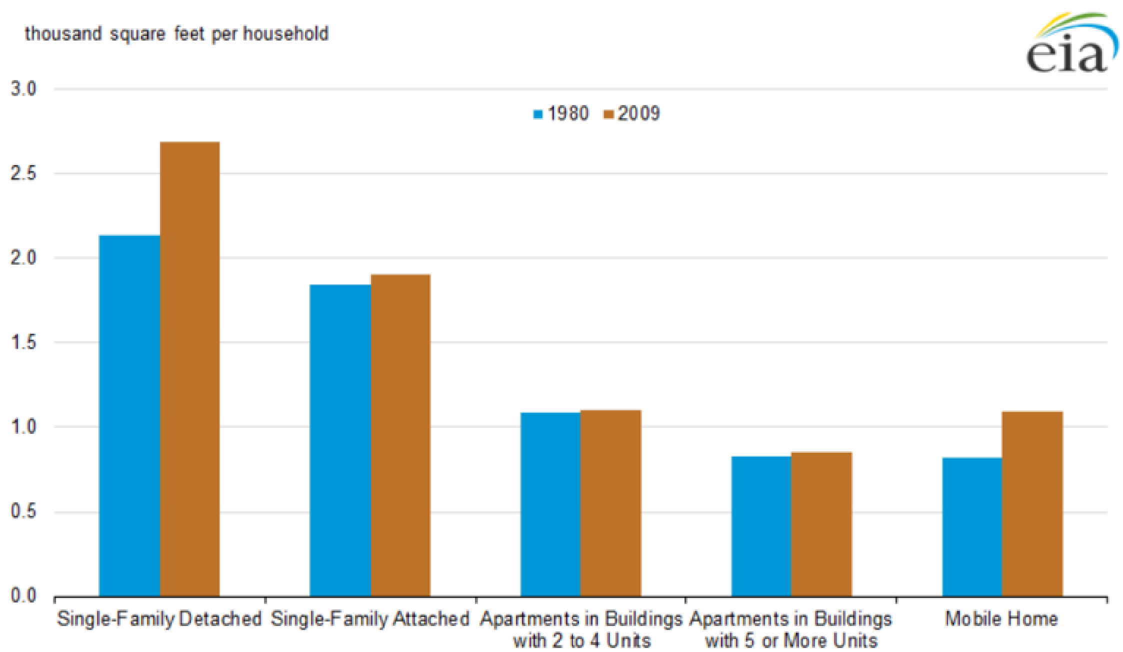

Figure 2.1: Size of residences by home type for 1980 and 2009 [ft^2] [1]

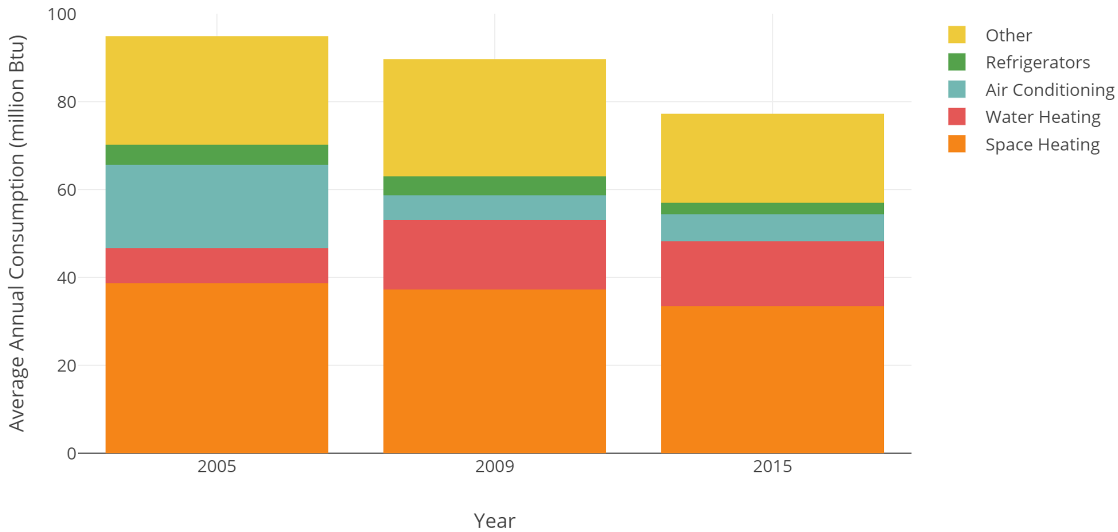

Figure 2.2: Annual household energy use by year [MMBTU]

energy intensity, especially in space heating. The U.S. population shifted from the Northeast and Midwest to the less heating-intensive West and South. Weather changed nationwide, with less cooling degree days and heating degree days per year in 2009 than 1980 [1] leading to natural reductions in energy use.

Additionally, advances in engineering and energy efficiency standards promoted the reduction in energy intensity of household appliances, see Figure 2.2. For example, the annual fuel utilization efficiency (AFUE) of a standard furnace increased from 78% to 97%. The largest change in consumption occurred from 1990-2001, at the same time federal efficiency programs, like ENERGY STAR, were enacted. The combined effect of all contributors in the 30-year period, was an increase in energy consumption over time, with U.S. homes consuming 9.1 quads per year [1].

2.1.2 Standards for Thermal Comfort

Despite the changes in the housing sector, the one constant is thermal comfort. ASHRAE Standard 55 states that buildings must be conditioned so "the mind expresses satisfaction with the thermal environment" [16]. The standard, first published in 1966, specifies the fraction of occupants that find a space comfortable. ASHRAE 55 is based on the predicted mean vote (PMV), first developed by Povl Ole Fanger. In his studies, Fanger varied air temperature, mean radiant temperature, relative humidity, air speed, metabolic rate, and clothing levels for participants to determine the satisfaction percent [17]. ASHRAE 55 requires at least 80% of occupants are satisfied, resulting in two comfort regions, one for summer and one for winter clothing levels.

2.2 Temperature Control in Buildings

Temperature, controlled using a thermostat and control system, is used to meet the comfort requirements in homes. The thermostat measures the indoor air temperature and compares it to the setpoint temperature while the control system manages how the HVAC system returns the indoor air to the setpoint. The interaction and programming of these two components determines how effectively the heating and cooling is accomplished.

2.2.1 Thermostats

Manual, programmable, and "smart" thermostats are the three main categories of thermostats. In manual thermostats, the setpoint is a single temperature. Programmable thermostats set temperatures by time of day or week [18], allowing temperatures to be setback during night-time hours or during daytime vacancies. Endorsed by ENERGY STAR at their 1995 release, initial demonstrations showed that programmable thermostats could reduce heating and cooling bills by 10-30%. Ultimately, the EPA suspended ENERGY STAR certification of programmable thermostats in 2009 because no energy savings materialized [18]. Investigations revealed that 30% of households had not set them and over 89% had not set separate weekend and weekday schedules [3]. Because of their complexity, most programmable thermostats were operated manually, negating their energy savings potential.

Connected, or smart thermostats, emerged in response to consumer aversion to programmable thermostats. Connected thermostats are internet-connected. Like programmable thermostats, connected thermostats create a setpoint schedule but the operation is designed to be user-friendly. Products like the Nest, Honeywell, or Ecobee thermostat can be controlled by phone, web interface, or a touchscreen. The system comes with a preset schedule that modifies itself based on the user's manual adjustments during initial use [19]. First introduced in 2011, the Nest Thermostat catalyzed the market, with over 100% market growth per annum. In 2015, 40% of the thermostats sold were connected thermostats [19]. The main features of most connected thermostats include extensive data tracking, remote accessibility, local sensors to track occupancy, and web-enabled weather forecast data. All are designed to enhance temperature control. Because connected thermostats are automatically programmed through initial use, they have an opportunity to create setbacks that save energy.

2.2.2 HVAC Control Strategies

The HVAC control strategy determines how the heating and cooling systems are run. The two main controller types for residential buildings are discontinuous and continuous controllers. Discontinuous controllers turn on and off, allowing only two states of operation. The benefit of on/off controllers is the easy installation and operation. Disadvantages are the time delay; the thermal inertia of buildings produces large swings from setpoint temperatures [20].

Continuous controllers modulate the rate of heating and cooling to provide heat transfer at the pace it is needed. Continuous control is normally provided by proportional-integral-derivative (PID) controllers. This includes a feed-back loop that works to minimize undershoot, overshoot, rise time, settling time, and steady-state error [3]. This control type is not widely used in the residential sector.

Additional inputs, beyond indoor temperature, can be fed into the control system. Two we will consider in this thesis are future occupancy and weather. In this case, an advanced control strategy of model predictive control (MPC) is used to control the building. In MPC, the future state of the building is predicted by considering multiple inputs. For example, the weather forecast and current indoor temperature are fed into a model to predict how the building temperature will change over time. An optimization is performed on the model to determine which actions will achieve the required temperature while minimizing energy use. The optimal control action will then be sent to the HVAC system telling it how to run [21, 22]. MPC, through modeling a future state, acts as a real-time building control by rerunning the model at consistent time intervals. Known as receding horizon control, a new optimized control strategy is determined as temperature and weather forecasts are updated [23]. Because MPC is actively predicting the future indoor temperature state, it is proactive rather than reactive and can anticipate future needs.

MPC typically uses a first-principle linear model that models the building as a lumped element resistance-capacitance (RC) model [22]. This means the heat transfer in and out of the building is simplified to a linear expression, making the optimization convex and easier to solve. Buildings do not always act linearly, creating inherent errors [22]. The work of properly calibrating a model for individual buildings has kept MPC from widespread adoption [24]. MPC continues to show promise and is predicted to gain traction with research showing residential energy savings of 28% on average and cost savings of 16% [24, 25].

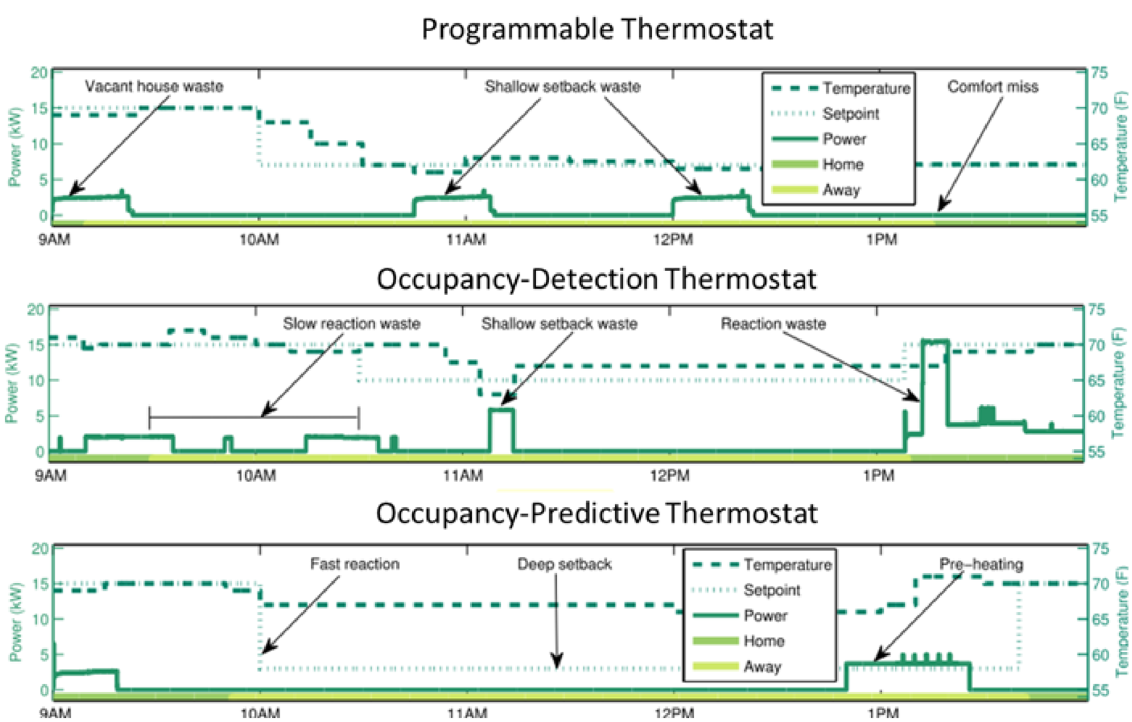

Figure 2.3: Daily power, temperature, and setpoint for programmable and occupancy-based thermostats [kW, °F] [2]

When occupancy is included in HVAC control, the setpoint temperature is used only when occupied. When vacant, the HVAC system is turned off and the indoor temperature drifts, see Figure 2.3. This minimizes total energy use. Occupancy-based controls can be reactive or predictive. In reactive control, the system turns on when an occupant is detected. The temperature may be uncomfortable until the system returns the indoor temperature to its setpoint. In predictive control, the arrival time of the occupant is predicted and the system preheats or cools so the indoor temperature reaches the setpoint right when the occupant arrives, minimizing energy while maintaining comfortable temperatures during occupied hours [2]. Correctly predicting occupancy to use predictive control is the main challenge facing industry today.

2.2.3 Energy Savings Potential of HVAC Systems

Industry has been working to ascertain the energy savings potential for thermostats and control systems. Nagele et al. conducted a survey of 30 households in southern Germany over 14 months. Using the setpoint temperature and house characteristics, the energy use for ten households was simulated for eight thermostat and control strategies. The on/off controller was used as a reference case. Results show that PID controllers, setback temperatures, MPC and occupancy-based HVAC control all have the ability to save energy when programmed correctly, see Figure 2.4.

Occupancy detection offers the largest energy savings but increases unmet comfort hours. To gain consumer adoption, unmet hours need to be low enough that consumers do not turn the control off. This makes occupancy-prediction the preferred control choice. In simulation studies of occupancy-prediction control, savings are estimated from 6-48% depending on climate, insulation levels, and occupancy schedule [2, 26]. These studies will be discussed further after the different occupancy-prediction models are introduced in Section 2.4.1.

Beyond simulations, utilities have measured the energy savings of connected thermostats, see Figure 2.5 [27]. In reviewing 35 studies from 2007 to 2016, the DOE reported energy savings ranged from 1% to 15% [19]. Definitive values are hard to determine due to the variety of hardware,

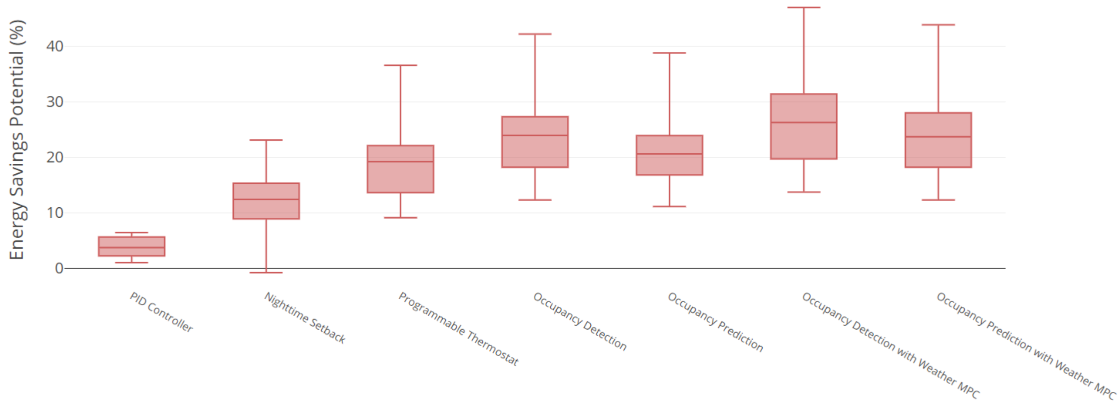

Figure 2.4: Boxplot of potential savings by control strategy for ten households [%]. Adapted from [3]

software, buildings, occupant behavior, and local weather.

Although pre-post assessments are industry protocol, additional variables can change, skewing results [28]. With a large sample size, individual variables should be negated. Overall, studies show a large variation in energy use and savings. This raises the question: What is its cause and can it be solved?

2.3 Occupancy in Buildings

A 2017 international survey of professionals in the Journal of Building Performance Simulation lists occupant behavior as the largest contributor in energy modeling errors, see Figure 2.6 [4]. Human behaviour is varied and stochastic, with behaviour changing by person and by day. This makes correctly predicting occupant behaviour, and its effect on buildings, extremely difficult.

2.3.1 Why Occupancy Matters

In 1978, Robert Socolow published a 5-year observational study on occupant behavior. The group tracked gas consumption of 205 identical townhouses in Twin Rivers, New Jersey. They found a 33% variation in consumption [29]. This revealed that identical buildings vary due to occupants' setpoint temperatures and hot water use. Similarly, a study in Kuwait showed that

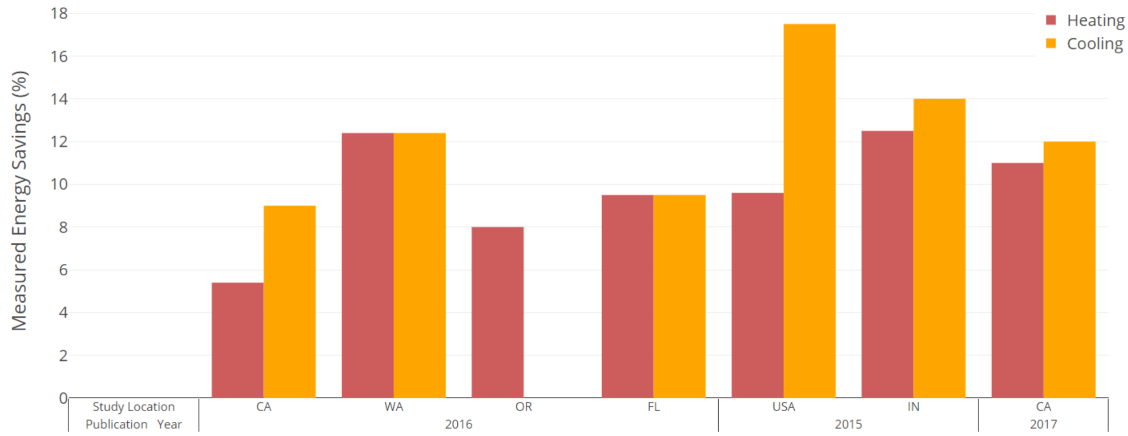

Figure 2.5: Average measured energy savings in eight connected thermostat studies [%]

residents used setpoint temperatures between 19 °C and 25 °C for air-conditioning, with electricity use increasing 21% with a 2 °C change [30]. In Denmark, Rune Andersen collected four years of annual heating data from 290 identical townhouses. Again, a wide variation was found with annual heating consumption ranging from 9.7 kWh/m^2 to 197 kWh/m^2, a ratio of 20 to 1 [31]. These studies, carried out in different climates, continents, and decades, show the significant impact of occupant behavior.

2.3.2 Occupancy Behavior

Occupant behavior is categorized into two categories: adaptive and non-adaptive. Adaptive behaviors are classified as actions used to increase comfort, such as opening windows or changing a setpoint temperature. Non-adaptive behaviors are actions that do not relate to comfort, such as presence or use of home appliances. Both types affect energy use [32]. Research has increased dramatically in the last decade, see Figure 2.7, indicating interest in solving the human dimension in buildings [5]. In this thesis only the non-adaptive behavior of occupant presence will be studied.

Table 2.1: Summary of occupancy presence models. Table adapted from [33]

Level	Model	Complexity
Group	Schedule	•
	Deterministic	• •
	Non-probabilistic	• •
	Probabilistic	• • • • •
Individual	Agent-based	• • • • • • • • • •

2.4 Incorporating Occupancy Models

Modeling the impact of occupancy behavior on energy consumption is comprised of two steps: 1) create an accurate model of occupancy and 2) incorporate the model into a building performance simulation (BPS). Research in the past decades has investigated the best method for performing each step. The results of previous studies are presented in the following two sections. Occupancy modeling will be discussed first, in Section 2.4.1. BPS will be discussed in Section 2.4.2. Both steps are necessary to understand the impact of occupancy-based HVAC control on energy use.

2.4.1 Modeling Occupant Presence

An accurate occupancy model is critical. The 2017 industry survey results in the Journal of Building Performance Simulations showed that industry professionals believe current models over-simplify real behaviors. When assumptions are overly optimistic, the building consumption performs below expectations. Conversely, conservative assumptions oversize mechanical equipment [4]. While many occupant models have been published in scientific papers, an industry consensus has not been achieved.

Occupancy can be predicted at two levels: group or individual. In the group level, one model is created for the entire group occupying a building. This is the most common. Individual level models create a separate model for each occupant. A summary of the models are shown in Table 2.4.1 and a description of each model follows [33].

Schedules

Schedules are the current industry standard for modeling occupancy presence. The schedules are provided by ASHRAE by space type. The fraction of occupancy, see Figure 2.8, is multiplied against space density to determine the number of people during each hour. The schedules are static, with only three schedules used for the year: weekday, Saturday, and Sunday [11].

Studies have compared ASHRAE occupancy schedules to real occupancy data. In 2013, Duarte et al. published a study of a 195,000 ft^2 11-story multi-tenant commercial office building in Boise, Idaho, using occupancy sensors to record the occupancy status for two years. The analysis showed private offices had a maximum occupancy of 50% compared to the 95% used in ASHRAE 90.1-2004 schedules. Open offices reached a peak of 80%, also well below 95% [34]. Occupancy fractions differed by seasons and days of the week, with Mondays and Fridays having different schedules from the rest. The static ASHRAE schedules diverged by as much as 46%.

Deterministic Models

Deterministic models use a rule-based approach to represent occupancy behavior. For example, a rule that blinds will be closed when indoor air reaches a certain temperature models window blind operation [35]. Unlike schedules, deterministic models incorporate environmental triggers that can affect actions.

Non-Probabilistic Models

Non-probabilistic models use historical data to create a model. The aggregated data is averaged to create a probability profile, with each time interval having a probability between 0 and 1. If the probability is above a threshold, the building is predicted to be occupied; below the threshold, vacant. Multiple non-probabilistic models can be made if distinguishing between days of the week, for example, is desired [33]. Because the profile is created from a training set, the accuracy of the model depends highly on the data used. The model created does not include a stochastic term. To include changes in behavior, non-probabilistic models can use a moving training set to create individual daily profiles [36].

Probabilistic Models

Probabilistic or stochastic models incorporate the variability of human behavior by using randomization. Instead of using a fixed departure and arrival times, a random non-repeating daily profile is created [36]. Like non-probabilistic models, stochastic models use historical data to create a model. A probability profile is created and compared to a randomly generated number to classify the space as occupied or vacant. The most common algorithm used in stochastic modeling is a Markhov chain. In a Markhov chain, the probability profiles represent the likelihood of transitioning between the two states, see Figure 2.9. At each time step, a random number between 0 and 1 is generated and compared against the transition probability profile. This comparison determines whether the state changes or stays the same. Markhov chains use only the previous time step [37]. Because a random number is used, a different profile will result each time the model is generated. Stochastic models require multiple runs to achieve reliable results. Except for the Markhov chain, stochastic models decouple occupancy from time of day [34].

Agent-Based Models

Agent-based models model occupants individually, aggregating multiple prediction models to create a full building model. Because modeling is done on an individual basis, the complexity is extremely high. Agent-based models allow high resolution; the location and actions of each occupant is predicted within the building [38]. Agent-based models are useful when detailed actions within the building, such as window opening on a certain floor, are desired.

2.4.2 Modeling Building Performance

The second step to incorporating occupancy is loading the model into the building simulation. There are many simulation programs available. The International Building Performance Simulation Association (IBPSA) lists sixty-seven whole building energy simulation programs [39]. EnergyPlus, developed and distributed by the U.S. Department of Energy (DOE), is used most commonly in occupancy research [4]. EnergyPlus is a compiled physical model; the characteristics of the building such as insulation values, window size, and orientation are built to create a building model [2]. The mechanical equipment and schedules, such as occupancy, are included as inputs to the building operation. When conducted, the simulation simulates the heat and mass transfer for each time step [11]. Simulations are normally performed per annum to integrate both heating and cooling seasons [2]. ASHRAE occupant schedules are embedded within the standard EnergyPlus software but different occupant models can be incorporated by co-simulation. An external interface, like the customizable energy management system (EMS), feeds the deterministic or probabilistic occupant model into EnergyPlus and conducts the simulation with the updated occupancy [11].

2.4.3 Review of Commonly Used Occupant Models

Past studies have sought to answer the question of which occupant model works best to predict occupancy [21, 23, 26, 32, 34, 37, 38, 40–45]. Because the published studies were conducted using individually collected occupancy data and climate and building-specific BSP models, it has been difficult to compare different occupant models directly [33]. Individual analyses have sought to solve this by comparing different occupancy models made with the same occupancy data. A review of occupant presence comparison studies is summarized in this section.

Mahdavi and Tahmasebi 2015

This study compared three models: two stochastic models from literature (Reinhart 2004 and Page 2008) and a non-probabilistic model they had developed. Using data from eight workspaces, the three models were created using four weeks of training data then predictive schedules were created for the next 90 working days. Predicted occupancy was compared to the measured ground-truth to analyze the prediction model's capability. The model was evaluated by comparing the arrival time, departure time, duration of occupancy, fraction of correct occupancy state, and number of transitions to the ground-truth data. Analysis showed that the two stochastic models performed similarly, while the non-probabilistic model performed best. Mahdavi and Tahmasebi conclude that while probabilistic models are suitable for annual simulations, non-probabilistic models are more effective in providing short-term occupant presence predictions [36]

Tahmasebi and Mahdavi 2017

Following the 2015 study, Tahmasebi and Mahdavi input their occupancy models into a building simulation program to determine the effect of the occupancy prediction on building performance. Plug load, occupancy, and light use data collected from the university office building for the previous three years was used for analysis. Six occupancy models were created. The first used the ASHRAE 90.1 office schedule. The second used the average group occupancy data for the year, while the third used the average individual occupancy data for the year. A stochastic model for each of the three was created to generate a total of six models.

EnergyPlus was performed with each occupancy model to simulate energy use; stochastic models were executed using 100 Monte Carlo runs to find the average performance. Performance was evaluated by looking at the key performance indicators of annual and peak heating and cooling loads, see Figure 2.10. ASHRAE schedules performed poorly in all metrics. Stochastic models of individual and grouped occupancy performed better simulating heating loads while the fixed models performed better simulating cooling loads. Tahmasebi and Mahdavi conclude that known occupancy data is critical for accurate building performance simulation; stochastic models are not [46].

Duarte, Van Den Wymelenberg, and Rieger 2013

Duarte et al. perfomred an occupancy study on a multi-tenant 11 story office building in Boise, Idaho. Using data from 223 private offices over two years, probabilistic and ASHRAE 90.1 schedules were compared to a non-probabilistic model. For the non-probabilistic model, occupancy changed by month and day of the week, with Mondays having the highest occupancy and Friday the lowest. Comparing the different occupancy models, the ASHRAE 90.1 schedule overestimated occupancy by as much as 46%. Using data from 10 offices for training, the stochastic model matched the training data but not the overall measured occupancy. Duarte, Van Den Wymelenberg, and Rieger recommend using a low and high non-probabilistic model because it represents occupancy well without increasing modeling complexity [34].

In all comparison studies, the authors agree that the best model is case specific [33]. Most models were developed using single data and building sets and do not transfer effectively to different building types or occupant behaviors [47]. Despite this, some general conclusions can be drawn. ASHRAE occupancy schedules are not reflective of actual behavior. Model complexity, such as stochasticity, does not always improve results. Finally, models perform best when applied to the case study used to derive them [33]. Because there is no universal occupant prediction model, the Annex 66 consortium recommends choosing a model that matches the complexity levels of the occupant model to the case study.

Because this thesis aims to evaluate the possible energy savings on short-term occupancy-based HVAC predictive control, a non-probabilistic model, which was shown to have the best short term presence prediction, will be used [34, 36].

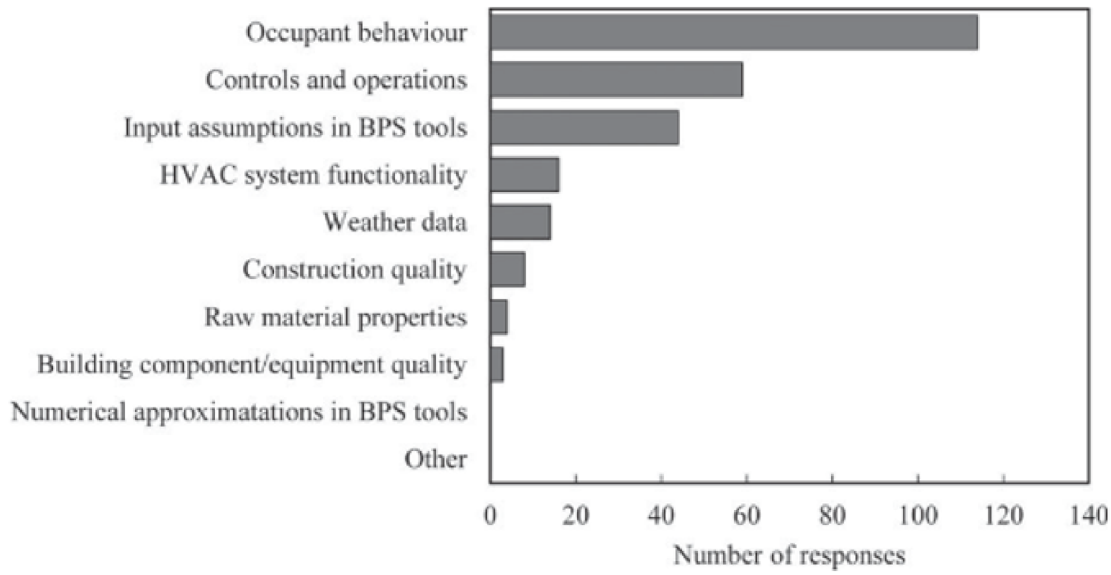

Figure 2.6: Participant response for main cause in discrepancy between simulation software and measurements [4]

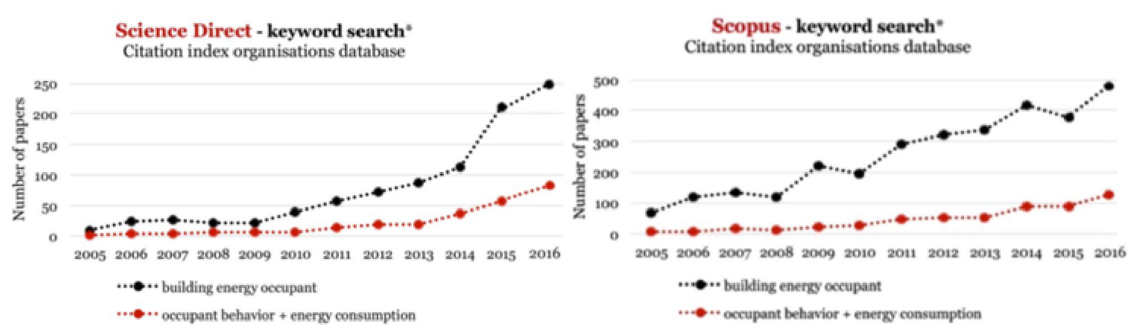

Figure 2.7: Number of occupant behavior papers [5]

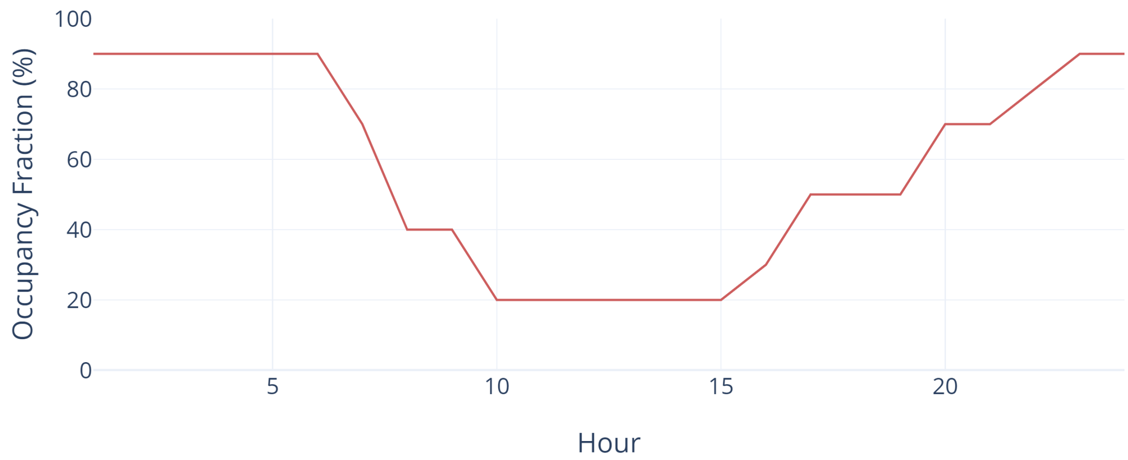

Figure 2.8: ASHRAE occupancy schedule for hotel on a weekday [%]

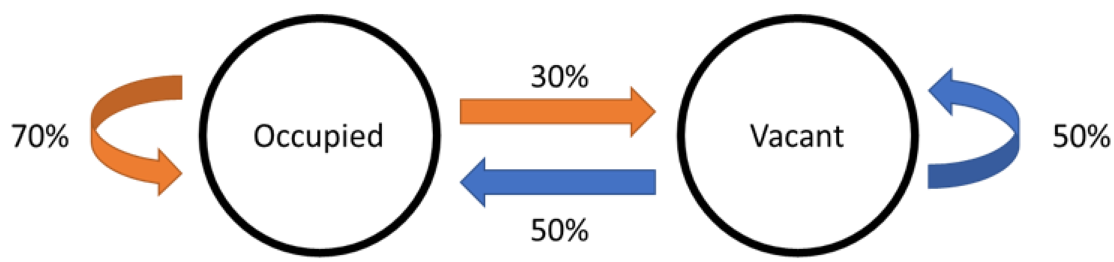

Figure 2.9: Graphic representation of Markhov chain transitions

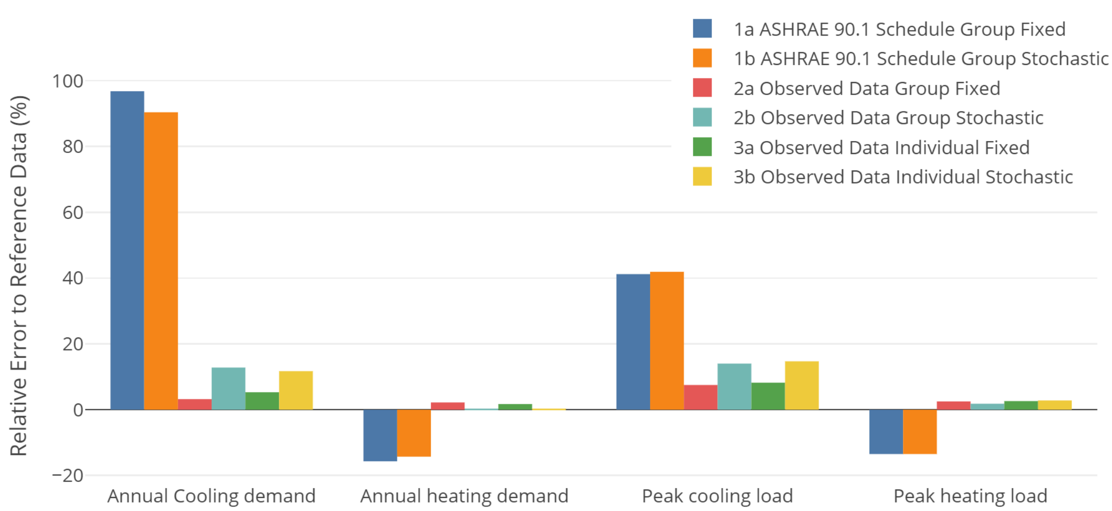

Figure 2.10: Results of BSP model accuracy from 2017 study [%]. Adapted from [6]

Chapter 3

Occupancy Model Generation and Discussion

To evaluate the impact of occupancy-based HVAC control, occupancy had to be determined. Real occupancy data was collected from multiple homes located in Boulder, Colorado and used to create non-probabilistic occupancy prediction models. Effectiveness was determined by comparing the prediction model against actual occupancy. The details of this process are explained in the following sections.

3.1 Ground Truth Data Collection

Occupancy and eight different physical modalities (e.g. CO_2 and VOC) data was collected from six homes for 4-9 weeks. Occupancy data was collected using a sign-in sheet and a geo-fencing phone tracking application. Two collection methods were cross-referenced to confirm correctness. Occupancy arrival and departure times were recorded for each occupant. Individual occupant data was combined to determine the binary occupancy state of the residence. General occupancy information for each residence is shown in Table 3.1. Average occupancy ranged from 52% to 86%.

Residences used for the study were chosen from volunteer participants in Boulder, Colorado. Five of the six residences included occupants that were university graduate students. Occupancy patterns of the collected homes may differ from occupancy patterns for different age groups and professions. Some factors that impacted occupancy patterns for the studied group were long absences during spring, summer, and Thanksgiving break.

While the average occupancy for the testing period ranged from 52% to 86%, the daily

Table 3.1: Summary of residences measured during study

House #	Occupant Count	House type	Days Measured	Avg. Occupancy
1	4	single-family residence	64	86%
2	1	apartment	45	56%
3	3	single-family residence	71	75%
4	3	apartment	29	82%
5	2	apartment	27	81%
6	1	apartment	63	52%

occupancy of each home varied. Distribution of daily occupancy rates for each residence is shown in Figure 3.1. Home 6 had the largest distribution in daily occupancy, while home 5 had the lowest.

Daily departure and arrival times for each residence are shown in Figures 3.2 and 3.3. Arrival is defined as the transition of the residence from an unoccupied state to occupied state; departure is the inverse transition. In residences with more than one occupant, the arrival and departure times identified only indicate cases where the residence transferred to or from vacant. For all residences measured, the arrival and departure times for weekdays differ significantly from weekend days. House 4 had the largest variation in departure times, house 1 the smallest. House 2 has the largest distribution in arrival times, house 1 the smallest. Overall, house 1 had the smallest time distribution for arrival and departure times.

Multiple arrival and departures occurred in a single residence each day. Daily transitions can be used as a metric to understand the number of occupancy and vacancy states. A transition includes both arrivals and departures. Figure 3.4 shows the distribution of transitions for each residence.

Distributions in daily occupancy, arrivals, departures, and occupancy state transitions demonstrate the stochasticity of human behavior. The range of homes tests the adaptability of the occupancy prediction models to different occupancy schedules.

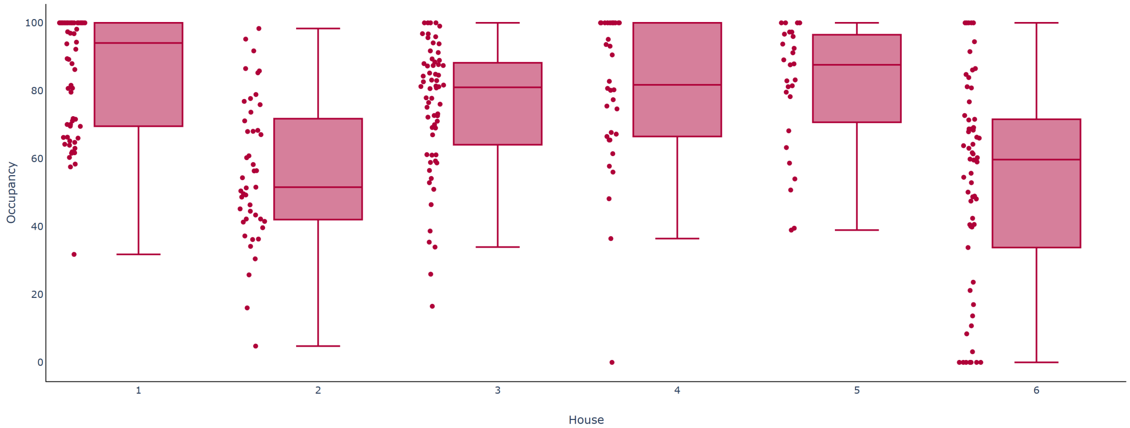

Figure 3.1: Daily occupancy rate for each measured day by house [%]

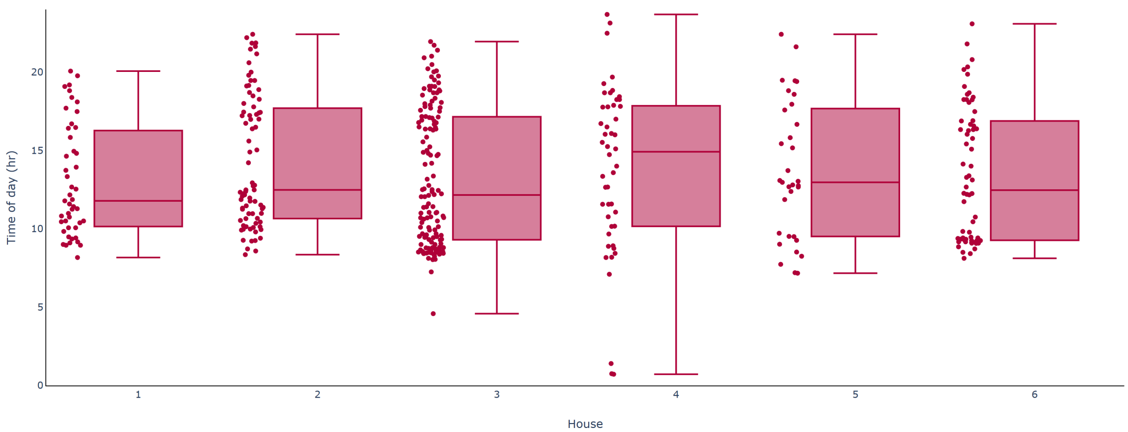

Figure 3.2: Distribution of departure times for each residence [hr]

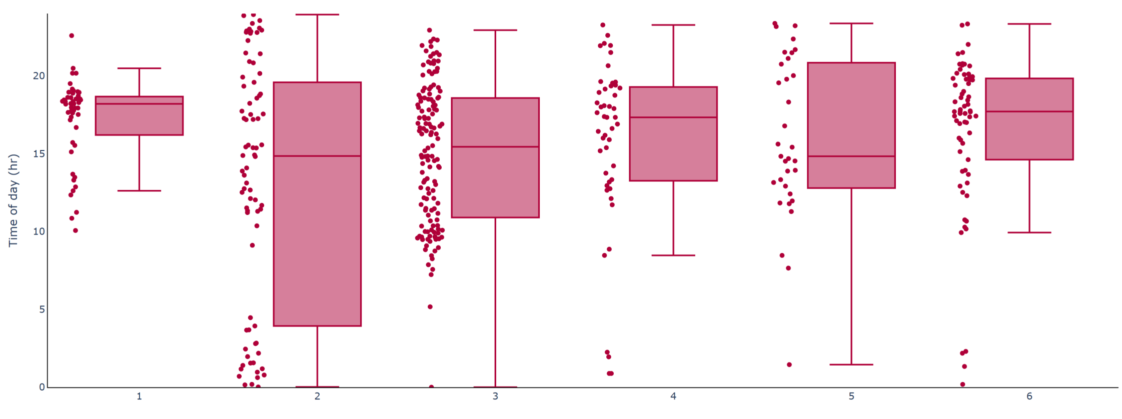

Figure 3.3: Distribution of arrival times for each residence [hr]

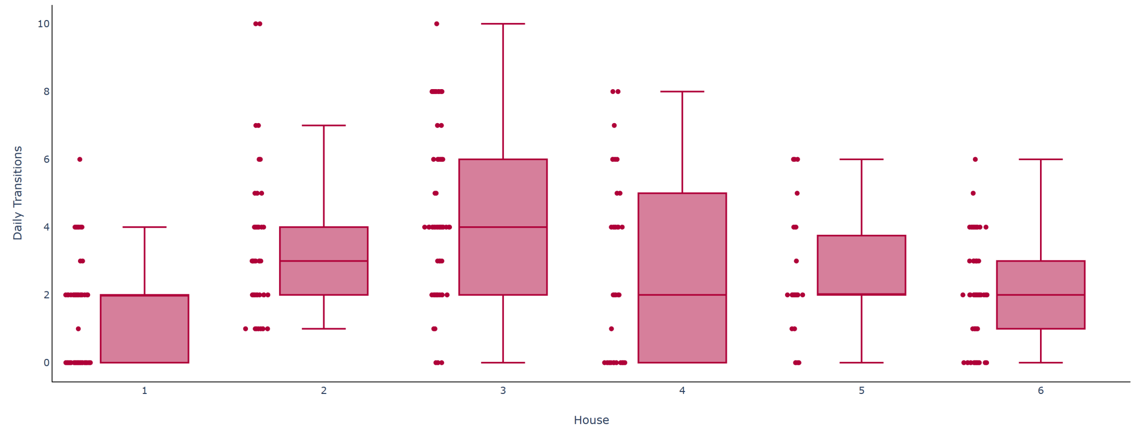

Figure 3.4: Number of daily transitions for each house [count]

3.2 Occupancy Model Generation

Based on the literature review, a non-probabilistic model was chosen to model occupancy. Because non-probabilistic models use past data to create an occupancy probability, the model can be optimized by establishing what training data to include. Model optimization was done by splitting the collected occupancy data in half. The first half of the measured data was used to create the model, except in the case of moving training mode which used a receding horizon. The second half was used as ground truth data and compared to the resulting models to analyze model performance.

Ninety-six non-probabilistic models were created for each house. Each model used a different subset of training data to create occupancy probabilities. The data used was determined by four parameters. For each parameter, two to four values were available. Table 3.2 shows the different values used to create the models.

The four group parameters are defined as follows:

- Day Categorization: this determines how each day of the week is categorized. For example, in “day of the week”, only training data that matches the day being predicted is used. In “week/end” day categorization, all Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays are used to predict occupancy for weekdays.

Table 3.2: Variables used for model creation

Day categorization	Training time	Training Mode	Time resolution
day of week	1	fixed	1
week/end	2	moving	5
m, f, week, end	3		15
	4		60

- Training time: this determines how many weeks are used. This can range from 1 week to 4 weeks. In residences where only 4 weeks of total occupancy data were recorded, only 2 weeks of training data was available.

- Training mode: this determines whether the training time is used in a fixed mode (static training set) or in a moving mode (where a trailing horizon is used). For example, in a 1 week moving mode, only the last seven days are used to predict occupancy for that day.

- Time resolution: this determines how often occupancy is sampled. Time is shown in minutes.

After input data was reduced to a training set, the average occupancy was determined for each interval of the day. This resulted in a occupancy probability between 0 and 1. To create a binary occupancy schedule, a threshold was set for each day. The threshold was determined by finding what threshold value produced the same minutes of predicted occupancy as the summed occupancy minutes in the input data. An example of a single day for House 1 is shown in Figure 3.5. Occupancy prediction models were created for a time period of two to five weeks, depending on the length of total measured data.

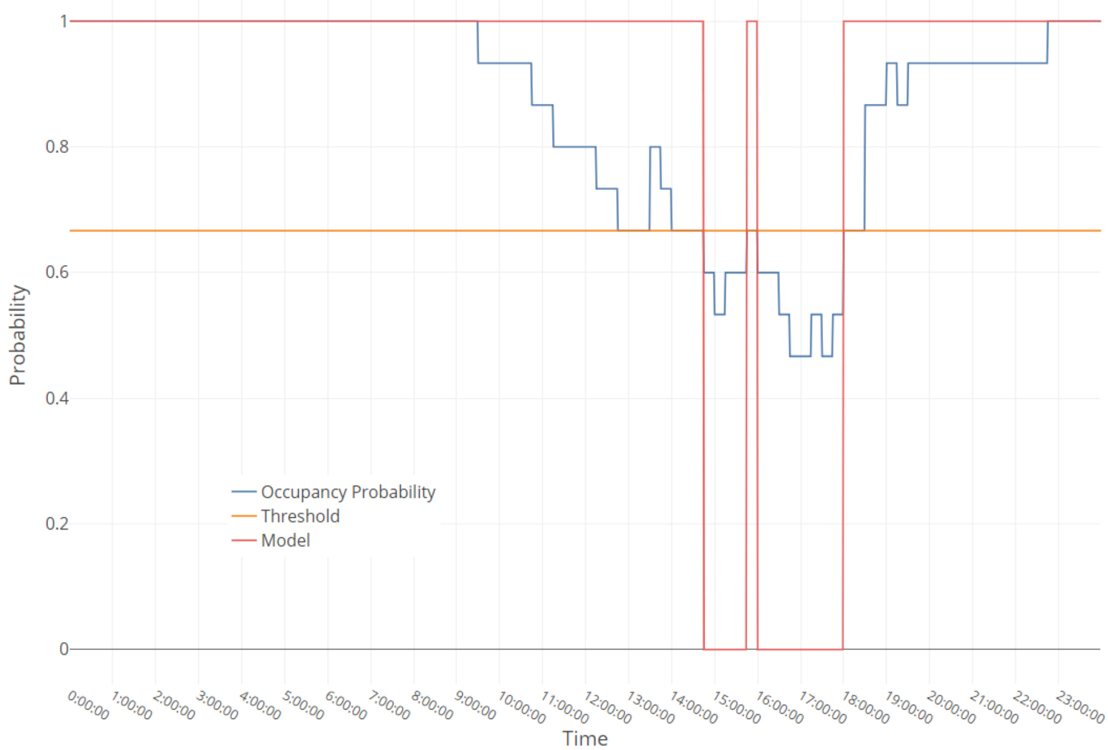

Figure 3.5: Occupancy probability, threshold, and resulting model for single day, by time

3.3 Occupancy Model Accuracy

Once the non-probabilistic models were generated for each home, the resulting predicted occupancy state was compared against the actual occupancy state. To evaluate the effectiveness of the occupancy prediction models, three metrics were used:

- False Negative Rate: Percentage of minutes that model incorrectly predicted house was vacant when it was occupied
- False Positive Rate: Percentage of minutes that model incorrectly predicted house was occupied when it was vacant
- State Matching Error: Percentage of minutes that model incorrectly predicted occupancy. This is the inaccuracy of the model. The state matching error is the sum of the false negative and false positive rate.

All the metrics are error rates, so the best models had low values for false negative rate, false positive rate, and state-matching error. False negative errors and false positive errors have two different impacts. When a false negative occurs, the house is occupied when predicted vacant. Indoor temperature may not be in the comfort range because the control has been changed to a setback temperature. When a false positive occurs, the house is vacant when predicted occupied. In this case, the HVAC system may be running to maintain a tighter temperature setpoint range, resulting in higher energy use.

Table 3.3: Summary of model training and evaluation for each house

House #	Number of models created	Training period (weeks)	Evaluation period (weeks)
1	96	4	5.1
2	96	4	2.4
3	96	4	6.1
4	48	2	2.1
5	48	2	1.9
6	96	4	5.3

Training and evaluation times varied by house due to differences in data collection periods. Table 3.3 lists how the collected data was used generating the prediction models. Training period shows how many weeks were set aside for determining occupancy probability. Evaluation period shows how many weeks were used to evaluate the generated model. During the evaluation period, the predicted occupancy was compared against the measured occupancy at one-minute intervals. For houses with less data, a two-week training period was used. In these cases, 48 models were generated because the 3-week and 4-week training time was unavailable.

The model configuration with the lowest state matching error is shown for each house in Table 3.4. House 5 had the lowest state matching error with an error rate of 8%. House 6 had the highest at 35%. Each house, due to its occupancy pattern, had a different optimal occupancy model. This indicates the value in tailoring the model to the specific use case.

To understand the effect of each parameter on the resulting prediction model, the compiled results were compiled and analyzed. Figures 3.6 to 3.11 show the distribution of state-matching error by each parameter. Figure 3.6 shows the error by house. Results show that each house, with

Table 3.4: Best occupancy prediction model for each residence

House	Day Catego-rization	Training Time	Training Mode	Time Resolu-tion	False Negative	False Positive	State Matching Error
1	mfweekend	4	fixed	15	12%	4%	16%
2	weekend	3	moving	5	13%	13%	26%
3	weekend	1	moving	1	28%	6%	35%
4	weekend	1	moving	60	8%	2%	10%
5	weekend	1	moving	60	5%	3%	8%
6	dayofweek	2	moving	15	7%	30%	37%

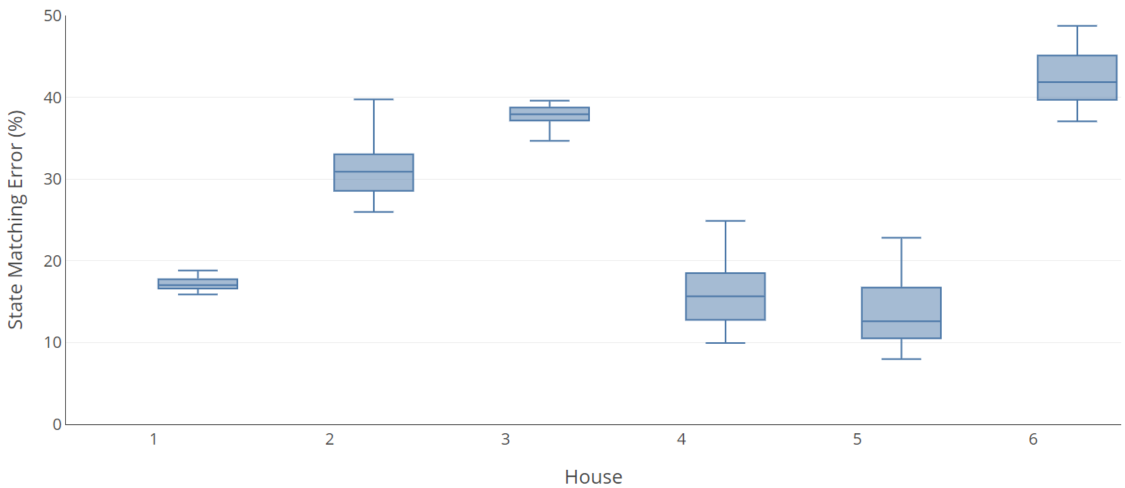

Figure 3.6: State matching error (%), by house

its different occupancy patterns, has a strong effect on the effectiveness of creating a prediction model. People have the largest variance and can drastically affect prediction models. Like previous studies, which show the range of human behavior, these result indicate human behavior varies widely.

Due to the strong impact of occupant behaviour, the results of the four parameters are shown by house to remove its large influence. Figure 3.7 shows the effect of day categorization on the occupancy prediction models. For Houses 2, 4, and 5, using days of the week had the least accurate prediction models. This is likely due to the limited training data resulting in only using the same day of the week. For House 6, where the occupant had a part time job two days a week, the day of the week method increased the prediction accuracy.

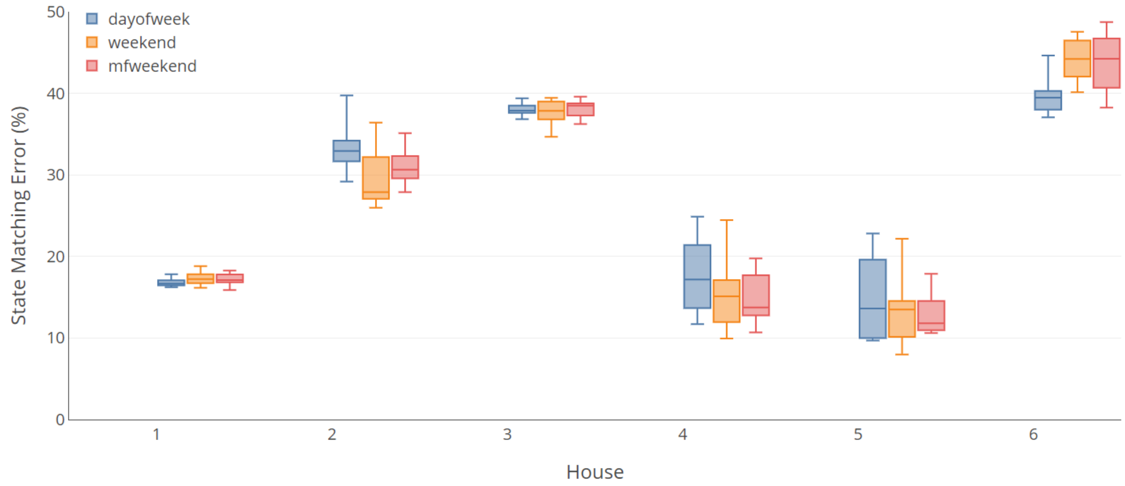

Figure 3.7: State matching error (%), by house and day categorization

Figure 3.8 shows the effect of training time on the occupancy prediction models. For five of the six houses, using more training data improved the accuracy of models. House 3, the exception, had a shift in occupancy halfway through data collection. This indicates that when new occupants join a household, the previous training data will not effectively predict the new occupancy pattern. To explore this theory further, Figure 3.9 shows the resulting state-matching error when the training time is extended to seven weeks. With additional training weeks, the error is reduced, indicating that the longer the training data is accumulated, the more error can be reduced.

Figure 3.10 shows the effect of using either a fixed or moving training mode. In this case, the moving training mode had improved prediction accuracy for five of the six houses. By allowing the model to adjust over time, the moving training mode adapts to shifting behavior.

Lastly, Figure 3.11 shows the effect of time resolution on the six houses. Results are nearly identical for each time resolution variable signaling that sampling time does not play a large role in reducing prediction accuracy.

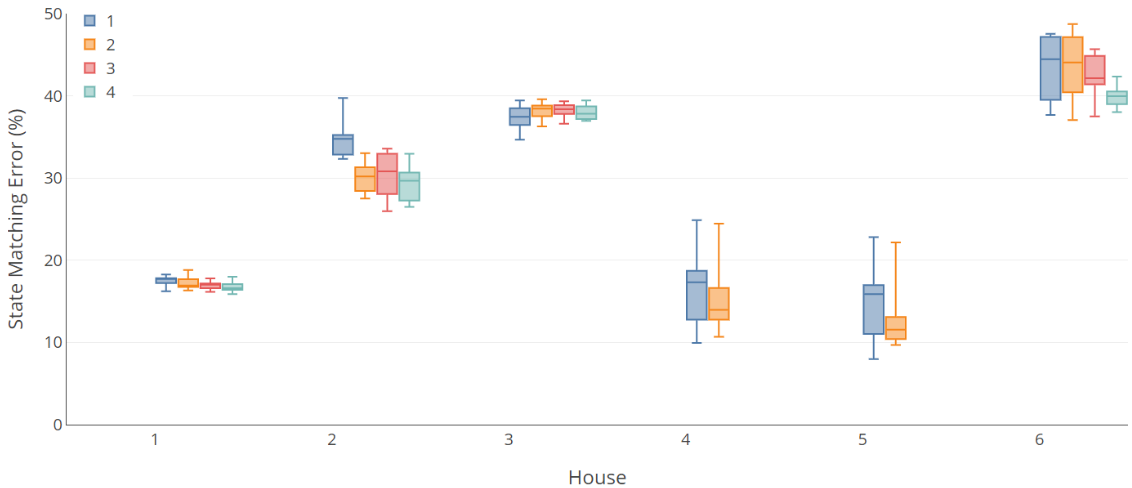

Figure 3.8: State matching error (%), by house and training time

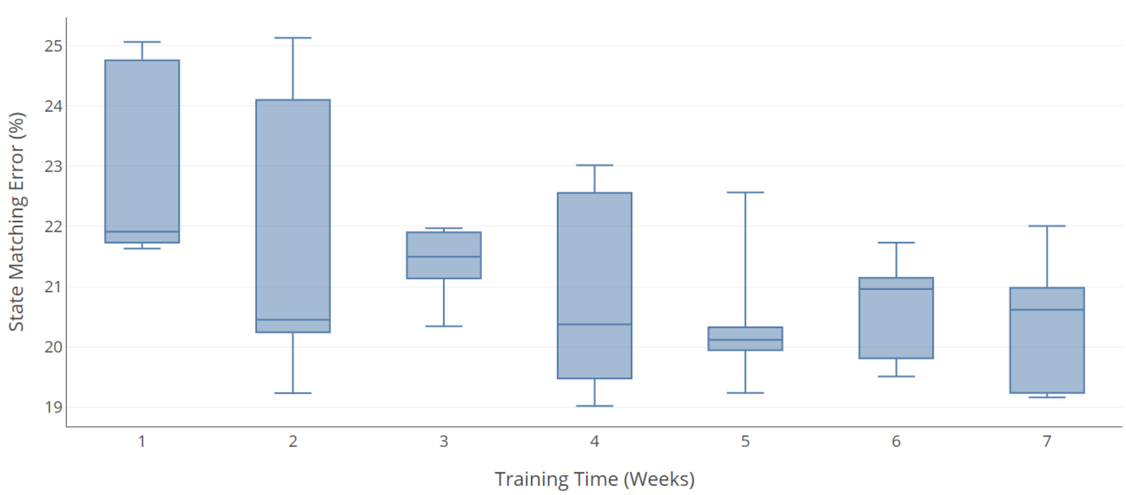

Figure 3.9: State matching error (%) for House 1, by week

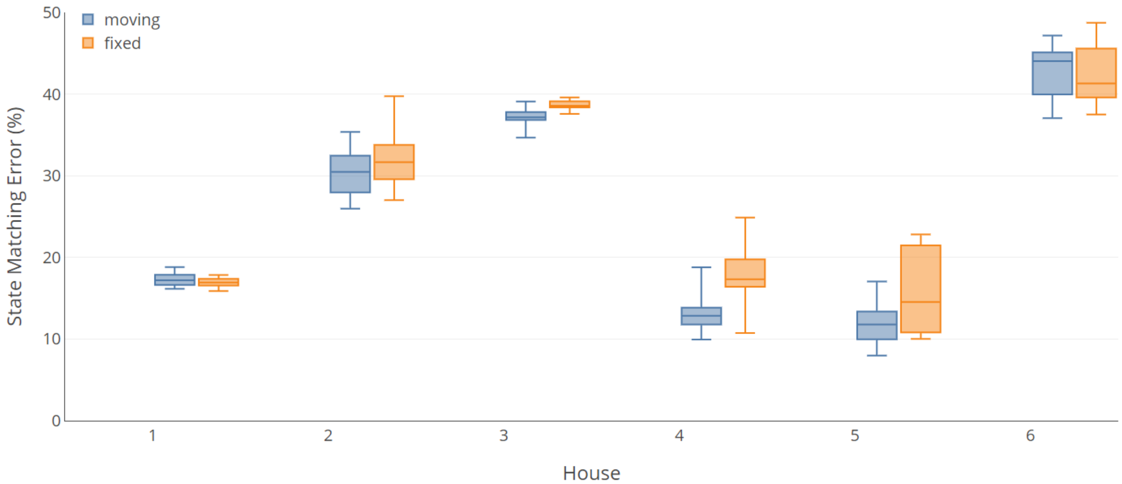

Figure 3.10: State matching error (%), by house and training mode

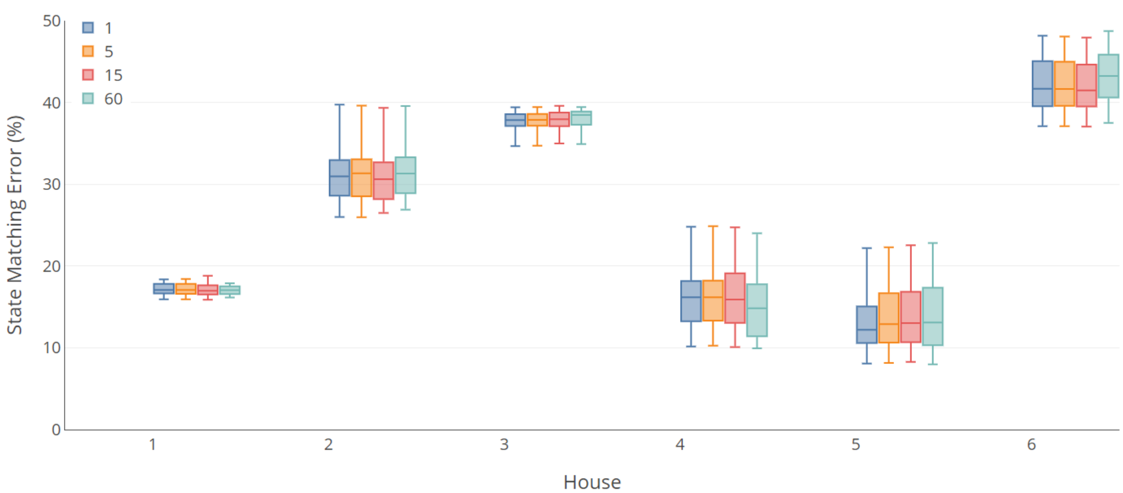

Figure 3.11: State matching error (%), by house and time resolution

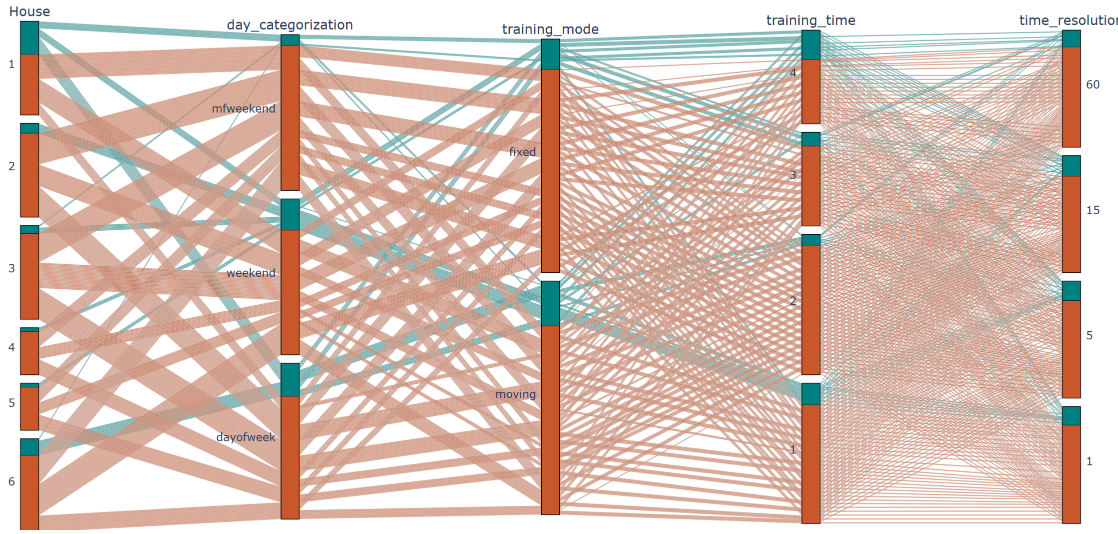

Figure 3.12: Parallel category plot of occupancy models

While optimizing the non-probabilistic occupancy prediction model to the specific house will reduce the state matching error, a universally applicable model would be able to easily deploy in houses without the need for individual setup. To find the best occupancy model for all the houses surveyed, the state matching error results were normalized by dividing the results by the lowest state-matching error achieved by that house. Using the lowest error for each house was done to ensure that all of the houses were considered equally. Figure 3.12 shows a parallel category plot of the results. Occupancy models that were within 5% of the best model for that house are shown in green, higher values are red.

Results show that the "mfweekend" day categorization does not work well for the houses surveyed. "Day of the week" and "weekend" both work well, with "weekend" performing better for most houses, suggesting the six houses survey do not have distinct Monday or Friday schedules. Four and one week training times were the best training periods, with two and three week training periods performing slightly worse. Moving training mode produced the highest number of low state matching errors. This is likely due to its ability to adjust to occupants over time. Time resolution produced equivalent results, with the 15 minute time resolution yielding a slightly higher number of low normalized state matching errors. Based on these results, the optimal universal model across all houses was a 1 week, moving model using weekend categorization and a 15 minute time interval. The state matching error for each house with the universal prediction model is shown in Table 3.5.

Table 3.5: Results of universal occupancy prediction model for each residence

House	Day Categorization	Training Time	Training Mode	Time Resolution	False Negative	False Positive	State Matching Error
1	weekend	1	moving	15	15%	2%	17%
2	weekend	1	moving	15	22%	10%	32%
3	weekend	1	moving	15	28%	6%	35%
4	weekend	1	moving	15	8%	2%	10%
5	weekend	1	moving	15	5%	3%	8%
6	weekend	1	moving	15	11%	33%	44%

Chapter 4

Building Simulation Evaluation and Discussion

Building performance simulations were conducted in EnergyPlus Version 9.1 to understand the impact of residential HVAC control. The goal of the building simulations was to determine the energy use and zone temperatures for a number of representative home scenarios. Multiple homes, climates, seasons, and occupancy patterns were used to understand the range of possible outcomes and more globally represent the range of possible outcomes.

4.1 Building Performance Simulation Settings

The six occupancy datasets from Chapter 3 were used as the six possible occupancy scenarios. Data from these homes was assumed to represent occupancy patterns for all climates and seasons the simulations were performed.

Five prototype homes provided by NREL were used for building models [48]. Each home had a different climate and building construction that was representative of national housing stock. The five locations used were Boston, Phoenix, Atlanta, Seattle, and Houston. House sizes averaged 2,000 ft^2 with typical home construction and vintages for each region were used (Table 4.1).

Occupancy prediction models developed in Section 3.3 ranged from 13 to 36 days depending on the house. A two-week period with one-minute timestep intervals was used for the building performance simulations. In houses where more data was available, the first two weeks were used. House 5, which only had 13 days, used the first day at the end to create a two-week period. Building simulations were performed twice, for the first two weeks of January and for July. Winter

Table 4.1: Summary of house model constructions

	Boston, MA	**Phoenix, AZ**	**Atlanta, GA**	**Seattle, WA**	**Houston, TX**
Climate	Cold 5A	Hot-Dry 2B	Mixed-Humid 3A	Marine 4C	Hot-Humid 2A
Vintage	<1950s	1970s	1970s	<1950s	1970s
House Size	2589 ft^2	2013 ft^2	2013 ft^2	1938 ft^2	2013 ft^2
Envelope					
Attic	Uninsulated	Ceiling R-13, Vented	Ceiling R-19, Vented	Ceiling R-13, Vented	Ceiling R-13, Vented
Wall Cavity	Uninsulated	Uninsulated	Uninsulated	Uninsulated	Uninsulated
Foundation	Uninsulated	Uninsulated	Uninsulated	Uninsulated	Uninsulated
Windows	Clear, Double, NM, Air	Clear, Double, Metal, Air	Clear, Single, Metal	Clear, Single, Metal	Clear, Double, Metal, Air
Air Leakage	15 ACH50	15 ACH50	15 ACH50	15 ACH50	15 ACH50
HVAC					
Heating	Gas Boiler, 80% AFUE	Gas Furnace, 80% AFUE	Gas Furnace, 80% AFUE	Gas Furnace, 80% AFUE	Gas Furnace, 80% AFUE
Cooling	Room AC, EER 10.7	Central, SEER 13	Central, SEER 13	None	Central, SEER 13

and summer runs allowed the impact of HVAC control to be ascertained for both heating and cooling modes.

Figures 4.1 to 4.5 show the outdoor air temperature during the simulation periods. Boston and Atlanta have cold winters, near freezing, and hot summers. Outdoor air temperature is not within the comfort range for the majority of the simulation. Houston has moderately cold winters and hot summers. Phoenix has mild winters and extremely hot summers. During the Phoenix winter, the outdoor air temperature oscillates within the comfort range. Seattle experiences cool winters and summers, requiring some heating year-round.

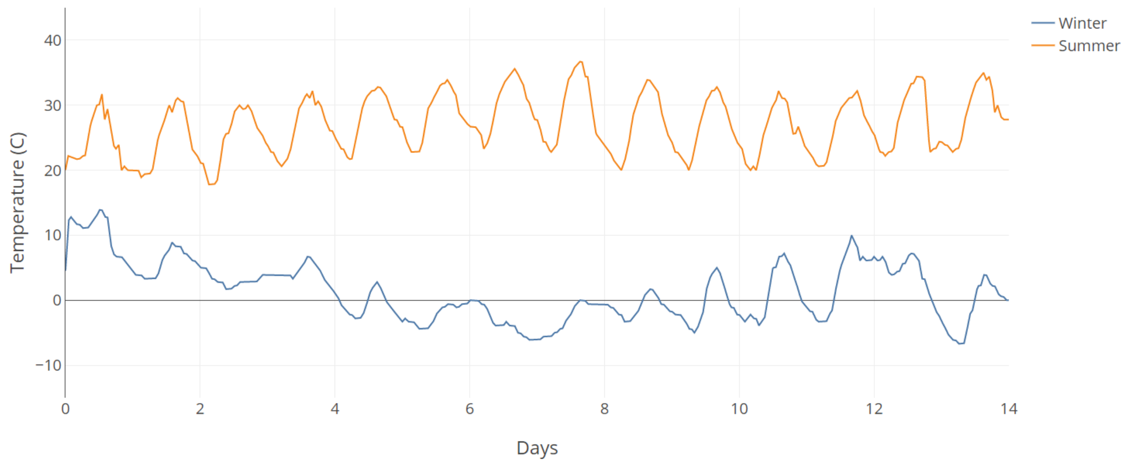

Figure 4.1: Outdoor temperatures during simulation runs for Atlanta, GA [°C]

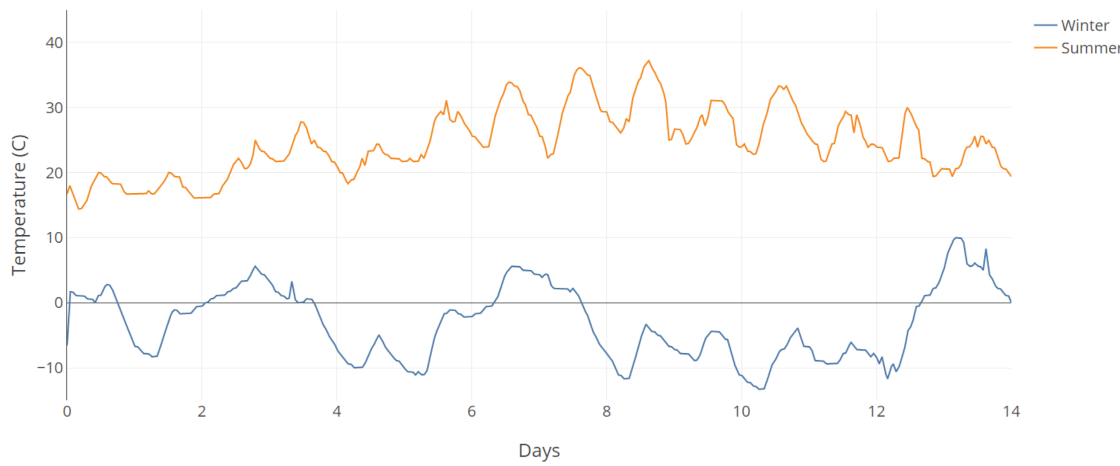

Figure 4.2: Outdoor temperatures during simulation runs for Boston, MA [°C]

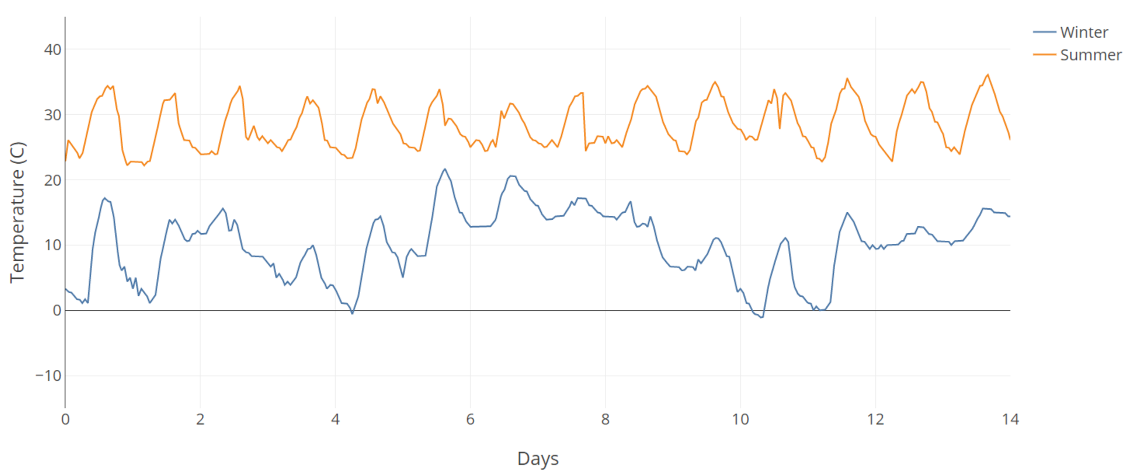

Figure 4.3: Outdoor temperatures during simulation runs for Houston, TX [°C]

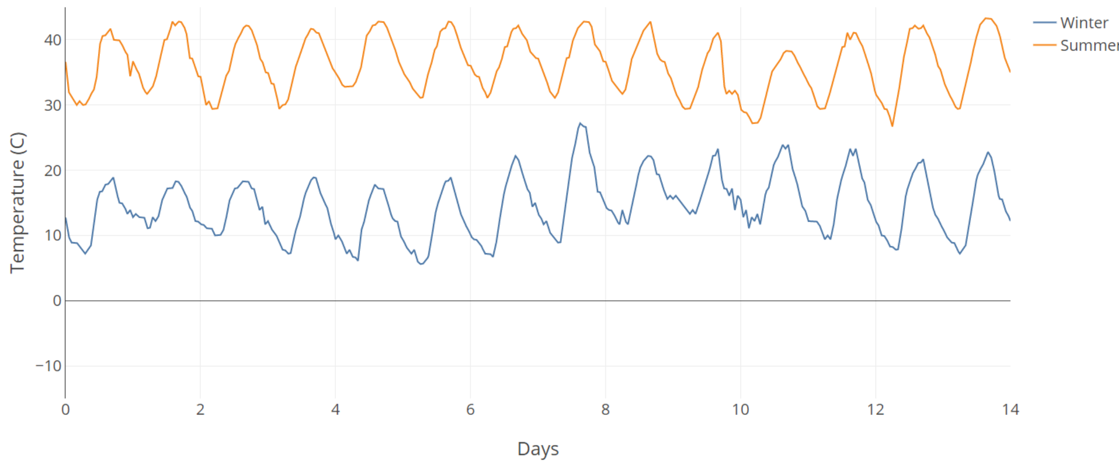

Figure 4.4: Outdoor temperatures during simulation runs for Phoenix, AZ [°C]

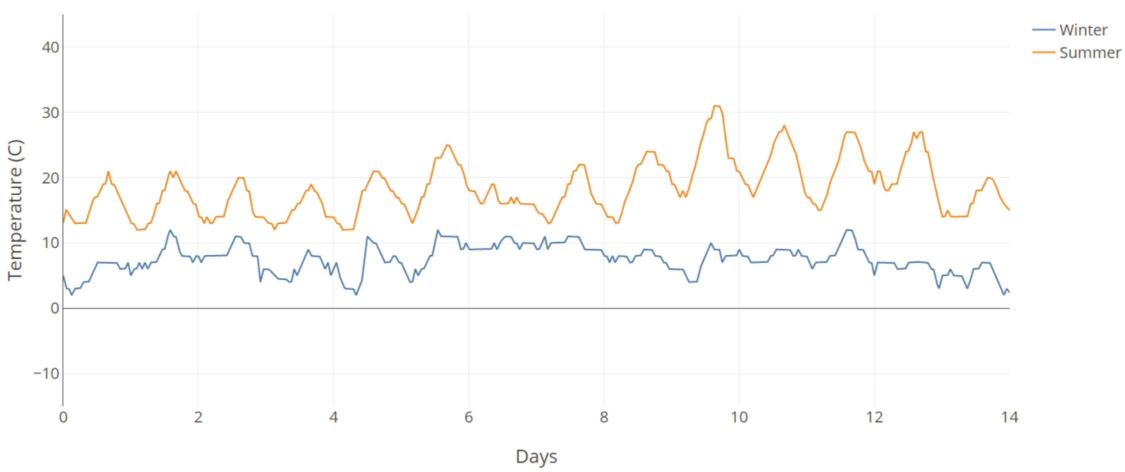

Figure 4.5: Outdoor temperatures during simulation runs for Seattle, WA [°C]

Heating and cooling setpoints were established using the predicted mean vote. ISO EN 7730 establishes three comfort categories using operative temperature. Categories are shown in Table 4.2. Class A and B were used as the defined comfort range to maintain a PMV below ± 0.5.

Table 4.2: Three categories of thermal comfort (ISO EN 7730, 2005)

Category	Thermal state of the body as a whole		Operative temperature °C	
	PPD %	PMV	Summer (0.5 clo)	Winter(1 clo)
A	<6	-0.2 <PMV <+ 0.2	23.5 - 25.5	21.0 - 23.0
B	<10	-0.5 <PMV <+ 0.5	23.0 - 26.0	20.0 - 24.0
C	<15	-0.7 <PMV <+ 0.7	22.0 - 27.0	19.0 - 25.0

While operative temperature or PMV is used to define thermal comfort, HVAC systems are controlled using ambient air temperatures. For the building simulations, ambient air temperature setpoints were used to control the heating and cooling in the homes. Setpoint temperatures used are summarized in Table 4.3.

Table 4.3: Setpoint and setback temperatures used in building simulation

	Setpoint Temperature °C	Setback Temperature °C
Heating	22	18
Cooling	24.5	28

Three HVAC control scenarios were modeled in the building performance software:

- conventional setpoint operation
- occupancy-based HVAC control
- occupancy-based MPC

Conventional operation, which used a constant heating and cooling setpoint, was used as the baseline. Occupancy-based HVAC control and occupancy-based MPC were simulated and compared to the baseline. Results for each strategy will be discussed in the subsequent sections.

4.2 Conventional Control Results

Each home and season was run with conventional control. Boston winters used the highest amount of energy during the simulation period, likely due to the cold environment, older vintage house, and larger size.

Beyond energy use, the indoor operative and ambient air temperature was simulated to evaluate thermal comfort. Comfort was defined using operative temperature. Temperatures between 20 °C and 26 °C were considered comfortable to accommodate 0.5 to 1.0 clo clothing. Comfort using conventional HVAC control is shown in Figure 4.6. With a constant setpoint temperature within the comfort range, the percentage of time within the comfort range was expected to be high. As seen in the graph, comfort was high except for Phoenix and Seattle in summer.

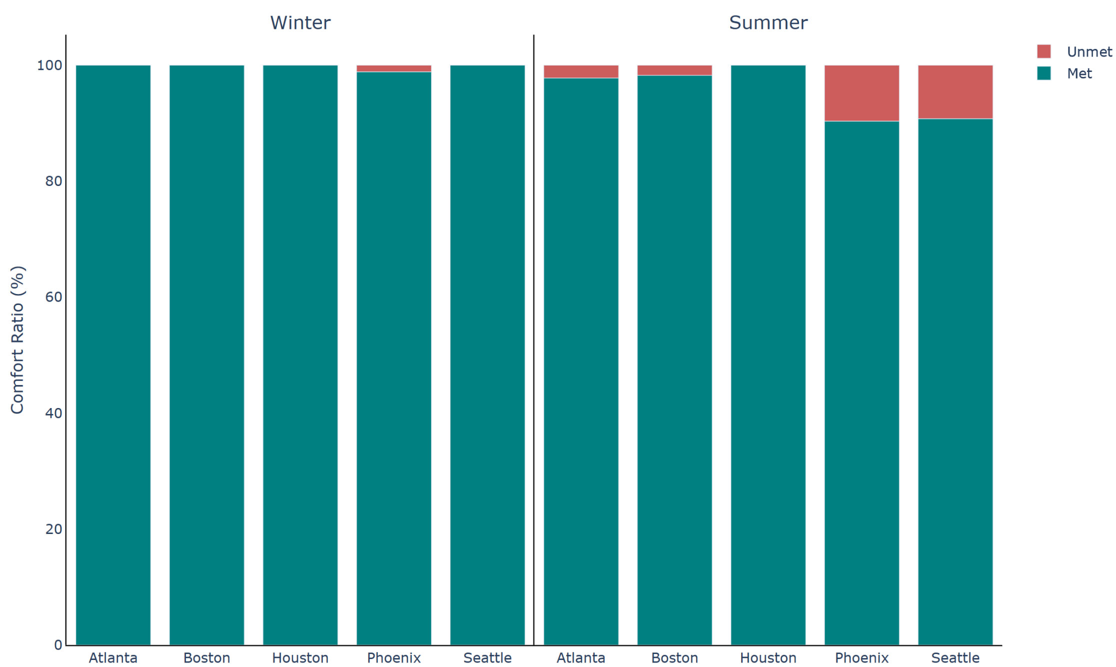

Figure 4.6: Comfort percent for simulation period, by city and season [%]

Further examination of the Phoenix and Seattle summer results can be seen in Figures 4.7 and 4.8. Phoenix achieved a constant ambient air temperature of 24.5 °C. However, the operative temperature was higher, creating uncomfortable temperatures. This case demonstrates the impact of using air temperature rather than operative temperature. In Seattle, both the indoor air and operative temperature were too high for comfort. High temperatures were caused by the hot outdoor temperatures which could not be mitigated without an air conditioner.

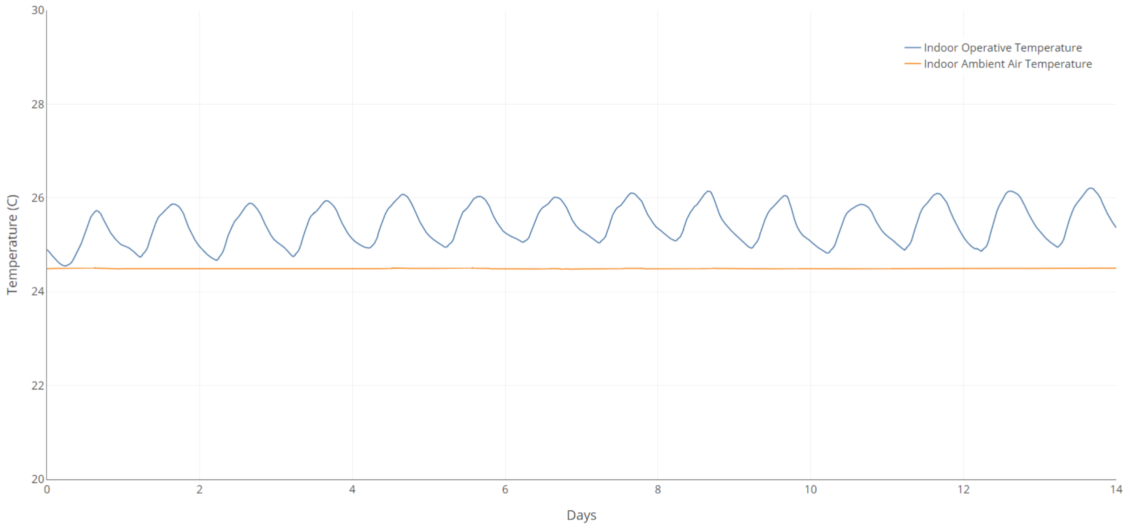

Figure 4.7: Comparison of operative temperature and ambient indoor air temperature for Phoenix summer [°C]

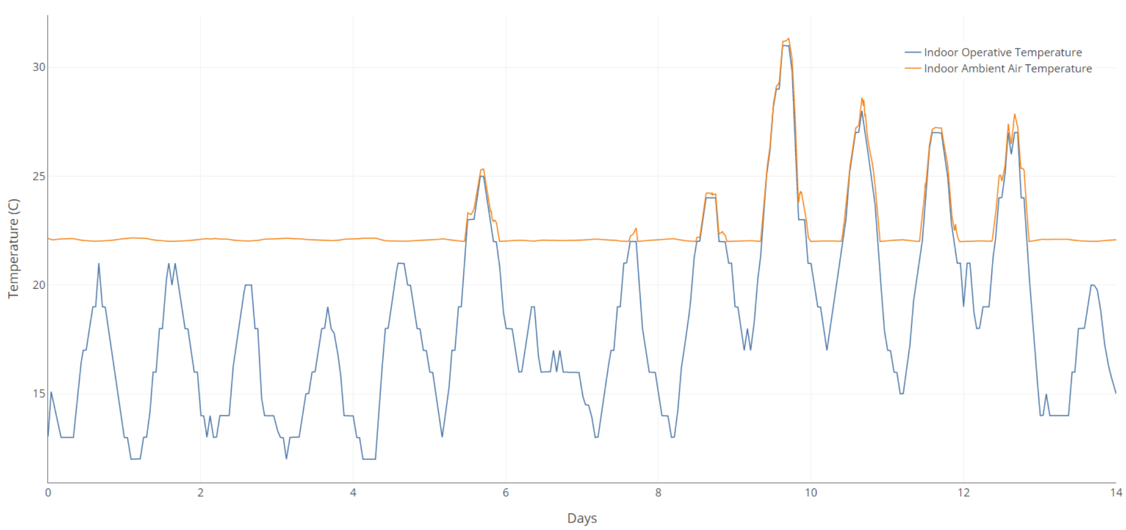

Figure 4.8: Comparison of outdoor temperature and ambient indoor air temperature for Seattle summer [°C]

4.3 Occupancy-based HVAC Control

Occupancy can be incorporated into HVAC control in a multitude of ways. For this study, occupancy is used to set the temperature for the HVAC control. In all cases simulated, either prediction or detection of occupancy is used to establish the setpoint temperature. When the space is vacant, the temperature is allowed to drift to a more relaxed setback temperature, reducing use of the HVAC system when unoccupied. When occupied, the temperature is kept at the setpoint, which maintains temperatures within the comfort zone. Three main strategies were considered and simulated:

(1) Reactive: Occupancy is detected and setpoint temperatures are adjusted accordingly. In this case, occupancy is sensed and no prediction is used.

(2) Predictive: Occupancy is predicted using the non-probabilistic models developed in Section 3.3

- Universal model - this prediction model performed best for all houses. Model used a 1 week, 15 minute, weekend/weekday categorization, moving training set.
- Individually tuned - this prediction model performed best for the specific house. Models used are listed in Table 3.4.

(3) Hybrid: A hybrid of predictive and reactive occupancy models. Occupancy is predicted using the non-probabilistic models developed in Section 3.3. If an occupancy change from vacant to occupied is detected that wasn't predicted, the control will react and reset the temperature control to occupied settings. The control method does not react to changes from occupied to vacant states to maintain the predictive aspect, keeping it from becoming a purely reactive control.

- Universal hybrid - this prediction model performed best for all houses. Model used a 1 week, 15 minute, weekend/weekday categorization, moving training set.

Table 4.4: Average energy savings by control method

Control Method	Energy Savings
Reactive	9.1%
Universal Model	10.9%
Individually Tuned	9.6%
Universal Hybrid	4.3%
Individual Hybrid	5.7%

- Individual hybrid - this prediction model performed best for the specific house. Models used are listed in Table 3.4.

Building simulations were conducted for each occupancy-aware control method for all homes, climates, and seasons. Energy savings of each method relative to the conventional baseline model is shown in Figure 4.9 with the average savings show in Table 4.4. All control strategies reduce the total energy used during the simulation. The two prediction models had the largest energy savings potential of the five control methods, with 10.9% and 9.6% savings respectively. Reactive control has similar energy savings to the prediction models, with an average of 9% of total energy consumption. This method was particularly helpful in homes where long periods of unpredictable vacancy such as vacation occurred. The hybrid approach, which used both predictive and reactive occupancy, saved the least energy with an average of 3-5% savings.

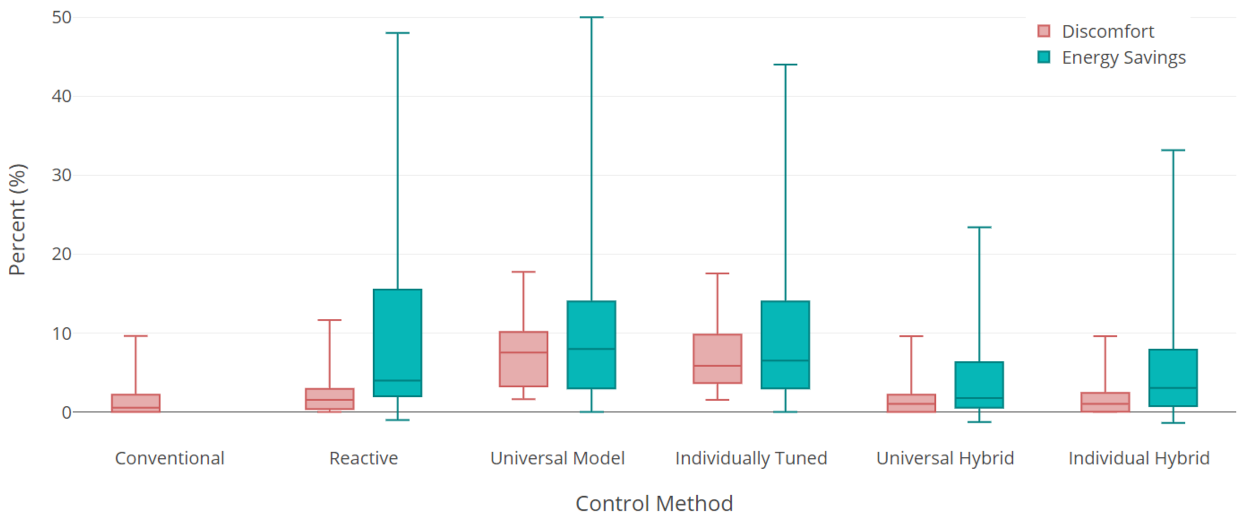

Figure 4.9: Energy savings and discomfort percent for simulation period, by control method [%]

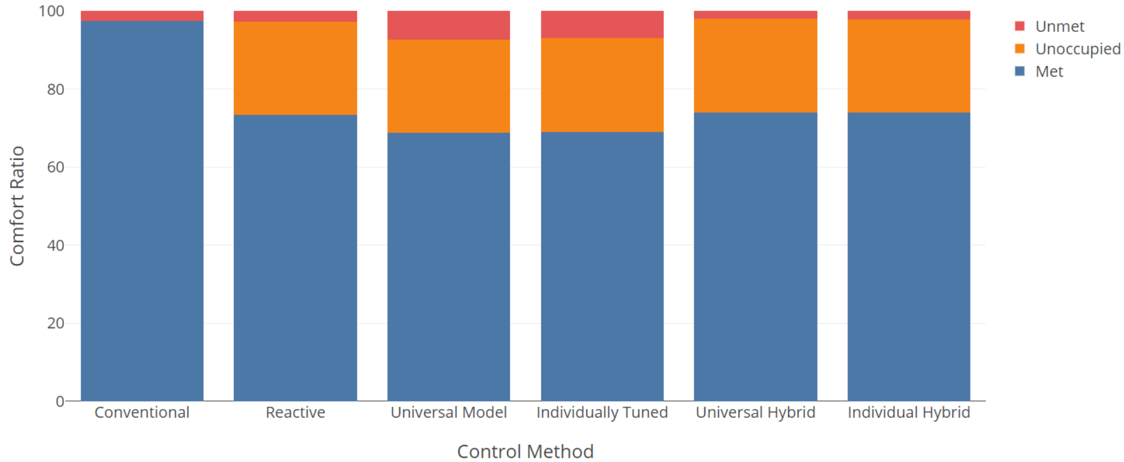

Figure 4.10: Distribution of comfort for simulation period, by control method [%]

In parallel with energy savings, the comfort of the occupant with each control method was assessed using operative temperature. Figure 4.9 shows the percentage of time the occupant was uncomfortable. Discomfort is classified as time spent occupying the residence when it is not within Class A or B comfort temperatures. Distribution of time spent in each comfort category is shown in Figure 4.10 with the three classifications of “Met”, “Unoccupied”, and “Unmet”. “Met” and “unmet” classifications designate times when the residence is occupied, with operative temperature designating the difference between met and unmet comfort. “Unoccupied” classification can include any temperature but does not count as discomfort because there is no occupant to experience it.

Table 4.5: Average discomfort by control method

Control Method	Unmet %
Conventional	2.4%
Reactive	2.6%
Universal Model	7.3%
Individually Tuned	6.9%
Universal Hybrid	2.0%
Individual Hybrid	2.1%

The average percent of time spent in discomfort for each control method is shown in Table 4.5. In this study, the hybrid predictive models achieved the most comfortable indoor conditions.

Time the occupant spent in discomfort was similar to comfort achieved using the conventional setpoint control. The purely predictive models, which did not react to incorrect predictions, led to large discomfort percentages.

Achieving energy savings without disrupting occupant comfort is the goal of effective HVAC control. Because energy savings is achieved by allowing the temperature to drift to uncomfortable conditions, unmet comfort must be optimized in parallel with energy savings. Figure 4.11 compares the discomfort against the energy savings for each simulation. Distribution of results is widely varied for all control strategies, signaling that additional variables beyond control strategies affect comfort and energy savings.

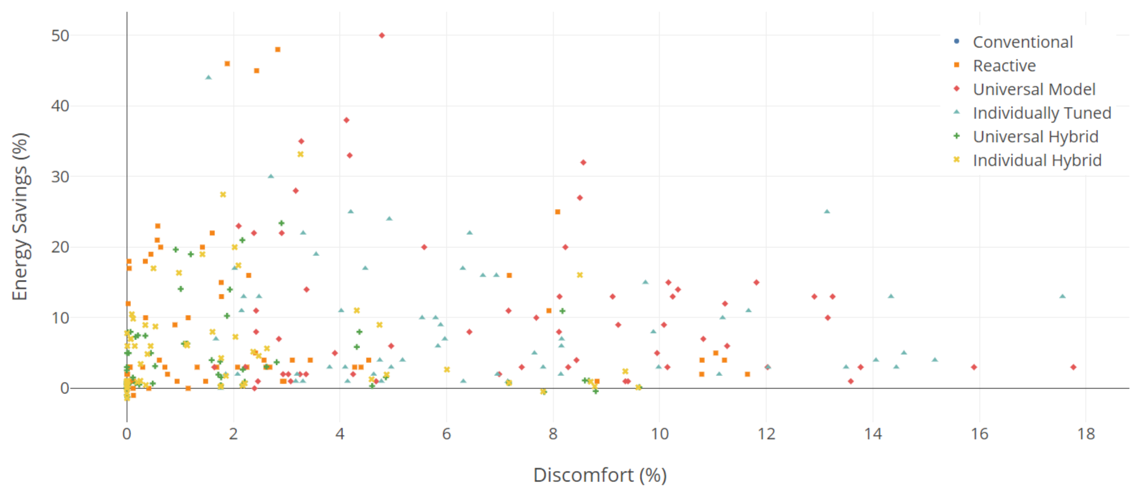

Figure 4.11: Unmet comfort vs. energy savings for simulation period, by control method [%]

To understand which additional variables affected the simulation results, the multivariate response is shown in Figures 4.12 and 4.13. In both figures, the combination of variables that achieved the best results is shown as the green. Less desirable results are shown in red. For energy, city and occupancy patterns also affect savings. For unmet comfort, city, season, and occupancy patterns all affect comfort. The impact of each is these variables is explored in Figures 4.14 to 4.17.

Figure 4.14 shows the energy saving and comfort for each city. Phoenix and Seattle have high unmet comfort ratios as reviewed in the conventional control. Energy savings also vary by city, indicating that the vintage of the home and the climate determines the relative savings potential.

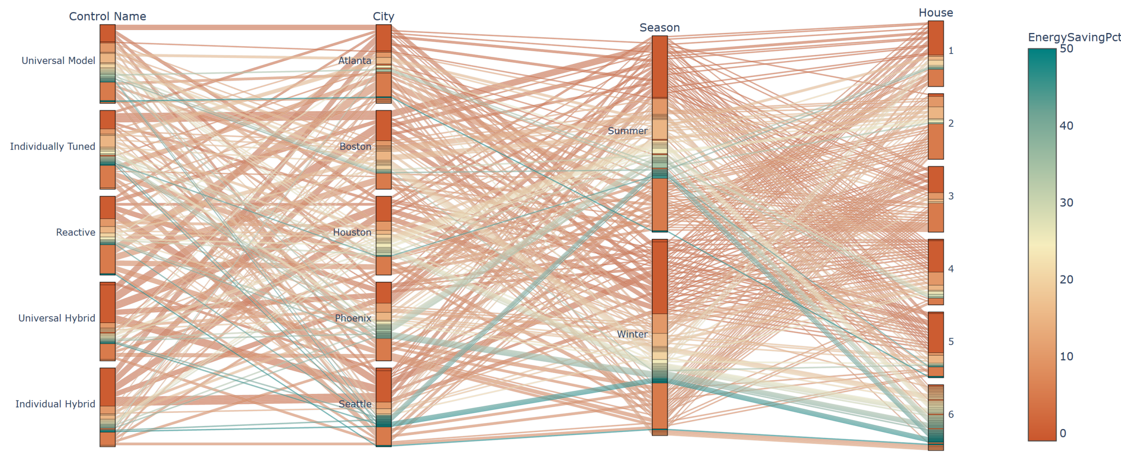

Figure 4.12: Parallel category plot of energy savings

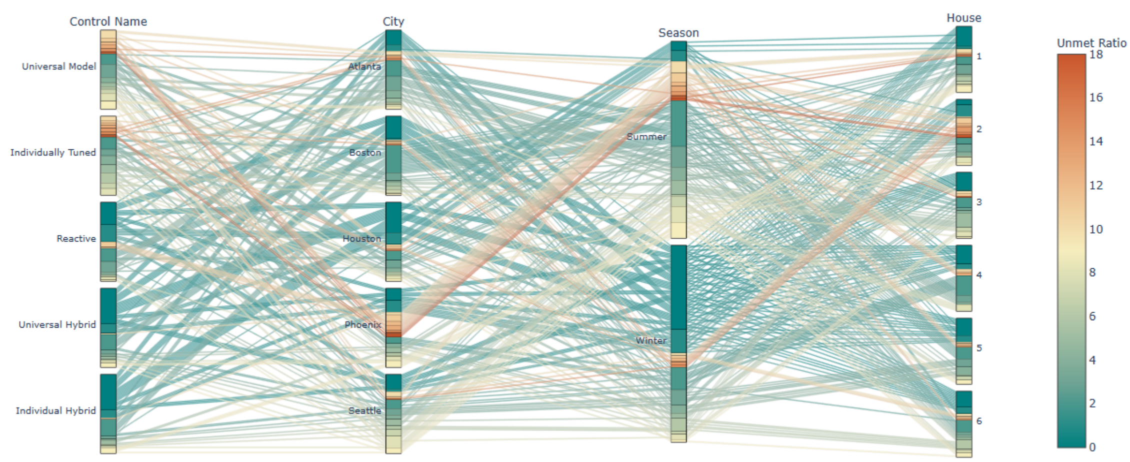

Figure 4.13: Parallel category plot of unmet comfort ratio

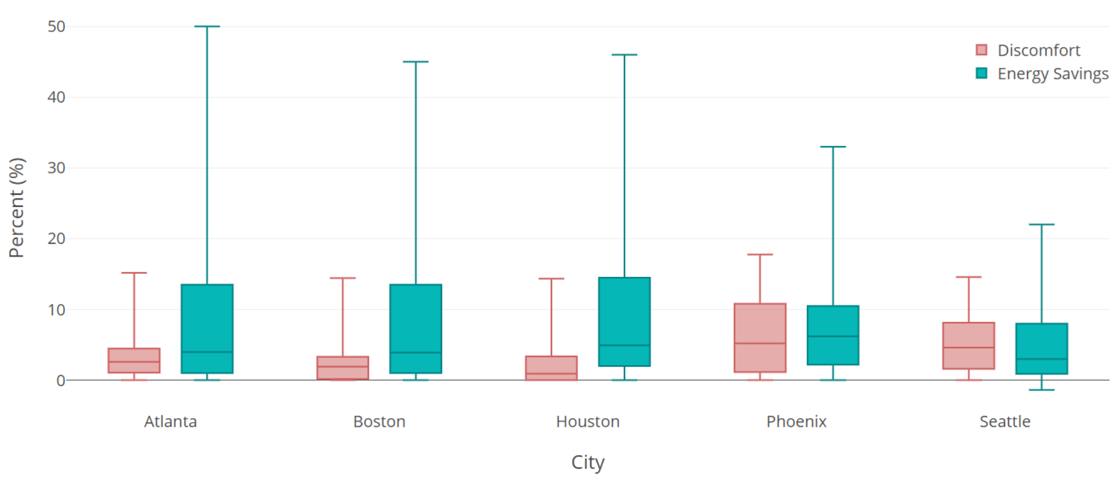

Figure 4.14: Unmet comfort and energy savings for simulation period, by city [%]

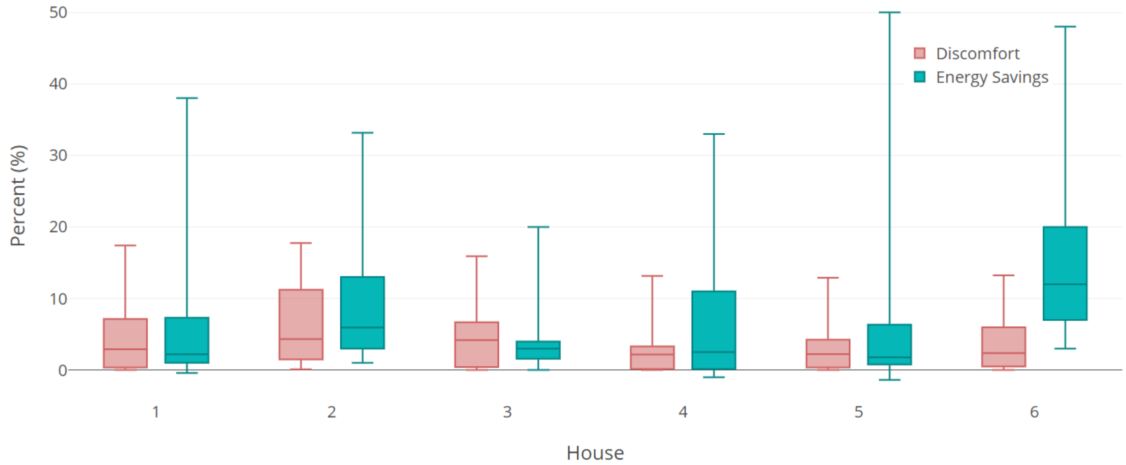

Figure 4.15: Unmet comfort and energy savings for simulation period, by house [%]

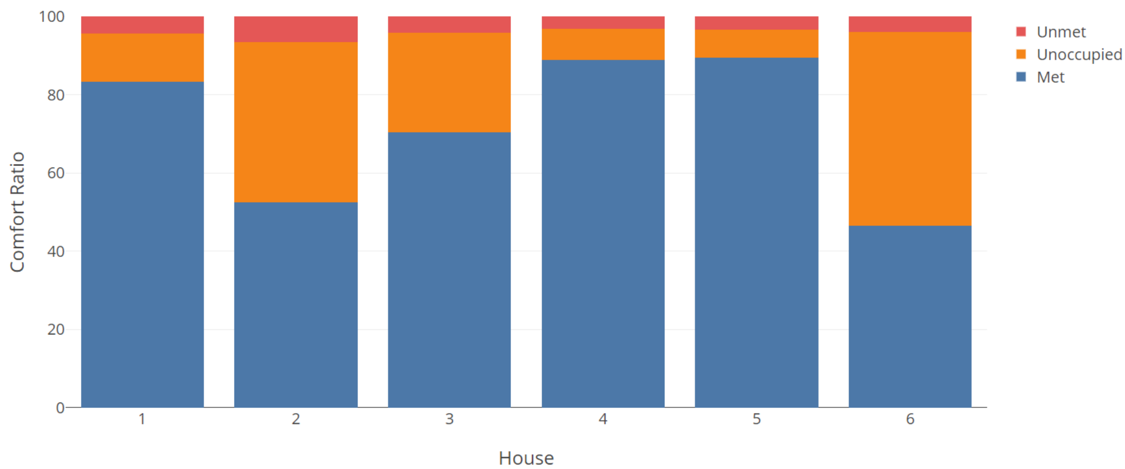

Figure 4.16: Distribution of comfort for simulation period, by house [%]

Figures 4.15 and 4.16 show the effect of the different house occupancy patterns on energy and comfort. Both energy savings potential and discomfort were affected by occupancy patterns. Houses 2 and 6 which had the highest vacancies had the highest median energy savings. This indicates that the higher the vacancy rate, the higher the energy savings potential. In the cases of predictive models, the energy savings potential is also dependent on how well the vacancy is predicted. The savings can only be realized when the house is vacant and it was correctly predicted to be so. Despite low vacancy times, House 5 had the highest energy savings in some simulations. House 5 also had the most accurate prediction model. This indicates that the better the prediction model is,

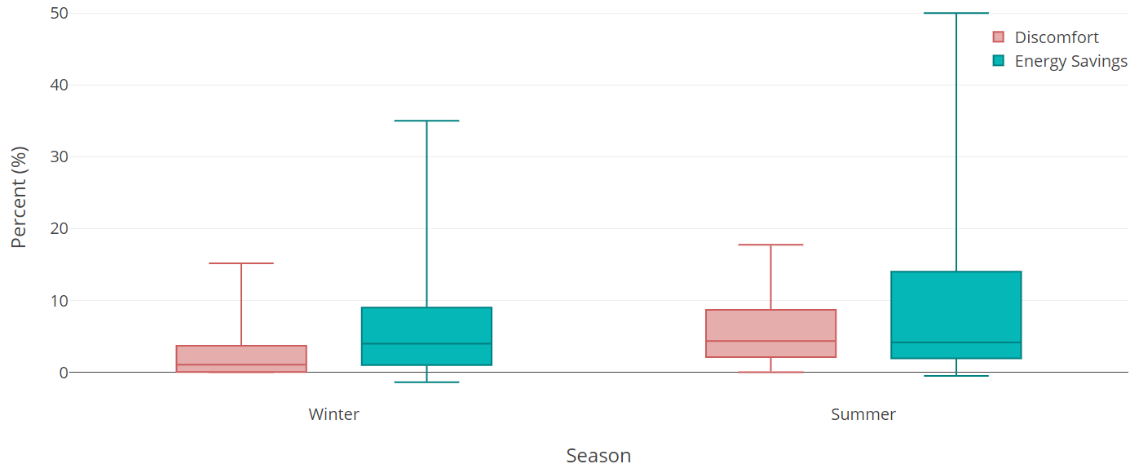

Figure 4.17: Unmet comfort and energy savings for simulation period, by season [%]

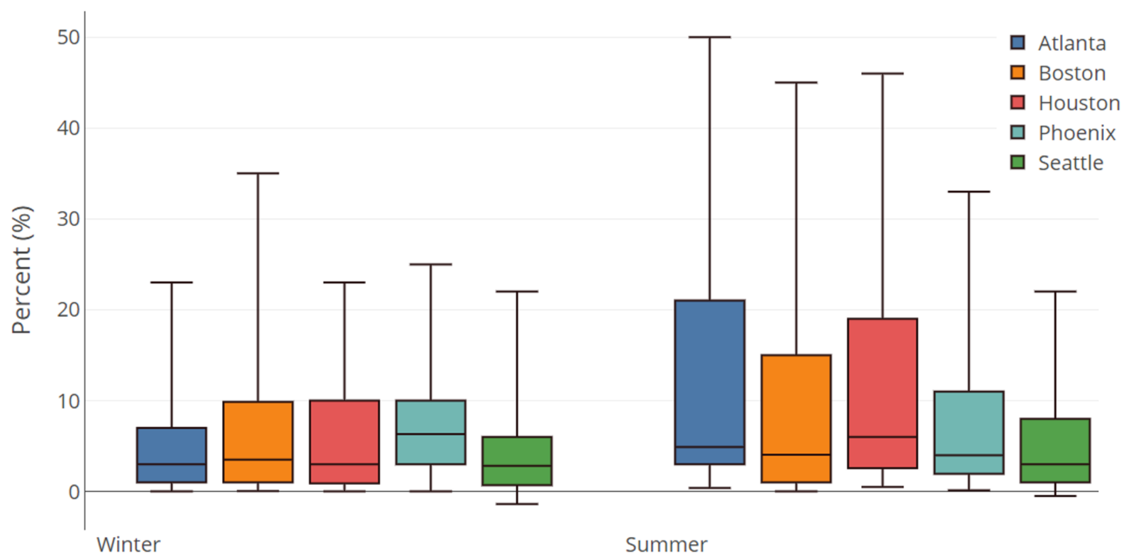

Figure 4.18: Unmet comfort and energy savings for simulation period, by city and season [%]

the higher the possible energy savings can be with time heating and cooling an incorrectly predicted occupied house reduced. Looking at comfort, House 2 had the highest discomfort portion and a high prediction inaccuracy signalling the importance of accurate prediction to comfort. Overall, these figures indicate that although large energy savings are possible, the prediction model needs to be accurately calibrated to achieve energy savings and comfort (Figure 4.16).

In Figure 4.17, the energy and comfort were evaluated by season. Two seasons were simulated:

summer and winter. Summer discomfort is higher due to discomfort in all Phoenix and Seattle summer simulations. Discomfort during winter is low, revealing that when the equipment and thermostat are tuned, comfort is readily achieved.

In Figure 4.18, energy savings and comfort values are shown by city and season simultaneously. Energy savings are highest in Atlanta, Boston, and Houston, especially during summer. All three climates are extremely hot and humid during the summer months so reducing air-conditioning results in large energy benefits.

Due to large ratio of discomfort, prediction-only models were not an effective HVAC control strategy. Both hybrid models were able to achieve comfort at the same level of conventional control. Because occupant comfort is not degraded, these methods are more likely to be used by occupants. Conventional, reactive, and hybrid methods all achieve discomfort below 3% on average. Reactive control has the highest average energy savings at 5.3% but at the cost of reduced comfort in comparison to conventional control. Conventional control, which is the baseline, has no energy savings. The two hybrid controls, universal and individual, achieve an average energy savings of 3.4% and 4.7% respectively. Individual hybrid is able to achieve the highest energy savings while maintaining comfort levels at or below conventional control. Therefore, the individual hybrid control is the recommended occupancy-based HVAC control.

4.4 Model Predictive HVAC Control

Model predictive control (MPC) was the final HVAC control considered. In MPC, an algorithm is used to predict and proactively react to upcoming temperature disturbances or setpoint changes. MPC has been used in the past to optimize a number of parameters in building control, from incorporating weather forecasts for temperature control to shifting peak loads for the power grid [20, 22, 24, 25, 49–60]. MPC has the advantage of being a proactive rather than reactive control strategy. By predicting the effect, for example, of outdoor temperature before it overheats or overcools the space, the amount of energy used by the HVAC system can be minimized.

In this project, MPC was used in conjunction with weather and the occupancy-prediction models from Section 3.3 to optimize the temperature setpoints. The optimization algorithm was performed to find the setpoint temperature that minimizes both energy use and discomfort. In Section 4.3, occupancy prediction models were used to change the setpoint temperature. In that case, four total temperature setpoints were allowed: the heating setpoint temperature, the cooling setpoint temperature, the heating setback temperature, and the cooling setpoint temperature (Table 4.3). MPC optimization not only considered those four temperatures but also the temperatures between them to find the optimal solution. To run MPC, the optimization algorithm and cost function, the prediction horizon, the execution horizon, and the model were chosen.

4.4.1 Model

Most commonly, MPC is performed utilizing reduced-order linear system models. This allows the optimization to be performed more quickly and easily. Whole building simulations, like EnergyPlus, allow the calculation of radiant heat balances which simplified models cannot capture [55]. In cases where thermal comfort is being evaluated, these calculations are essential. Therefore, MPC was completed using the EnergyPlus model simplified as much as possible to reduce computation time. To reduce the computation time of the simulation, the HVAC equipment was hard-sized using the TMY3 data for January and July, the simulation timestep was increased from

1 to 15 minutes, and the reported variables were reduced [55, 61].

4.4.2 Optimization Parameters

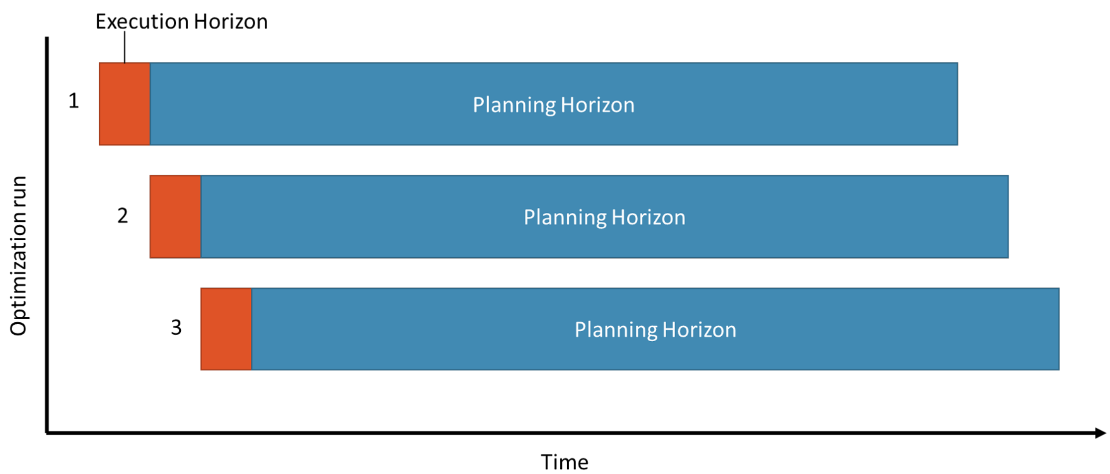

Figure 4.19: Graphical representation of MPC process

Model predictive control requires a prediction horizon to designate how far into the future the model is predicting and optimizing. In this study, a 24-hour prediction horizon was used to account for diurnal temperature swings and internal gains from daily occupancy patterns and equipment use. The execution period, how often the optimization is conducted, used a one-hour horizon to adjust to actual occupancy and indoor temperature values. Once an optimum was found, the optimum temperature setpoint was set, EnergyPlus was run, and stepped forward one hour in time (Figure 4.19). A new optimization process was then restarted with the current state values and a new 24-hour prediction horizon.

The MPC simulation setup consisted of a particle swarm optimization run in Matlab using the EnergyPlus model. Particle swarm optimization (PSO) uses a group of candidate solutions as beginning values. These seeds are run yielding initial results for the cost function. As the particles are evaluated and move towards an optimum solution, they “swarm” towards the optimum solution. By using a swarm, the possibility of finding a local minimum rather than global minimum is reduced [62].

4.4.3 Objective Function

The cost function combines all the factors that are to be evaluated and optimized into a formula [20]. How the cost function is configured determines which values are considered most important in finding the best solution. In this study, the cost function minimized energy use and occupant discomfort. Optimization of this function was limited by one constraint: an allowed temperature band. Discomfort was calculated using predicted mean vote (PMV). Since thermal comfort was not always achieved in the baseline model, the predicted mean vote (PMV) from the baseline was used as the maximum allowed PMV in the optimization run. This prevented the optimization algorithm from penalizing solutions that provided equal comfort performance as the baseline model. By minimizing energy use and occupant discomfort concurrently, MPC could reduce energy use without sacrificing occupant comfort using Equation 4.1.

$$min\left(\sum_k E_k + P\right) \tag{4.1}$$

$$\text{subject to: } T_{lower,k} \leq T_{optimal,k} \leq T_{upper,k}$$

where P is the occupied discomfort, E_k is hourly HVAC consumption at timestep k, T is the temperature setpoints at timestep k, and k is the number of timesteps in the evaluation.

Occupant discomfort is calculated as shown in Equation 4.2.

$$P = C\sum_k (\mid PMV_k \mid - PMV_{max}) \tag{4.2}$$

where C is the comfort penalty slope, PMV_k is predicted mean vote during timestep k, and PMV_{max} is the PMV comfort threshold. Because the PMV during the baseline run may exceed 0.5, such as in the case of Seattle in summer, the threshold is adjusted to allow the optimized MPC to use an equitable PMV during optimization. This threshold is calculated with Equation 4.3.

$$PMV_{max} = max(0.5, \mid PMV_{base,k} \mid) \tag{4.3}$$

where PMV_{max} takes the higher value between 0.5 and the PMV from the baseline run at timestep k.

The goal of MPC is to optimize for all the factors within the cost function. In this thesis two factors were used: energy consumption and thermal discomfort. Relative importance between the two factors was controlled by C, the comfort penalty slope. This value determines how strictly the PMV is controlled. With a smaller C value, uncomfortable setpoint temperatures do not increase the cost function as much, allowing some thermal discomfort to occur in favor of energy savings. Therefore, the comfort slope allows the cost function flexibility. It can be tuned to meet the goals of the occupant depending on how much discomfort they are willing to experience. To determine an appropriate comfort slope, MPC simulations were completed for two days for House 1. The house was vacant once each day for an average of six hours, allowing the comfort slope to be optimized during both occupancy states. Figure 4.20 shows the summarized results. Discomfort was evaluated by classifying hours within ± 0.5 PMV as comfortable, ± 0.7 PMV as Class C discomfort, and beyond ± 0.7 PMV as excessive discomfort. As the comfort slope value increased, hours of discomfort decreased while energy use increased. A comfort slope of was chosen that allowed a few hours of Class C discomfort. Using a value of 1000, allowed energy savings while still maintaining comfort most of the time. For all simulations in this study, a comfort penalty slope of 1000 was used.

4.4.4 Thermal History Management

When starting a whole building simulation, the temperatures of all the surfaces and the heat balance between them needs to be established. EnergyPlus achieves this by completing a warm-up period, in which the first day is repeatedly simulated until the temperatures and heat balance converge. This, however, assumes that the past thermal history is the exact same as the first day. Previous studies have shown that when simulations are completed in EnergyPlus with the preceding simulation time period varied, the HVAC cooling energy for the same day changes, indicating that the thermal history during the simulation needs to be long enough that the building has achieved its

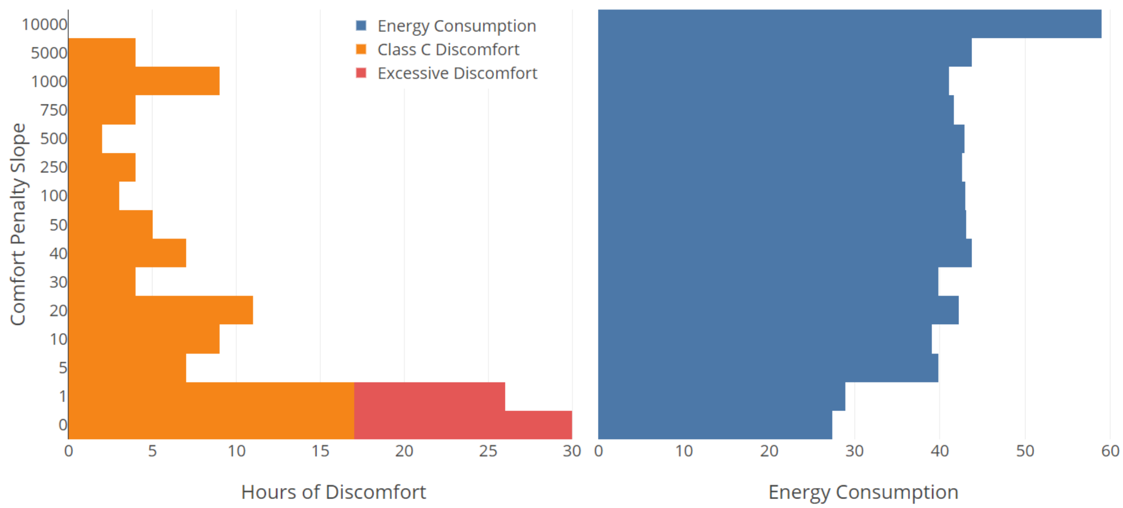

Figure 4.20: Resulting comfort and energy use, by comfort slope value

true thermal state before the intended simulation period begins [55]. Termed pre-conditioning, the length of time necessary to negate the simulation thermal history effects depends on the building [55]. A thermal history experiment was conducted on all five home models by varying the start day between 0 to 14 days prior to a constant, measured day. HVAC energy transfer for the last day was measured in all simulations to determine how many pre-conditioning days were required until this value remained steady.

Due to the thermal history affecting the model, the energy transfer for the measured day was not always constant (Figure 4.21). Boston in winter was affected by pre-conditioning the most. Since this is the biggest house and has the highest temperature differential between the inside and outside, this is expected. All other house models had differences between 0% and 1%. In literature, differences as large as 30% occurred [55]. Because the models in this thesis are light frame houses rather than large commercial buildings with high thermal capacitance, the thermal history has less of an effect. The Boston winter model, which needed to most thermal conditioning, stabilizes with one conditioning day. To avoid the slight differences seen for the first six days in the Atlanta and Houston summer models, a pre-conditioning horizon of eight days was used for all simulations.

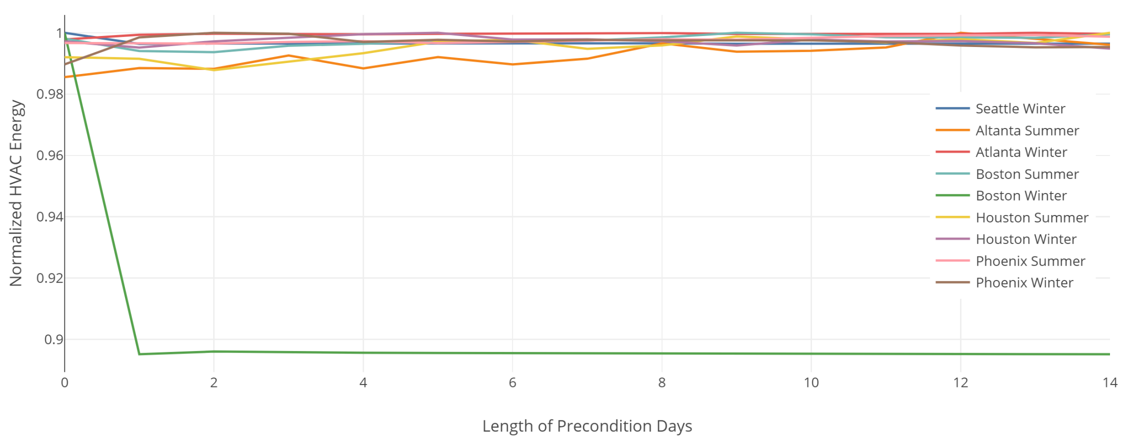

Figure 4.21: Comparison of HVAC energy by pre-conditioning days for all house models

4.4.5 Simulation

MPC was used to simulate three scenarios. In the first scenario, the individual hybrid prediction model developed in Section 3.3, was used for House 2 and 5 in Houston in summer. In the second scenario, the individual hybrid prediction model was used for House 1 and 2 in Atlanta in summer. Lastly, perfect occupancy forecasting was used for House 1 and 2 in Atlanta in summer. These three scenarios each yielded different results and insights.

MPC Case 1: Houston with occupancy prediction

In this simulation, MPC was used for Houston in summer. EnergyPlus model settings, optimization parameters, objective function, and thermal history management developed in Sections 4.4.1 to 4.4.4 were used for this simulation. A summary of these settings is provided in Table 4.6. Computation time allowed for each optimization 30 minutes, allowing for at least 300 optimization runs for each execution horizon. MPC simulation for each house took 60 hours to complete the five-day run period.

Table 4.6: Settings used for MPC optimization in Case 1

Parameter	Value
City	Houston
Season	Summer
Houses	2 & 5
Prediction model	Individual Hybrid
Run period	5 days
Timestep	15 minutes
Planning horizon	24 hours
Execution horizon	1 hour
Occupied allowed temperatures	$22°\mathrm{C} \leq \mathrm{T}_{optimal,k} \leq 24.5°C$
Unoccupied allowed temperatures	$18°\mathrm{C} \leq \mathrm{T}_{optimal,k} \leq 28°C$
Temperature increments	0.5°C
Comfort penalty slope (C)	1000
Optimization time per execution horizon	30 minutes

Results for Case 1 are shown in Table 4.7. Discomfort is measured by exceedance above Class A and B comfort. This value was calculated by summing the operative temperature deviation above 26.0°C or below 23.0°C for all occupied hours. For both houses, the energy saved is very low, with an average savings of 1%, but achieved little to no discomfort. Figures 4.23 and 4.22 show the electricity use and temperatures for two days of the simulated run period. In Figure 4.23, which shows the electricity consumption for the building, MPC optimized consumption shows reduced consumption during unoccupied hours as the allowed maximum temperature is changed from 24.5°C to the higher 28°C. When the occupant returns, the electricity consumption increases above the conventional electricity consumption as it has to reduce the indoor temperature rather than simply maintain it.

The MPC optimization results found that the highest allowed temperatures provided the lowest resulting cost function. Although any temperature within the band was allowed, the optimal temperature ending up matching the values used in occupancy-based setpoint control, signaling that for the Case 1 setpoint control and MPC optimization yielded the same temperatures.

Table 4.7: Results for Case 1

House	Energy Savings	Discomfort
2	2.1%	3.7 Kh
5	0.2%	0 Kh

Figure 4.24 shows a duration curve of the temperature deviation for all hours of the simulation. Temperature deviation was determined by calculating the absolute value of the difference between the operative temperature and 24.5°C. All values within 1.5°C are considered comfortable by ISO 7730 standards. Most hours of the simulation are within the comfort region, yielding high comfort but low energy savings. Hours in which the temperature was allowed to drift above comfortable temperatures were few due to the small number hours when the prediction model accurately predicted the house to be vacant.

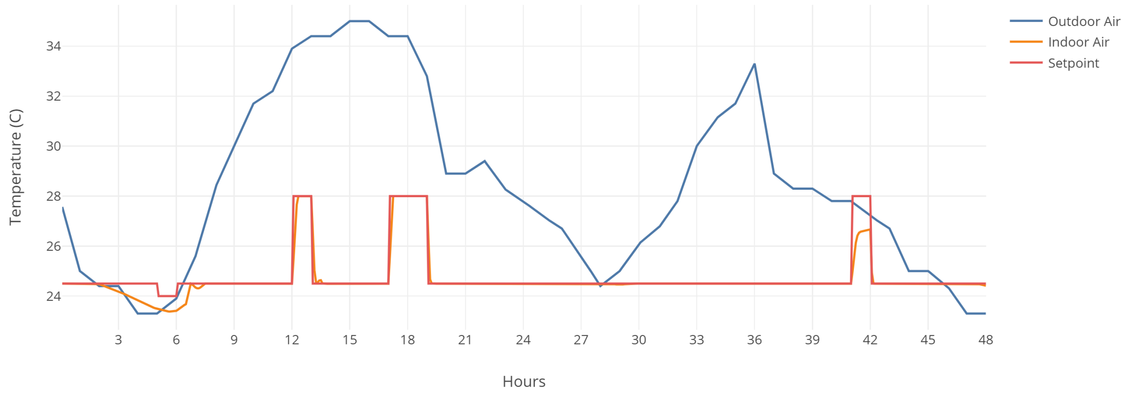

Figure 4.22: Case 1 temperatures for House 2 using individualized hybrid prediction model, by hour

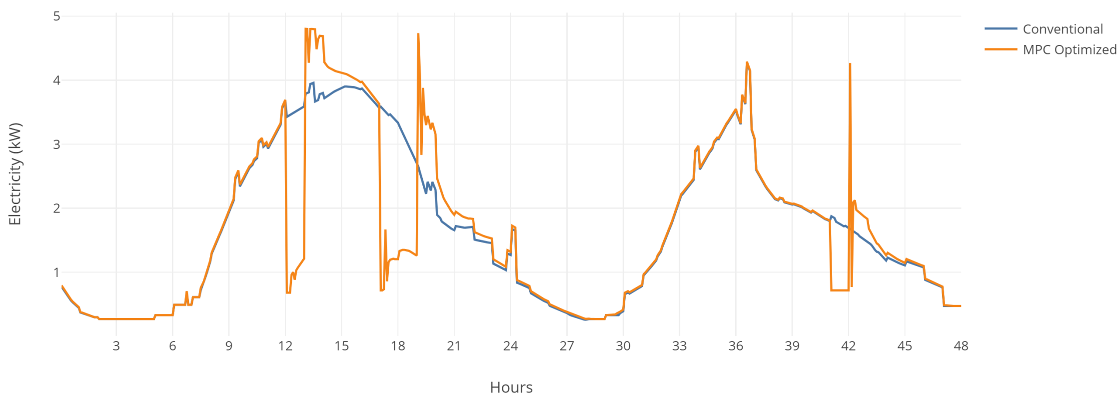

Figure 4.23: Case 1 electricity consumption for House 2 using individualized hybrid prediction model, by hour

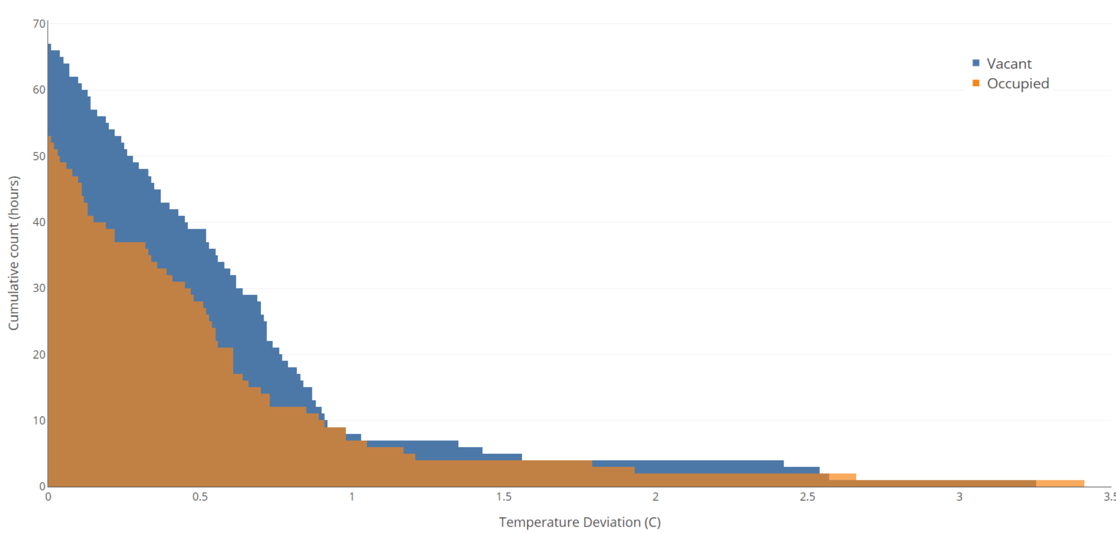

Figure 4.24: Case 1 duration curve of deviation from 24.5°C operative temperatures

Table 4.8: Settings used for MPC optimization in Case 2A

Parameter	Value
City	Atlanta
Season	Summer
Houses	1 & 2
Prediction model	Individual Hybrid
Run period	1 week
Timestep	15 minutes
Planning horizon	24 hours
Execution horizon	1 hour
Occupied allowed temperatures	19 °C $\leq$T$_{optimal,k}$ $\leq 27°C$
Unoccupied allowed temperatures	19 °C $\leq$T$_{optimal,k}$ $\leq 27°C$
Temperature increments	0.5 °C
Comfort penalty slope (C)	1000
Optimization time per execution horizon	30 minutes

MPC Case 2A: Atlanta with occupancy prediction

In Case 2A, MPC optimization was performed for Atlanta in summer. Individual hybrid occupancy prediction models developed in Section 3.3 were used. Unlike Case 1, the maximum allowed temperatures during occupied hours were kept at the same values used during unoccupied hours. By allowing the large temperature band at all times, the constraints within the cost function were reduced. Temperatures which produced the smallest cost function were used, rather than the being restricted by the temperature band. A summary of all the settings used for the simulation is shown in Table 4.8.

Results of these simulations are summarized in Table 4.9. For this simulation, the two houses saved an average of 9.0% in energy use. Exceedance is higher than Case 1 with an average of 35.3 Kh. Total occupied hours for both houses for the week was 185 hours, with House 1 being occupied for 100 hours and House 2 being occupied for 85 hours. With the average exceedance of 35.3 Kh, an average distribution of thermal discomfort would yield 0.4°C above the ideal temperatures.

Figures 4.25 and 4.26 show temperatures for each house for two days of the simulation. Chosen temperatures ranged from 19°C to 27°C, with the average setpoint temperature at 25.4°C and 26.2°for House 1 and 2, respectively. With the expanded temperature range, temperature

Table 4.9: Results for Case 2A

House	Energy Savings	Discomfort
1	7.5%	30.8 Kh
2	10.4%	39.8 Kh

values selected did not always conform to setpoint temperatures as seen in Case 1. While setpoint values allowed an 8°C range, ambient air temperatures occurring within the building had a 4.5°C to 4.7°C range. More extreme setpoint temperature only lasted for an hour, preventing temperature within the building from reaching the setpoint and maintaining a comfortable space despite the setpoints used.

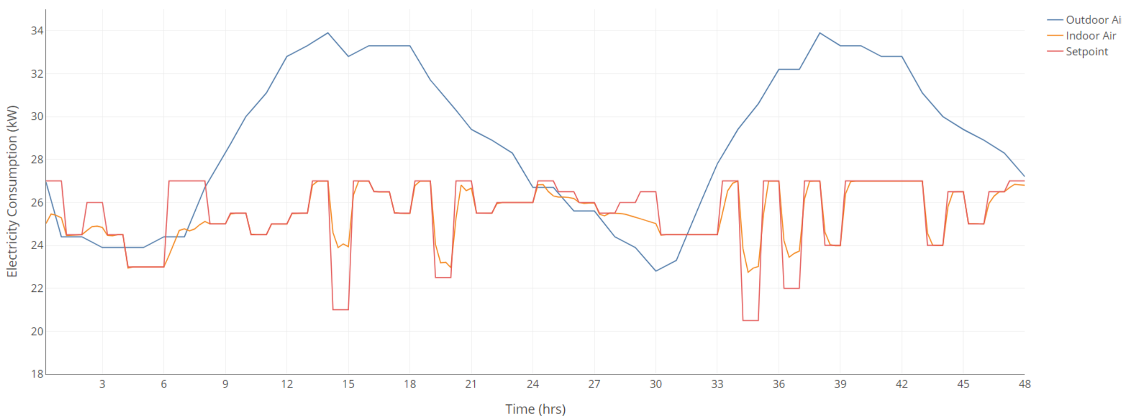

Figure 4.25: Case 2A temperatures for House 1 using individualized hybrid prediction model, by hour

Figures 4.27 and 4.28 show the electricity consumption resulting from the setpoint temperatures used. Due to changing setpoint temperatures, electricity consumption jumped in hours using low setpoint temperatures as more cooling occurred. In other hours, however, electricity was significantly less than the conventional constant temperature. Over the week simulated, electricity consumption was reduced to allow 7.5% and 10.4% energy savings for House 1 and 2 respectively.

Figures 4.29 and 4.30 show the duration curve of operative temperatures for the two houses. Both figures show that allowed temperature deviation was higher for unoccupied hours. In hours that the house was occupied, operative temperature was kept closer to the centerpoint temperature

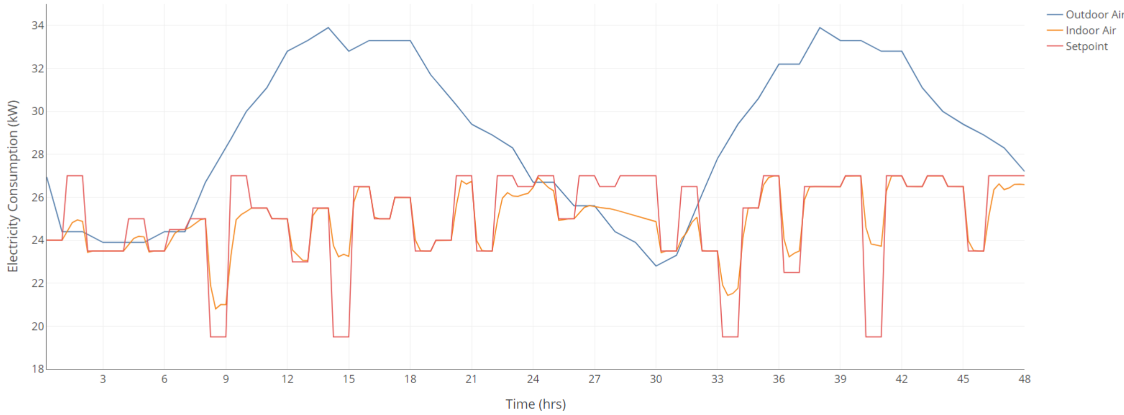

Figure 4.26: Case 2A temperatures for House 2 using individualized hybrid prediction model, by hour

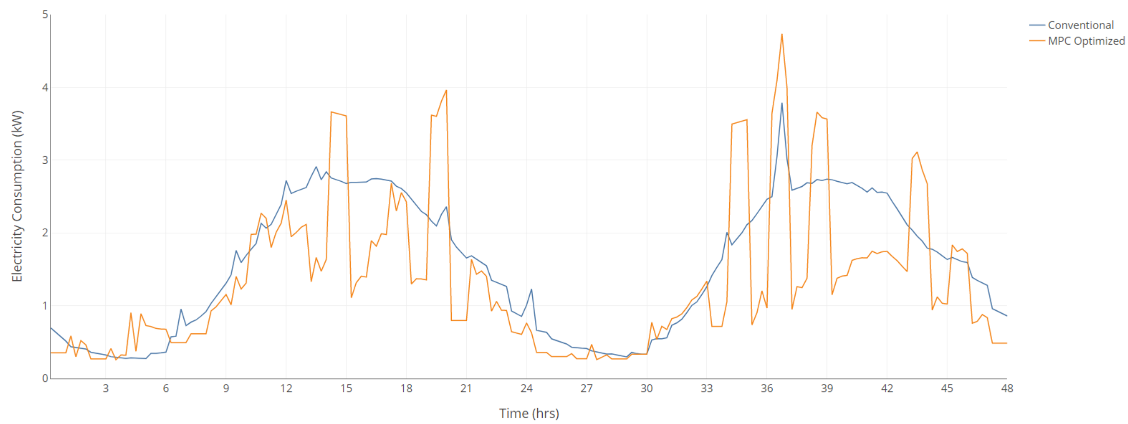

Figure 4.27: Case 2A electricity consumption for House 1 using individualized hybrid prediction model, by hour

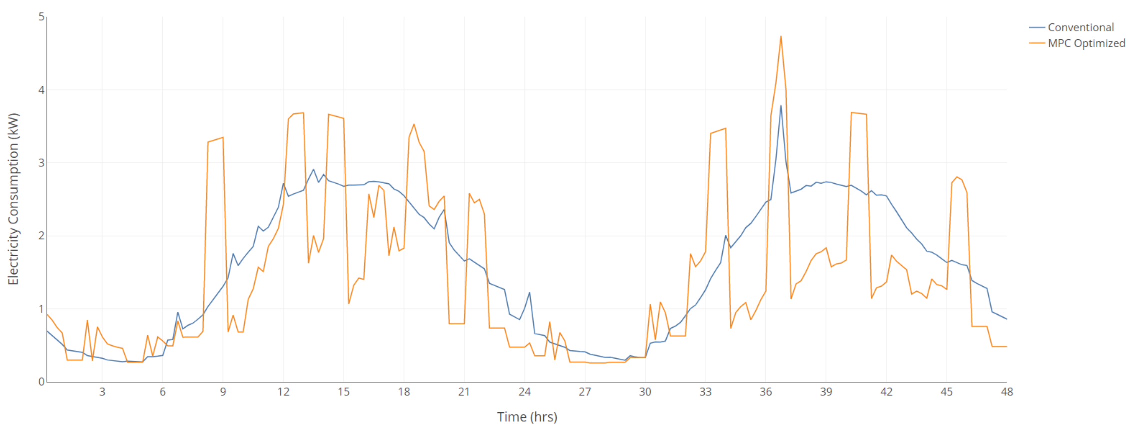

Figure 4.28: Case 2A electricity consumption for House 2 using individualized hybrid prediction model, by hour

of 23.5°. With the chosen cost function, some temperature deviation was allowed to achieve higher energy savings. Unlike Case 1, which had tight occupied temperature constraints, temperature deviation in Case 2A is higher. Changes to the comfort penalty slope would change how much deviation is allowed and, in result, how much energy was saved.

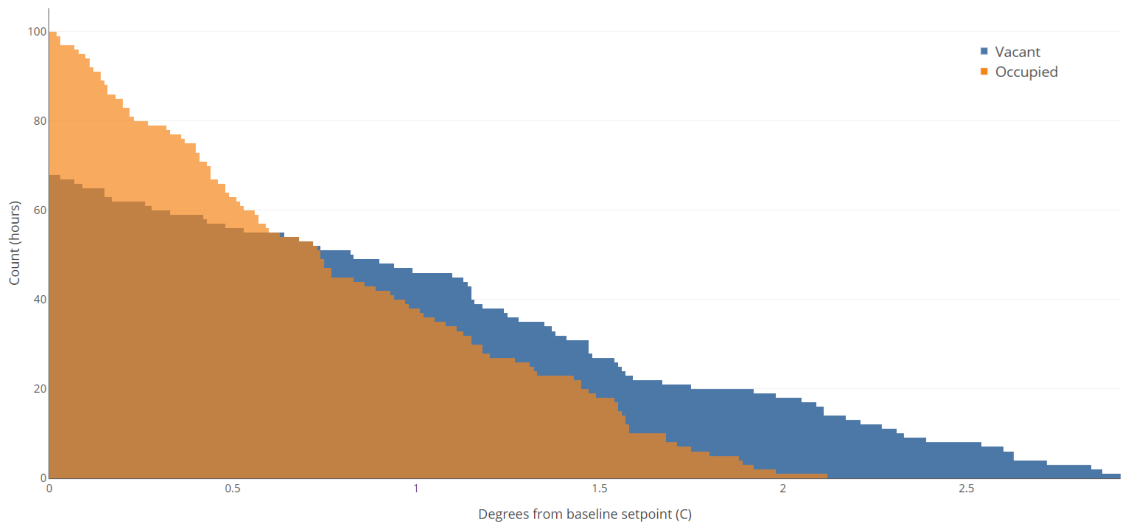

Figure 4.29: Case 2A duration curve of deviation from 24.5°C operative temperatures

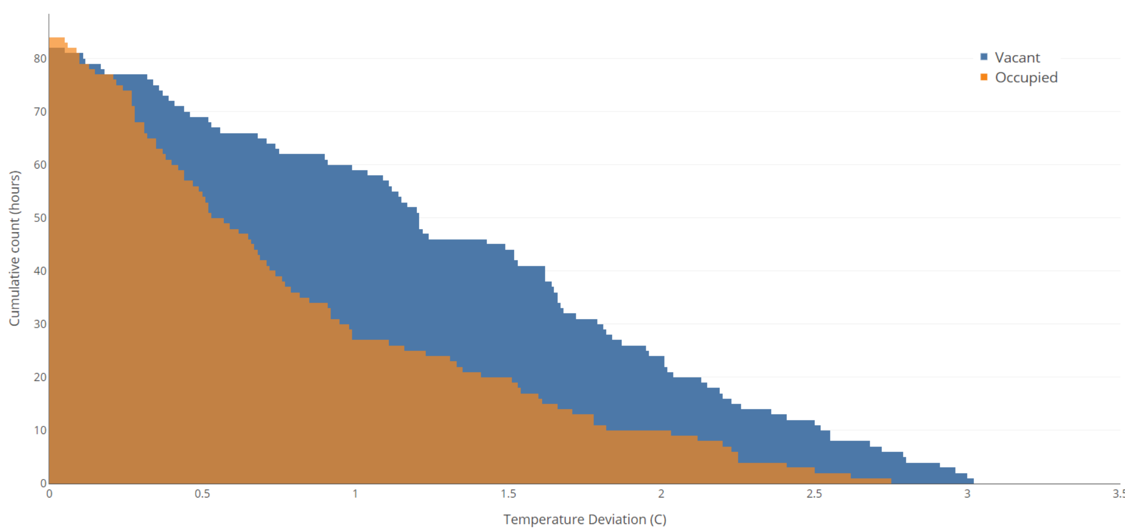

Figure 4.30: Case 2A duration curve of deviation from 24.5°C operative temperatures

Table 4.10: Settings used for MPC optimization in Case 2B

Parameter	Value
City	Atlanta
Season	Summer
Houses	1 & 2
Prediction model	Perfect forecasting
Run period	1 week
Timestep	15 minutes
Planning horizon	24 hours
Execution horizon	1 hour
Occupied allowed temperatures	19 °C $\leq$T$_{optimal,k}$ $\leq 27°C$
Unoccupied allowed temperatures	19 °C $\leq$T$_{optimal,k}$ $\leq 27°C$
Temperature increments	0.5 °C
Comfort penalty slope (C)	1000
Optimization time per execution horizon	30 minutes

MPC Case 2B: Atlanta with perfect occupancy forecasting

In a third scenario, all settings used from Case 2A were kept the same except for the occupancy prediction model. In this scenario, actual occupancy data was used to imitate perfect occupancy forecasting to allow an exploration of how imperfections in the occupancy prediction impact MPC. A summary of all used settings are shown in Table 4.10.

Results from the two MPC optimizations are shown in Table 4.11. Energy savings for House 1 and 2 increased by 5.4% and 2.9% respectively. Comfort exceedance decreased. Figures 4.31 and 4.32 show resulting temperatures found during MPC optimization. Like Case 2A, setpoint temperatures range from 19°C to 27°C. Unlike Case 2A, the number of times low setpoint temperatures are used to achieve a quick temperature change is less. Resulting ambient air temperature ranged from 22°C to 27°C.

Table 4.11: Results for Case 2B

House	Energy Savings	Discomfort
1	12.9%	21.0 Kh
2	13.3%	21.0 Kh

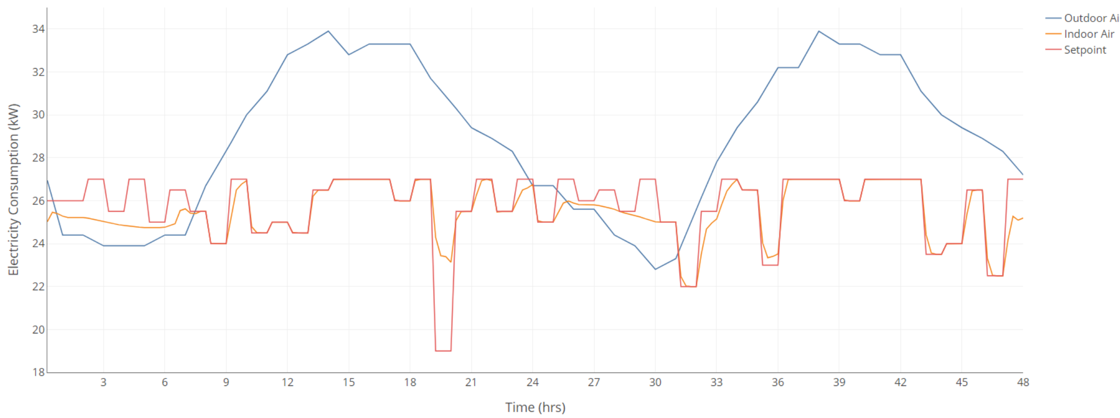

Figure 4.31: Case 2B temperatures for House 1 using perfect occupancy forecasting, by hour

Figures 4.33 and 4.34 show electricity consumption during the same two-day period. While higher power consumption occurs during certain hours, overall consumption is less.

Resulting temperatures within the home are shown in Figures 4.35 and 4.36. With an accurate occupancy forecast, MPC optimization is able to allow less energy-intensive temperatures during vacant periods without the penalty of high energy consumption and discomfort when an occupant unexpectedly returns. This allows improvements in both energy savings and thermal comfort. With accurate occupancy forecasting, energy savings above 10% are possible with the occupancy patterns recorded. Accurate occupancy prediction, therefore, is essential in improving HVAC control.

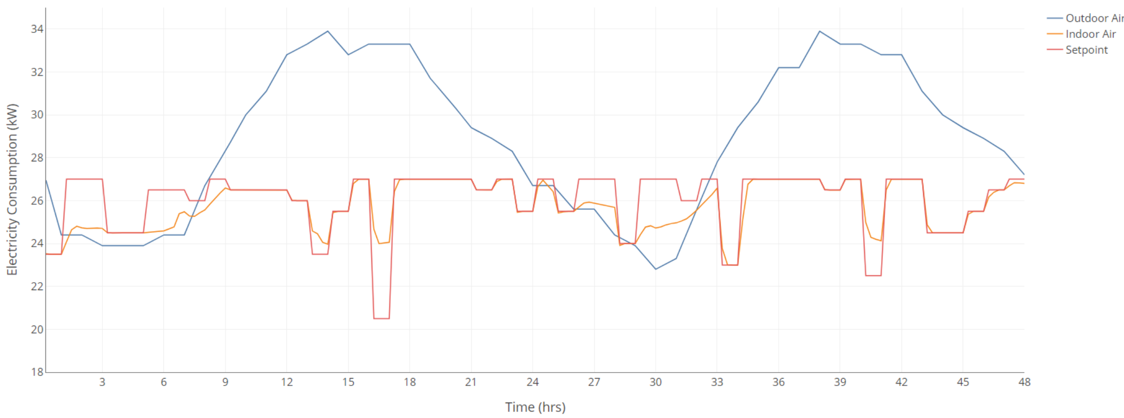

Figure 4.32: Case 2B temperatures for House 2 using perfect occupancy forecasting, by hour

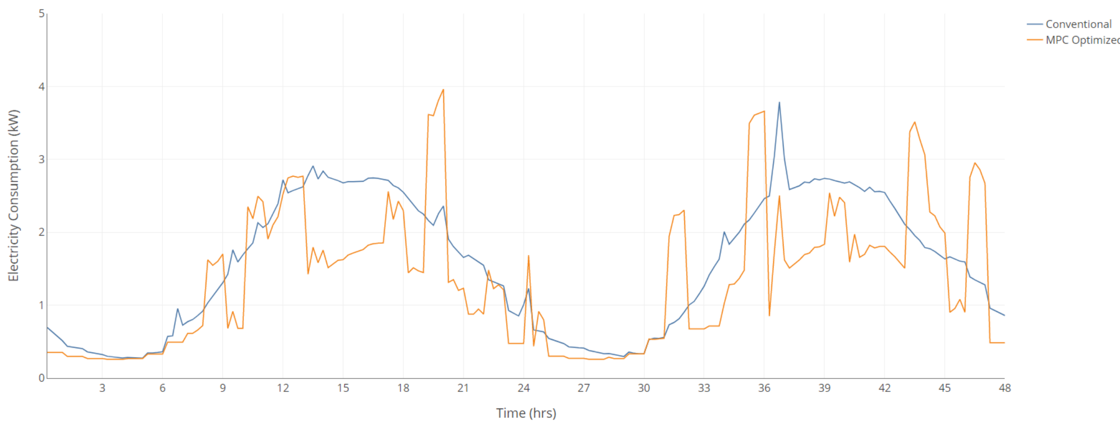

Figure 4.33: Case 2B electricity consumption for House 1 using perfect occupancy forecasting, by hour

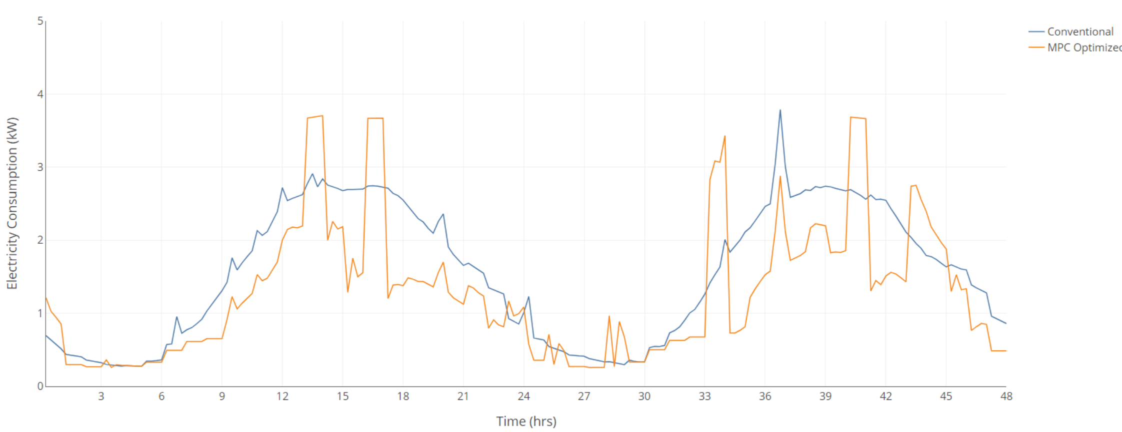

Figure 4.34: Case 2B electricity consumption for House 2 using perfect occupancy forecasting, by hour

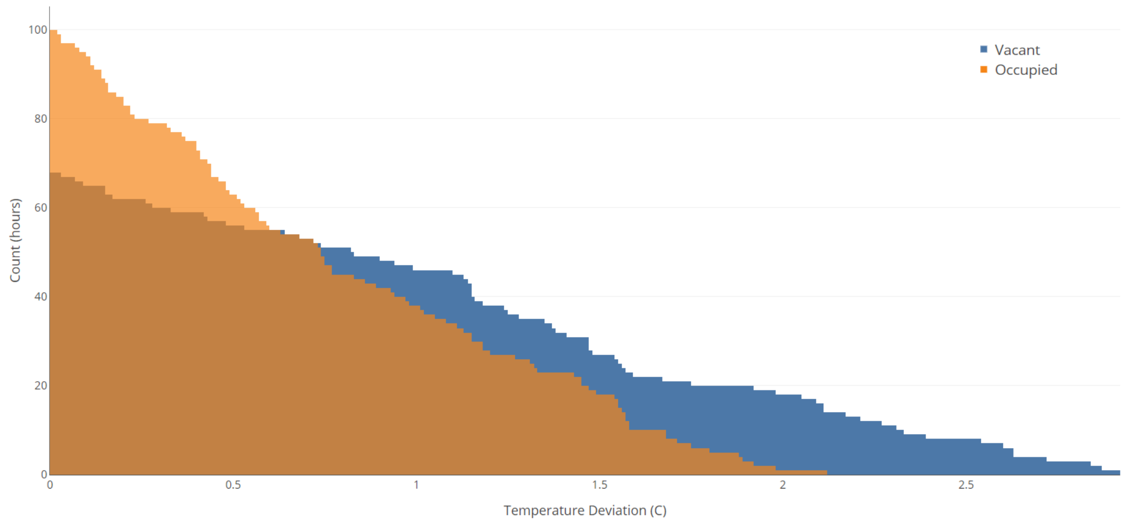

Figure 4.35: Case 2B duration curve of deviation from 24.5°C operative temperatures

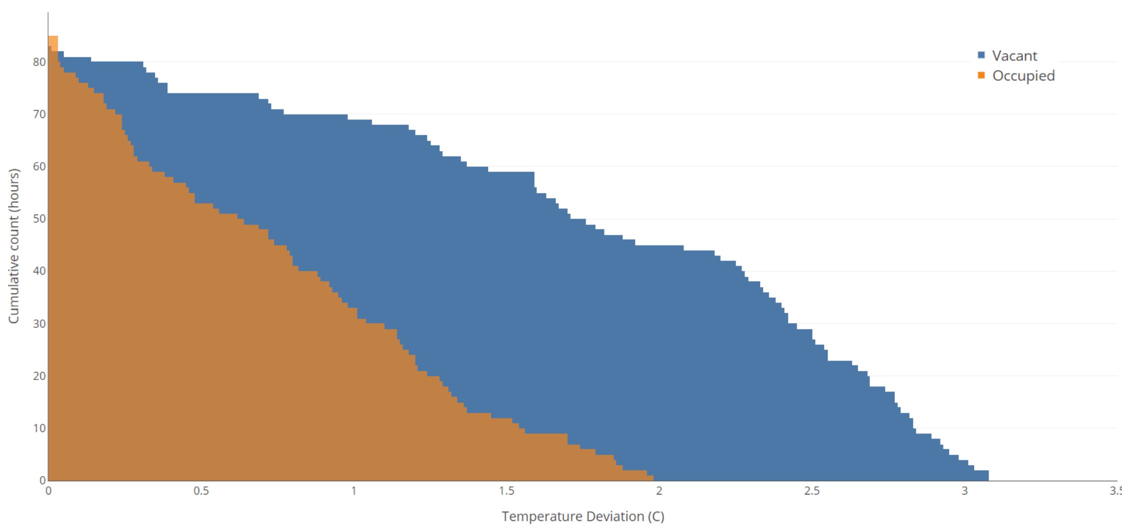

Figure 4.36: Case 2B duration curve of deviation from 24.5°C operative temperatures

Chapter 5

Summary and Conclusions

Residential heating and cooling accounts for a large portion of annual energy consumption in the United States. Reducing energy use can contribute to furthering the goals of the Paris Agreement. By accounting for occupancy in HVAC control, energy savings can be realized without negatively impacting thermal comfort. Maintaining occupant control is essential to allow widespread adoption.

The goal of this thesis was determine the prospective impact of incorporating occupancy into residential HVAC control. To accomplish this, a review of historical residential energy use, housing characteristics, typical temperature control and occupancy prediction was completed. Using conclusions from previous studies, a non-probabilistic occupancy prediction model was chosen for the basis of this data study. Analysis was completed in two parts. First, real occupancy data from six homes in Boulder, Colorado was used to generate and evaluate occupancy prediction models. Secondly, the generated occupancy prediction models were incorporated into HVAC control strategies and evaluated by a building simulation program. HVAC control strategies simulated were occupancy-based setpoint control and occupancy-based model predictive control. Results were evaluated using energy use and thermal comfort as the success criteria.

5.0.1 Conclusions

The literature review completed in Section 2.4.1 revealed that non-probabilistic models performed best for short-term occupancy prediction in previous studies. Collected occupancy data from six homes was used to to generate and evaluate the best non-probabilistic model and verified against the collected data. Prediction inaccuracy, termed state-matching error in the study, ranged from 8.0% to 48.7%. Model training data that used a moving, multi-week training set worked the best. Occupancy pattern was the highest contributor to prediction accuracy. Examination of increased training time indicates that models can improve over time as more data is collected and included into the prediction model.

Once the occupancy prediction models were generated, they were incorporated in occupancy-based setpoint control and occupancy-based model predictive control in a building performance simulation. Five EnergyPlus home models were used to simulate the energy use and temperatures for a two-week period in summer and in winter. Occupancy-based setpoint control showed possible energy savings from 0% to 50.0% over a constant setpoint temperature depending on climate, occupancy pattern, and control strategy. Non-probabilistic prediction models achieved the highest energy savings, with an average of 10.0%, but with the disadvantage of high thermal discomfort. By including an override, in which the occupancy prediction model can sense the actual occupant presence and react to it, thermal discomfort was reduced. In these hybrid occupancy models, the energy savings averaged 5.0%.

Model predictive control showed that energy savings is highly dependent on how the cost function and constraints are parameterized. In Case 1, where the temperature constraints were much stricter during occupied hours, little to no energy savings was achieved. However, in Case 2A, where temperature constraints were relaxed during occupied hours, energy savings increased with only a slight impact on discomfort. In Case 2B, where occupancy was perfectly predicted, both energy savings and thermal comfort improved. This leads to two implications. First, a cost function that combines both energy consumption and thermal discomfort allows for flexibility to the occupant to determine what tradeoff between energy savings and discomfort is appropriate. Secondly, accurate occupancy prediction improves both performance aspects in the cost function. As occupancy prediction improves, the ability for occupancy-aware HVAC control to maintain comfort and increase energy savings improves.

In this study, the cost function optimized energy savings and thermal comfort. Future work could use alternative cost functions to achieve different goals. Monthly utility costs are a main concern for many users. The cost function could be changed to use estimated cost rather than energy consumption. Both energy consumption and energy cost would reduce energy use. However, energy cost would allow time of use (TOU) rates to be incorporated. With TOU rates, reducing energy use during certain hours of the day is more more beneficial than others. This would assist utilities by shifting loads to reduce stress on the electric grid while saving money and energy for the consumer.

Figure 5.1 shows the energy savings from this study in comparison to Nagele's 2017 study which analyzed ten houses in southern Germany using the occupancy data and house characteristics of each individual home to estimate annual energy savings [3].

Overall, energy savings simulated in this thesis are lower than those estimated in the previous study. This was found in all control strategies modeled. In comparing occupancy data for the ten houses in the previous study and the six houses in this study, occupancy rate differs. Table 5.1 shows occupancy rates for the two studies. Median occupancy rate in the literature study was 62%, while in this study the rate was 78%. With 16% difference in occupied times, the amount of time setback temperatures could be used was much less. This likely contributed largely to the differences seen. Additionally, the two studies used different house models, climates, and simulation periods, which could lead to further differences.

Table 5.1: Comparison of occupancy rate for two studies

Study	Lowest Occupancy	Median Occupancy	Highest Occupancy
Literature	55%	62%	82%
This study	52%	78%	86%

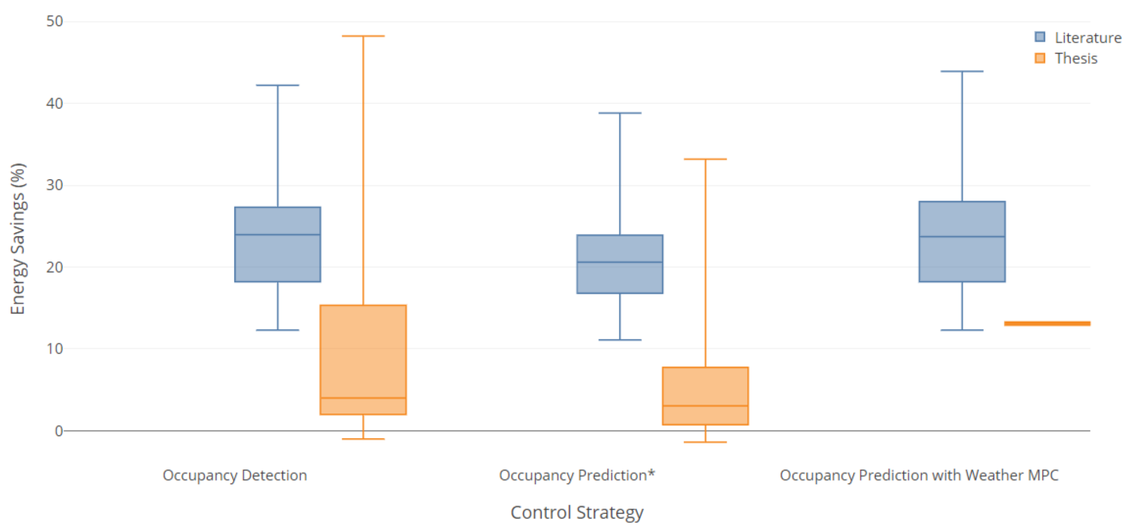

Figure 5.1: Comparison of energy savings to previous studies, by control strategy

5.0.2 Future Work

Smart thermostat technology already exists and is in use in homes nationwide. This existing infrastructure is a platform that could use occupancy-prediction HVAC control with a simple software update. With their ability to detect occupancy, they could act as a tool for measuring occupancy and subsequently use the collected data to create prediction models. Connected thermostats offer a large opportunity to easily implement changes in new and existing homes alike.

In order to reduce negative occupant experiences, further work to improve occupancy prediction models is proposed. Some questions that should be answered with further research are:

- How well can other residential occupancy data be predicted? In this thesis, data was collected from occupants enrolled or involved with the university. How would this change with occupants who have a more "typical" schedule?
- Do other types of prediction models work better? The literature review suggests non-probabilistic models work best for short-term occupancy prediction, however, other types were not analyzed in this study.
- How does occupancy data improve with longer periods of collected data? For this study, the occupancy was recorded for 4-9 weeks. How would the prediction model accuracy change with more training data?

By improving occupancy prediction, the full extent of the energy savings potential can be realized. As seen in the building performance simulation, the maximum potential depends highly on the occupancy pattern. Future work could refine the energy savings estimation by exploring the following items:

- What are the energy savings potential with additional occupancy patterns? In comparing the results of this thesis to past literature, the estimated results in this study are smaller. Additional occupancy patterns with a higher vacancy rate could improve energy savings and be more representative of the average savings potential.

- If an occupancy-aware HVAC control was implemented in a building, what would be the actual realized savings? The building models in this simulated used constructions typical for their climate. By verifying with a real building, the accuracy of the building simulation could be established.

Bibliography

[1] Independent Statistics. Drivers of U . S . Household Energy Consumption , 1980-2009. (February), 2015.

[2] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems - SenSys '10, page 211, 2010.

[3] Florian Nägele, Thomas Kasper, and Bastien Girod. Turning up the heat on obsolete thermostats: A simulation-based comparison of intelligent control approaches for residential heating systems. Renewable and Sustainable Energy Reviews, 75(January):1254–1268, 2017.

[4] William O'Brien, Isabella Gaetani, Sara Gilani, Salvatore Carlucci, Pieter Jan Hoes, and Jan Hensen. International survey on current occupant modelling approaches in building performance simulation†. Journal of Building Performance Simulation, 10(5-6):653–671, 2017.

[5] Simona D'Oca, Tianzhen Hong, and Jared Langevin. The human dimensions of energy use in buildings: A review. Renewable and Sustainable Energy Reviews, 81(May 2017):731–742, 2018.

[6] Ardeshir Mahdavi and Mahnameh Taheri. An ontology for building monitoring. Journal of Building Performance Simulation, 10(5-6):499–508, 2017.

[7] TCEP: Buildings, Sep 2019. [Online; accessed 19. Sep. 2019].

[8] Luis Pe. A review on buildings energy consumption information . 40:394–398, 2008.

[9] Residential Energy Consumption Survey (RECS) - Data - U.S. Energy Information Administration (EIA), Sep 2019. [Online; accessed 19. Sep. 2019].

[10] A sense of units and scale for electrical energy production and consumption, Sep 2019. [Online; accessed 19. Sep. 2019].

[11] Kaiyu Sun and Tianzhen Hong. A simulation approach to estimate energy savings potential of occupant behavior measures. Energy and Buildings, 136:43–62, 2017.

[12] Yan Zhang, Xuemei Bai, Franklin P. Mills, and John C.V. Pezzey. Rethinking the role of occupant behavior in building energy performance: A review. Energy and Buildings, 172:279–294, 2018.

[13] What is the Paris Agreement? | UNFCCC, Sep 2019. [Online; accessed 19. Sep. 2019].

[14] Vahid Vakiloroaya, Bijan Samali, Ahmad Fakhar, and Kambiz Pishghadam. A review of different strategies for HVAC energy saving. Energy Conversion and Management, 77:738–754, 2014.

[15] United States Employment Rate | 2019 | Data | Chart | Calendar | Forecast, Sep 2019. [Online; accessed 19. Sep. 2019].

[16] ASHRAE. ASHRAE Standard 55-2010. ASHRAE Inc., 2010:42, 2010.

[17] P. O. Fanger. Thermal confort analysis and applications in environmental engineering. Mcgraw-Hill, 1972.

[18] Todd Malinick, Nate Wilairat, Jennifer Holmes, and Lisa Perry. Destined to Disappoint: Programmable Thermostat Savings are Only as Good as the Assumptions about Their Operating Characteristics. ACEEE Summer Study on Energy Efficiency in Buildings, (7):162–173, 2012.

[19] DOE. Overview of Existing and Future Residential Use Cases for Connected Thermostats. Technical report, 2016.

[20] Abdul Afram and Farrokh Janabi-Sharifi. Theory and applications of HVAC control systems - A review of model predictive control (MPC). Building and Environment, 72:343–355, 2014.

[21] Mengda Jia, Ravi S. Srinivasan, and Adeeba A. Raheem. From occupancy to occupant behavior: An analytical survey of data acquisition technologies, modeling methodologies and simulation coupling mechanisms for building energy efficiency. Renewable and Sustainable Energy Reviews, 68(June 2016):525–540, 2017.

[22] Muhammad Aftab, Chien Chen, Chi Kin Chau, and Talal Rahwan. Automatic HVAC control with real-time occupancy recognition and simulation-guided model predictive control in low-cost embedded system. Energy and Buildings, 154:141–156, 2017.

[23] Zhenghua Chen and Yeng Chai Soh. Comparing occupancy models and data mining approaches for regular occupancy prediction in commercial buildings. Journal of Building Performance Simulation, 10(5-6):545–553, 2017.

[24] David Lindelöf, Hossein Afshari, Mohammad Alisafaee, Jayant Biswas, Miroslav Caban, Xavier Mocellin, and Jean Viaene. Field tests of an adaptive, model-predictive heating controller for residential buildings. Energy and Buildings, 99:292–302, 2015.

[25] Abdul Afram and Farrokh Janabi-Sharifi. Supervisory model predictive controller (MPC) for residential HVAC systems: Implementation and experimentation on archetype sustainable house in Toronto. Energy and Buildings, 154:268–282, 2017.

[26] Wilhelm Kleiminger, Friedemann Mattern, and Silvia Santini. Predicting household occupancy for smart heating control: A comparative performance analysis of state-of-the-art approaches. Energy and Buildings, 85:493–505, 2014.

[27] Southern California Edison. Residential Smart Communicating Thermostat. 2017.

[28] Marco Pritoni, Jonathan M. Woolley, and Mark P. Modera. Do occupancy-responsive learning thermostats save energy? A field study in university residence halls. Energy and Buildings, 127:469–478, 2016.

[29] Robert H. Socolow. The twin rivers program on energy conservation in housing: Highlights and conclusions. Energy and Buildings, 1(3):207–242, 1978.

[30] Adil Al-Mumin, Omar Khattab, and Gopisetti Sridhar. Occupants' behavior and activity patterns influencing the energy consumption in the Kuwaiti residences. Energy and Buildings, 35(6):549–559, 2003.

[31] Rune Andersen. The influence of occupants ' behaviour on energy consumption investigated in 290 identical dwellings and in 35 apartments. Abstract from 10th International Conference on Healthy Buildings, (January 2012):1–3, 2012.

[32] Tianzhen Hong, Sarah C. Taylor-Lange, Simona D'Oca, Da Yan, and Stefano P. Corgnati. Advances in research and applications of energy-related occupant behavior in buildings, 2016.

[33] Isabella Gaetani, Pieter Jan Hoes, and Jan L.M. Hensen. Occupant behavior in building energy simulation: Towards a fit-for-purpose modeling strategy. Energy and Buildings, 121:188–204, 2016.

[34] Carlos Duarte, Kevin Van Den Wymelenberg, and Craig Rieger. Revealing occupancy patterns in an office building through the use of occupancy sensor data. Energy and Buildings, 67:587–595, 2013.

[35] Da Yan, William O'Brien, Tianzhen Hong, Xiaohang Feng, H. Burak Gunay, Farhang Tahmasebi, and Ardeshir Mahdavi. Occupant behavior modeling for building performance simulation: Current state and future challenges. Energy and Buildings, 107:264–278, 2015.

[36] Ardeshir Mahdavi and Farhang Tahmasebi. Predicting people's presence in buildings: An empirically based model performance analysis. Energy and Buildings, 86:349–355, 2015.

[37] J. Page, D. Robinson, N. Morel, and J. L. Scartezzini. A generalised stochastic model for the simulation of occupant presence. Energy and Buildings, 40(2):83–98, 2008.

[38] Ki Uhn Ahn and Cheol Soo Park. Different occupant modeling approaches for building energy prediction. In Energy Procedia, volume 88, 2016.

[39] Best Directory | Building Energy Software Tools, Sep 2019. [Online; accessed 19. Sep. 2019].

[40] F Tahmasebi and A Mahdavi. Exploring the Predictive Potential of Probabilistic Occupancy Models. (September 2014), 2018.

[41] Ian Richardson, Murray Thomson, and David Infield. A high-resolution domestic building occupancy model for energy demand simulations. Energy and Buildings, 40(8):1560–1566, 2008.

[42] Joakim Widén and Ewa Wäckelgård. A high-resolution stochastic model of domestic activity patterns and electricity demand. Applied Energy, 87(6):1880–1892, 2010.

[43] Simona D'Oca, Valentina Fabi, Stefano P. Corgnati, and Rune Korsholm Andersen. Effect of thermostat and window opening occupant behavior models on energy use in homes. Building Simulation, 7(6):683–694, 2014.

[44] Xiaohang Feng, Da Yan, and Tianzhen Hong. Simulation of occupancy in buildings. Energy and Buildings, 87:348–359, 2015.

[45] Yoon Soo Lee and Ali M. Malkawi. Simulating multiple occupant behaviors in buildings: An agent-based modeling approach. Energy and Buildings, 69:407–416, 2014.

[46] Farhang Tahmasebi and Ardeshir Mahdavi. The sensitivity of building performance simulation results to the choice of occupants' presence models: a case study. Journal of Building Performance Simulation, 10(5-6):625–635, 2017.

[47] Da Yan, Tianzhen Hong, Bing Dong, Ardeshir Mahdavi, Simona D'Oca, Isabella Gaetani, and Xiaohang Feng. IEA EBC Annex 66: Definition and simulation of occupant behavior in buildings. Energy and Buildings, 156:258–270, 2017.

[48] Andrew Speake, Eric Wilson, Madeline Salzman, and Adam Hasz. Cost Benchmarks for Residential 50 % Energy Savings Packages. page 2012, 2012.

[49] S. Goyal, H. A. Ingley, and P. Barooah. Zone-level control algorithms based on occupancy information for energy efficient buildings. 2012 American Control Conference (ACC), pages 3063–3068, 2014.

[50] Frauke Oldewurtel, Alessandra Parisio, Colin N. Jones, Dimitrios Gyalistras, Markus Gwerder, Vanessa Stauch, Beat Lehmann, and Manfred Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. Energy and Buildings, 45:15–27, 2012.

[51] Arnab Roy and Lihui Bai. Control of HVAC System in Residential Buildings using Model Predictive Control : Plant Model Selection. pages 1538–1544, 2017.

[52] Radu Godina, Eduardo Rodrigues, Edris Pouresmaeil, João Matias, and João Catalão. Model Predictive Control Home Energy Management and Optimization Strategy with Demand Response. Applied Sciences, 8(3):408, 2018.

[53] Jan Široký, Frauke Oldewurtel, Jiří Cigler, and Samuel Prívara. Experimental analysis of model predictive control for an energy efficient building heating system. Applied Energy, 88(9):3079–3087, 2011.

[54] Radu Bălan, Joshua Cooper, Kuo Ming Chao, Sergiu Stan, and Radu Donca. Parameter identification and model based predictive control of temperature inside a house. Energy and Buildings, 43(2-3):748–758, 2011.

[55] Charles D. Corbin, Gregor P. Henze, and Peter May-Ostendorp. A model predictive control optimization environment for real-time commercial building application. Journal of Building Performance Simulation, 6(3):159–174, 2013.

[56] T. Zakula, P. R. Armstrong, and L. Norford. Modeling environment for model predictive control of buildings. Energy and Buildings, 85:549–559, 2014.

[57] Gregory S. Pavlak, Gregor P. Henze, and Vincent J. Cushing. Evaluating synergistic effect of optimally controlling commercial building thermal mass portfolios. Energy, 84:161–176, 2015.

[58] Gregor P Henze, David P Yuill, and Andrew H Coward. Development of a Model Predictive Controller for Tankless Water Heaters. 15(1):3–23, 2009.

[59] Brian Coffey, Fariborz Haghighat, Edward Morofsky, and Edward Kutrowski. A software framework for model predictive control with GenOpt. Energy and Buildings, 42(7):1084–1092, 2010.

[60] Peter May-Ostendorp, Gregor P. Henze, Charles D. Corbin, Balaji Rajagopalan, and Clemens Felsmann. Model-predictive control of mixed-mode buildings with rule extraction. Building and Environment, 46(2):428–437, 2011.

[61] Lawrence Berkeley National Laboratory, United States. Department of Energy. Office of Scientific, Technical Information, and Lawrence Berkeley National Laboratory. Environmental Energy Technologies Division. EnergyPlus Run Time Analysis. Lawrence Berkeley National Laboratory. Environmental Energy Technologies Division, Oak Ridge, Tenn;Berkeley, Calif;, 2008.

[62] Andrea E. Olsson. Particle swarm optimization: theory, techniques and applications. Nova Science Publishers, New York, 2011;2010;.