

Improving energy efficiency in office buildings through adaptive control of the indoor climate

by

H. Burak Gunay, M.A.Sc.

A thesis submitted to the
Faculty of Graduate and Postdoctoral Affairs
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy Engineering

Ottawa-Carleton Institute for Civil Engineering
Department of Civil and Environmental Engineering
Carleton University
Ottawa, Ontario
July, 2016

©Copyright
H. Burak Gunay, 2016

The undersigned hereby recommends to the
Faculty of Graduate and Postdoctoral Affairs
acceptance of the thesis

**Improving energy efficiency in office buildings through
adaptive control of the indoor climate**

submitted by **H. Burak Gunay, M.A.Sc.**
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy Engineering

Dr. José Candanedo

Dr. William O'Brien, Thesis Supervisor

Dr. Ian Beausoleil-Morrison, Thesis Supervisor

Dr. Matthew Johnson

Dr. Amir Hakami

Dr. John Gales

Professor Paul Van Geel, Chair,
Department of Civil and Environmental Engineering
Ottawa-Carleton Institute for Civil Engineering
Department of Civil and Environmental Engineering
Carleton University
July, 2016

Abstract

Zone level heating and cooling systems, lighting and blinds in offices have been either controlled by occupants or automated based on fixed setpoints and schedules. Occupants' control of these systems target maintaining a comfortable indoor climate with minimum number of control actions with little intention to save energy. This leads to the suboptimal utilization of daylight and passive solar heat gains. However, when they are automated to mitigate the extravagance in occupant behaviour, operators tend to choose conservative setpoints and schedules that maintain the comfort of the majority and minimize the frequency of complaints with little consideration to save energy again.

Although we have been treating the building controllers as our servants maintaining fixed setpoints and schedules, they represent great potential to implement distributed artificial intelligence. To this end, within the scope of this thesis, an adaptive indoor climate control tool was developed. The tool contains novel algorithms that recursively learn from the occupancy patterns, adaptive occupant behaviours, and develop an inverse model of the heat transfer problem in each zone. The algorithms were designed to be embedded inside local building controllers to undertake the learning process in real-time using a small number of low-cost building sensors. The information derived from the occupants and the building's temperature response was used to autonomously choose operating setpoints and schedules – tailored to exploit the nuances amongst subspaces in a building.

The algorithms were implemented and tested for over a year in a controls laboratory which was a shared-office space with a standalone controls network. Furthermore,

they were implemented and tested in private offices for over a year. Energy and day-lighting simulations were conducted to analyze alternative scenarios regarding the design of these algorithms. Implementation results indicate that the use of adaptive indoor climate control algorithms developed in this thesis can substantially reduce the space heating, cooling, and lighting loads in office buildings – without adversely affecting the occupant comfort. The simulation results indicate that adaptive control of the indoor climate not only reduces the expected load intensities but also reduces the risk of poor operational decisions leading to discomfort and excessive energy use.

Hofstadter's Law: It always takes longer than you expect, even when you take into account Hofstadter's Law.

Douglas Hofstadter

Acknowledgments

I would like to thank my supervisors Drs. Liam O'Brien and Ian Beausoleil-Morrison. They simply went above and beyond what one can ever expect from a supervisor. You can feel them being truly happy with your successes and sad with your setbacks. I wish to have them as my lifelong mentors and friends.

I would like to acknowledge my co-workers at Carleton Building Research Groups: Brent Huchuk, Jayson Bursill, Dr. Sara Gilani, Isis Bennet, Chris Baldwin, Zach Burgoyne, Cynthia Cruickshank, Shawn Shi, Adam Wills, Sébastien Brideau, Aly Abdelalim, Anthony Fuller, Austin Selvig, and Dr. Scott Bucking. Also, I would like to acknowledge Philippe Bisaillon for his friendship, critical assistance, and knowledge.

I would like to acknowledge the entire team of Carleton Facilities Management and Planning. Particularly, the support from visionaries like Darryl Boyce, Lisa Paterick, Scott Macdonald, and Philip Mansfield made this research project possible.

I would like to acknowledge Delta Controls and CopperTree Analytics for their financial and in-kind contributions. Particularly, the technical knowledge of Chris Kwong, Timo Kinna, and Carl Neilson in building controls was an essential anchor to ensure the practicality of the findings of this research project. I would like to thank Regulvar – particularly Stéphan Riffault, Gilles Lebrun – for training me and donating/installing automated-shades.

I would like to thank my colleagues from IEA Annex 66 for their critical assistance; in particular Dr. Ardeshir Mahdavi, Dr. Darren Robinson, Farhang Tahmasebi, Simona D'Oca, Dr. Tianzhen Hong, and Dr. Da Yan.

I would like to thank to my other industry collaborators: Autodesk Research (Rhys Goldstein, Azam Khan, Simon Breslav), NRCan (Phylroy Lopez, Meli Stylianou), and Green Power Labs.

I would like to thank NSERC for supporting my research project through doctoral scholarships (CGS and PGS) and through a research grant (CRD). I would like to thank ASHRAE for the Grant-in Aid Award, Carleton for the internally administrated awards, and IBPSA for the travel award.

I would like to acknowledge the staff at Civil and Environmental Engineering Department in Carleton University: Payal Chadha, Kay Casselman, Sara Seiler, Ben Griffin, Stan Conley, and Kristin Holtz.

Finally, thank you to my family for all their support and encouragement through not just this process but everything leading up until this point. Thanks so much Yagmur Babaoglu, your presence made everything easier.

Table of Contents

Abstract	iii
Acknowledgments	vi
Table of Contents	vii
List of Tables	ix
List of Figures	x
Nomenclature	xvii
1 Introduction	1
1.1 Background on indoor climate control	2
1.1.1 Control for thermal comfort and IAQ	2
1.1.2 Control for visual comfort	9
1.2 Motivation and research objectives	13
1.3 Research methodology	17
1.4 Document Structure	23
2 Learning from occupants' presence	25
2.1 Detecting presence	25
2.1.1 Literature review	26
2.1.2 Methodology	30
2.1.3 Results and discussion	31
2.2 Predicting presence	46
2.2.1 Literature review	46
2.2.2 Model forms	49
2.2.3 Diversity aspects	53
2.2.4 Occupancy-learning algorithm	58
2.3 Summary	62
3 Learning from adaptive behaviours	66
3.1 Literature review: adaptive behaviour modelling	66
3.2 Development of the thermostat learning algorithm	71
3.2.1 Analysis of the observational data	71
3.2.2 Thermostat learning algorithm	79

3.3	Development of the lighting and blinds learning algorithm	87
3.3.1	Analysis of the observational data	87
3.3.2	Lighting and blinds learning algorithm	102
3.4	Summary	106
4	Learning from the temperature response	109
4.1	Literature review	109
4.2	Objectives and scope	113
4.3	Methodology	114
4.3.1	Monitored offices and data	114
4.3.2	Control-oriented models	121
4.3.3	Recursive parameter estimation	130
4.4	Predictive accuracy of the models	135
4.5	Residual analysis and parameter estimates	141
4.6	Summary	149
5	Adaptive control of the indoor temperature	152
5.1	Literature review	152
5.2	Objectives and scope	155
5.3	Adaptive temperature control algorithm	155
5.4	Implementation results and comfort implications	165
5.4.1	Laboratory implementation	167
5.4.2	Field implementation	174
5.5	Estimating energy-savings potential	189
5.6	Summary	202
6	Adaptive control of the indoor illuminance	205
6.1	Literature review	205
6.2	Objectives and scope	209
6.3	Lighting and Blinds Control Algorithm	210
6.4	Implementation results and comfort implications	213
6.4.1	Laboratory implementation	213
6.4.2	Field implementation	219
6.5	Estimating energy-savings potential	223
6.6	Summary	236
7	Conclusions and future work	238
7.1	Future work	245
	List of References	250

List of Tables

1.1	An overview of the dataset used during the occupant model and data-driven greybox building model development process.	19
2.1	Characteristics of the sensory data used in occupancy detection study.	32
2.2	The parameters of the model shown in Eqn. 2.1. The parameters are shown in mean \pm standard deviation format.	43
2.3	The occupancy parameters in the monitored offices.	57
3.1	List of rooms from which the data were collected.	73
3.2	List of rooms from which the data were collected.	75
3.3	The logistic regression parameters for light switch behaviour in individual rooms as a function of the ceiling illuminance.	100
4.1	Characteristics of the data.	117
4.2	Model parameters and inputs.	129
4.3	The MAE of the model predictions.	139
4.4	Reciprocals of the model parameters x_1	147
5.1	Mean and standard deviation of the indoor temperature in CR1 to CR8 and R1 to R8.	185
5.2	The mean blind occlusion rates in R1 to R8 – rounded to the nearest multiple of 0.25.	192
6.1	Summary of the performance metrics of the self-adaptive lighting learning control algorithm during the field implementation.	224
6.2	Lighting and blinds use model parameters. $N(\mu, \sigma)$ stands for normal and u stands for uniform distributions.	229
6.3	Studied control scenarios.	230

List of Figures

1.1	A schematic of a simple VAV AHU system.	4
1.2	Weekday occupancy profiles in three different office buildings. The annotations occupancy with flexible work hours and nighttime setback-to-setpoint transition time do not emerge from neither of the three buildings. They illustrate the spread of operating hours due to flexible work hours and setback-to-setpoint transition periods.	5
1.3	A statistical model predicting the number of complaints.	8
1.4	An illustration of the slow temperature response of buildings.	8
1.5	Two examples illustrating the conservativeness in technician choices in an academic office building.	15
1.6	The controls laboratory.	18
1.7	The sensors and actuators in the controls laboratory.	21
1.8	The offices employed in the field implementation. The red lines represent the eight rooms used in the implementation for the temperature controls and the yellow dashed lines represent the five rooms used in the implementation for the lighting controls.	22
1.9	The sensors and actuators in offices whereby the algorithms were implemented.	23
2.1	An example where the occupant taped over the motion detectors (credit: Sara Gilani).	26
2.2	Sensor locations and their coverage ranges.	31
2.3	Movement detections over an illustrative 12 h period by (a) PIR 1, (b) PIR 2, (c) PIR 3; and (d) the ground-truth occupancy for the same period.	33
2.4	Accuracy of presence and absence detections as a function of the delay value using (a) PIR 1, (b) PIR 2, (c) PIR 3, (d) all three PIRs simultaneously.	35
2.5	Frequency of movement detections during intermediate occupancy by using (a) PIR 1, (b) PIR 2, (c) PIR 3, and (d) all three PIRs simultaneously.	38
2.6	Door position, the movement detections by PIR 2, and the ground truth.	39
2.7	Effect of door sensor on the accuracy of presence predictions at different PIR delay values.	40
2.8	Mass transfer processes affecting the CO ₂ concentration in a perimeter office space.	41

2.9	Indoor and supply air CO ₂ concentrations, door position, and number of occupants on weekdays.	42
2.10	The predictions made by the model shown in 2.1 using the parameter estimates shown in Table 2.2 for (a) CO ₂ 1 and (b) CO ₂ 2 sensors. . .	44
2.11	The influence of using CO ₂ -based occupancy detections together with PIR 2 over the absence detection accuracy.	45
2.12	Weekly occupancy schedules in the studied offices.	50
2.13	Discrete-time Markov model predicting the likelihood of a first arrival on a weekday in the next one hour.	51
2.14	Discrete-time Markov model predicting the likelihood of the last departures in a weekday.	52
2.15	Survival model predicting the duration of an uninterrupted intermediate occupancy period.	52
2.16	Survival model predicting the duration of an intermediate vacancy period. . .	53
2.17	Occupancy schedules in individual offices. Each plot represents one of the 16 occupants. The week starts on Monday.	55
2.18	Autocorrelation of the occupancy time-series in individual offices. . .	56
2.19	Effect of zone granularity on the occupancy parameters $t_{arr,early}$, $t_{arr,late}$, and $t_{dpt,late}$	59
2.20	Effect of zone granularity on the occupancy parameters p_{abs} (h) and D_{abs} (h).	59
3.1	Interactions between a building and its occupants.	67
3.2	Thermostat use behaviour during the heating season over an illustrative ten day period in room 4.	76
3.3	Distribution of the indoor temperatures just before the thermostat override instances.	77
3.4	Indoor and outdoor temperatures at thermostat keypress instances. .	78
3.5	Probability of (a) increasing and (b) decreasing the temperature setpoint in the next one hour. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bT_{in})}}$. The properties of the regression parameters were annotated in the figure.	80
3.6	Probability of (a) increasing and (b) decreasing the temperature setpoint at arrival. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bT_{in})}}$. The properties of the regression parameters were annotated in the figure.	81
3.7	An illustration of the recursive parameter update process in a logistic regression model.	82
3.8	The recursive evolution of the setpoint decrease and increase models for room 16 in time over a year. The darkness of the line plots increases as time progresses and the thick red line plots represent the last state of the model.	84

3.9	The recursive evolution of the parameter estimates a and b for the models shown in Figure 3.8.	84
3.10	A comparison of the recursively estimated model with the measured discrete likelihood densities and the traditionally developed logistic regression models (using a batch of observations and the MLE method) for room 16.	86
3.11	The weekly lighting schedules in the ten monitored offices.	89
3.12	Weekly lighting schedules in individual rooms plotted against weekly occupancy schedules.	90
3.13	The mean weekly lighting use with respect to the mean weekly occupancy when the peak daily solar altitude was higher than 40° (March to September) and lower than 40° (October to February).	91
3.14	The ratio of lights on during occupied hours as a function of the solar irradiance on the facade in individual offices.	92
3.15	Mean blind occlusion rate as a function of the solar irradiance on the facade.	93
3.16	Probability of closing blinds in the next 60 min as a function of the incident solar irradiance on the facade. The univariate logistic regression model is in the following form: $p = \frac{1}{1+e^{-(a+bS_{\text{rad}})}}$. The properties of the regression parameters were annotated in the figure.	94
3.17	A survival model for the lifetime of a blinds position and a Weibull distribution fit for this model. The Weibull model is in the following form: $p = 1 - e^{-(t/A)^B}$. The properties of the regression parameters were annotated in the figure.	95
3.18	The distribution of the blinds closing and opening increments at four discrete levels.	96
3.19	Probability of switching on the lights (a) in the next 15 min (discrete-time) and (b) at arrival (discrete-event) as a function of the ceiling illuminance. The univariate logistic regression model is in the following form: $p = \frac{1}{1+e^{-(a+bE_{\text{lux}})}}$. The properties of the regression parameters were annotated in the figure.	98
3.20	Probability of switching on the lights (a) in the next 15 min (discrete-time), (b) at arrival, (c) during intermediate occupancy in individual rooms. The univariate logistic regression model are in the following form: $p = \frac{1}{1+e^{-(a+bE_{\text{lux}})}}$. The thick red line represents the mean of the ten occupants. The regression parameters are tabulated in Table 3.3. For comparison purposes, Reinhart's light switch model for arrival and intermediate occupancy were included. Note that Reinhart's model was developed upon workplane illuminance measurements, whereas the models shown in this study were developed upon indoor illuminance measurements taken from the ceiling.	99

3.21	Probability of a light switch off at departure with respect to the duration of absence following the departure. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bD_{\text{abs}})}}$. The regression parameters are tabulated in the figure. For comparison, another light switch off model extracted from an independent study is shown in the figure.	101
3.22	The recursive update procedure for the logistic regression models predicting the light switch on or the blinds closing behaviour.	104
3.23	(a) The recursive evolution of the light switch on models in one of the ten offices over a year. The darkness of the line plots increases as time progresses and the thick red line plots represent the final state of the model. (b) A comparison of the recursively estimated model with the measured discrete likelihood densities and the traditionally developed logistic regression models (using a batch of observations and the MLE method) for the same room.	105
4.1	Photos of the monitored offices.	114
4.2	Position of the sensors and actuators in the monitored offices.	115
4.3	Measured indoor (y) and outdoor temperatures (u_1), indoor illuminance (u_2), and vertical solar irradiance on the facade (u_3).	118
4.4	Monitored occupancy (u_4), electric lighting state (u_5), and plug-in equipment load (u_6).	119
4.5	Measured radiant panel heater state (u_7), VAV terminal unit's discharge air temperature (u_8) and pressure (u_9).	120
4.6	(a) First and (b) second-order lumped parameter thermal network model of a perimeter office space.	122
4.7	An illustration of the prediction error (by one of the models) at two random two day prediction time horizons in room 1.	136
4.8	Histogram plots illustrating the distribution of the offline prediction errors over the two day time horizons (bin size 1°C) in room 1.	137
4.9	Scatter plots illustrating the models' prediction residuals as a function of the incident solar radiation and the outdoor temperature, when they were trained with the EKF. The solid lines represent the least squares line fit through these scatter plots.	143
4.10	Scatter plots illustrating the models' prediction residuals as a function of the incident solar radiation and the outdoor temperature, when they were trained with the PF. The solid lines represent the least squares line fit through these scatter plots.	144
4.11	Autocorrelation plots for the models' prediction residuals. (a) The single-state models trained by the EKF, (b) the two-state models trained by the EKF, (c) the single-state models trained by the PF, and (d) the two-state models trained by the PF.	145

4.12	The evolution of the model 4's parameters in time when it is trained by the EKF.	148
5.1	The temperature control algorithm and its relationship with the learning algorithms.	156
5.2	Illustration of the receding time horizons employed in the temperature control algorithm.	160
5.3	The three scenarios the weekday temperature schedule can follow in the temperature control algorithm.	162
5.4	Error in predicting the outdoor temperature and the indoor illuminance three hours ahead by using the persistence and the prior moving average methods.	164
5.5	Illustrative characteristic profiles for (a) daylight and (b) outdoor temperature sensor readings for different months of the year.	166
5.6	An infographic summarizing the implementation results in the controls laboratory.	168
5.7	Mean weekly occupancy schedule in the controls laboratory.	169
5.8	Distribution of the first arrivals and the last departures in the controls laboratory.	170
5.9	Distribution of the duration of the intermediate vacancy periods in the controls laboratory.	170
5.10	Discrete-time Markov thermostat use model predicting the likelihood of a thermostat override in the next 30 min for the controls laboratory.	171
5.11	Distribution of the indoor temperature at the keypress instances in the controls laboratory.	172
5.12	Number of thermostat override actions in 100 h of occupancy between different temperature ranges in the controls laboratory.	172
5.13	Comparison of the indoor temperatures with the default and the adaptive learning temperature control algorithms	175
5.14	An infographic summarizing the field implementation results. The order of the rooms R1 to R8 and CR1 to CR8 were shuffled to preserve the anonymity of the occupants.	177
5.15	An illustrative example of the setpoint evolution process inside the controller database serving R5.	179
5.16	Illustration of the free-running temperature response in R8 prior to implementation.	182
5.17	Diversification of the radiant panel heater profiles with the adaptive control algorithm.	184
5.18	Indoor temperatures during intermediate occupancy in (a) the heating season and (b) the cooling season. Indoor temperatures at first arrivals in (c) the heating season and (d) the cooling season.	187
5.19	CO ₂ concentration during the occupied hours in R1 to R8 – except the unoccupied office R6.	189
5.20	Sensitivity of the EnergyPlus results to the infiltration rate assumption.	192

5.21	An overview of the EnergyPlus model.	192
5.22	Average weekly occupancy, plug-in equipment load, and lighting schedules.	193
5.23	Correlation between the setpoint/scheduling decisions and the cooling load/discomfort in the default temperature control algorithm. The histogram plots present the distribution of the cooling load and discomfort metric results with individual run cases.	196
5.24	Correlation between the setpoint/scheduling decisions and the heating load/discomfort in the default temperature control algorithm. The histogram plots present the distribution of the heating load and discomfort metric results with individual run cases.	197
5.25	Pearson correlation coefficients between the setpoint/scheduling decisions and (a) the cooling/discomfort and (b) heating/discomfort. The axes captions were abbreviated as follows: heating setpoint (HSP), cooling setpoint (CSP), heating setback (HSB), cooling setback (HSB).	198
5.26	The correlation between the parameters affecting the performance of the adaptive temperature control algorithm and the performance of the algorithm.	200
5.27	Comparison of the simulation run cases with the default and the adaptive temperature control algorithms.	201
6.1	Two common configurations for daylight-integrated blinds and lighting automation with examples from the literature: (a) Closed-loop systems and (b) Open-loop systems	207
6.2	The relationship of the lighting and blinds control algorithm with the learning algorithms.	210
6.3	The cumulative distribution of the increase in the ceiling-mounted photodiode sensors' readings in response to switching on the lights in ten identical West-facing offices.	212
6.4	The ratio of a ceiling-mounted photodiode sensor's readings before and after closing the blinds in a Southwest-facing office space. The whiskers indicate the 10 th and 90 th percentiles.	212
6.5	An overview of the controls laboratory. The sensors and actuators employed in testing the lighting and blinds control algorithm. A summary of the implementation results.	214
6.6	Verification of the recursive learning algorithms implemented inside the controllers against those estimated offline using batch datasets. discrete-time Markov models predicting (a) a blinds closing action and (b) a light switch on action in the next 30 min.	216
6.7	Ratio of occupied duration the lights were on with respect to the incident solar irradiance on the facade.	218

6.8	Comparison of the discrete-time Markov models predicting the blinds closing actions in the next 30 min by using a ceiling-mounted photodiode and a pyranometer sensor measuring the transmitted solar irradiance on the interior surface of the window. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bE_{lux} \text{ or } S_{rad})}}$. The properties of the regression parameters were annotated in the figure.	220
6.9	Field implementation results. The ratio of the occupied duration the lights were on as a function of (a) the ceiling-mounted photodiode sensor readings and (b) the solar irradiance on the facade. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bE_{lux} \text{ or } S_{rad})}}$. The properties of the regression parameters were annotated in the figure.	222
6.10	The daylighting model geometry and material properties.	225
6.11	Stochastic occupancy time-series generation process.	228
6.12	Duration of occupancy for a synthetic population of 300 occupants. The log-normal model is in following form: $p(D_{occ}) = \frac{1}{D_{occ}scale\sqrt{2\pi}}e^{-\frac{(\ln D_{occ}-location)^2}{2scale^2}}$. The regression parameters are annotated in the figure.	230
6.13	The distribution of duration of lighting use by the simulated occupants with the control scenarios (1) to (5) listed in Table 6.3 are shown in (a) to (e), respectively. The log-normal models are in following form: $p(D_{occ}) = \frac{1}{D_{occ}scale\sqrt{2\pi}}e^{-\frac{(\ln D_{occ}-location)^2}{2scale^2}}$. The regression parameters are annotated in the figure.	233
6.14	Simulated mean blind occlusion rate during occupied hours with the manual blinds control (scenarios 1 to 4 in Table 6.3) and with the adaptive occupant-learning blinds control (scenario 5 in Table 6.3). The whiskers indicate 10 th and 90 th percentile in the distribution.	234
6.15	Simulated frequency of (a) light switch on and (b) blinds closing actions by different control algorithms.	236

Nomenclature

Abbreviation	Full Name
ACH	Air changes per hour
AHU	Air-handling unit
AIC	Akaike information criterion
BAS	Building automation system
BACnet	ASHRAE/ANSI/ISO standard protocol for building controls
BPS	Building performance simulation
EKF	Extended Kalman Filter
IAQ	Indoor air quality
HVAC	Heating, ventilation, and air conditioning
MAE	Mean absolute error
MLE	Maximum likelihood estimation
MPC	Model-based predictive controller
PIR	Passive-infrared
PF	Particle filter
VAV	Variable air volume

Symbols	Full Name
clo	Clothing insulation unit
C_{sa}	Supply air CO ₂ concentration (ppm)
C_{in}	Indoor CO ₂ concentration (ppm)
C_{out}	Outdoor CO ₂ concentration (ppm)
C_m	Measured indoor CO ₂ concentration (ppm)
C	Thermal capacitance (J/°C)
D_{abs}	Longest expected intermediate vacancy period (h)
D_{int}	Longest expected intermediate presence period (h)
E_{lux}	Indoor illuminance (Lux)
$E_{light,sp}$	Brightest expected conditions to cause a light switch on (Lux)
$E_{blinds,sp}$	Darkest expected conditions to cause a blinds closing (Lux)

H	Measurement model
N_{occ}	Number of occupants in a room
p_{sa}	Supply air pressure (Pa)
$p_{\text{da,sp}}$	Discharge air pressure setpoint of VAV (Pa)
p	Probability
p_{abs}	Probability of absence
\mathbf{p}	Presence
P	Model covariance
P	Proportional control loop
PI	Proportional integral control loop
PID	Proportional integral derivative control loop
Q_{sol}	Solar gains (W)
Q_{rads}	Heat input from radiant panels (W)
Q_{vav}	Heat input from the VAV (W)
Q_{int}	Casual heat gains (W)
\mathbf{Q}	Covariance of the model error
R	Thermal resistance ($^{\circ}\text{C}/\text{W}$)
\mathbf{R}	Thermal resistance ($^{\circ}\text{C}/\text{W}$)
S_{rad}	Solar irradiance ($\text{W}\cdot\text{m}^{-2}$)
SHGC	Solar heat gain coefficient
t	Time (s)
$t_{\text{arr,early}}$	Earliest expected arrival time of the day (h)
$t_{\text{arr,late}}$	Latest expected arrival time of the day (h)
$t_{\text{dpt,late}}$	Latest expected departure time of the day (h)
t_{arr}	Time of arrival (h)
t_{dpt}	Time of departure (h)
T_{f}	Free running indoor temperature ($^{\circ}\text{C}$)
T_{in}	Indoor temperature ($^{\circ}\text{C}$)
$T_{\text{da,sp}}$	Discharge air temperature setpoint of VAV ($^{\circ}\text{C}$)
T_{hsp}	Warmest expected temperature leading to a setpoint increase ($^{\circ}\text{C}$)
T_{csp}	Coldest expected temperature leading to a setpoint decrease ($^{\circ}\text{C}$)
\mathbf{T}_{sp}	Temperature setpoint ($^{\circ}\text{C}$)
u_1	Model disturbance outdoor temperature ($^{\circ}\text{C}$)
u_2	Model disturbance indoor illuminance (lux)
u_3	Model disturbance solar irradiance ($\text{W}\cdot\text{m}^{-2}$)
u_4	Model disturbance occupancy state
u_5	Model disturbance lighting state
u_6	Model disturbance plug loads (W)
u_7	Model disturbance heater state

u_8	Model disturbance VAV discharge air temperature (°C)
u_9	Model disturbance VAV pressure (Pa)
v	Measurement noise
w	Standard Wiener process
x	Parameter of a data-driven model
y	Measured indoor temperature (°C)
ΔU	Temperature change due to the heating/cooling loads (°C)
Δt	Timestep (s)
k	Timestep index
μ_{arr} (h)	Mean arrival time (h)
μ_{dpt} (h)	Mean departure time (h)
σ_{dpt} (h)	Standard deviation of the departure time (h)
σ_{arr} (h)	Standard deviation of the arrival time (h)

Glossary

Adaptive behaviours Actions that occupants undertake to restore their comfort (e.g., light switch, window opening).

Adaptive comfort Comfort approach in which occupants' comfort expectations change based on the changes in the indoor and outdoor environmental conditions

Adaptive control The control method in which the parameters adapt to a stochastic controlled system.

Artificial neural networks A class of models used to approximate the behaviour of an unknown system depending on a number of inputs and outputs.

Backpropagation A method commonly used in training artificial neural networks.

BACnet A data communication protocol for building automation and control networks.

Bayesian inference A statistical method used to update the probability for a hypothesis as more information becomes available.

Bernoulli random process Sequence of binary random variables.

Blackbox models Non-physical models in which a system's response is represented in terms of its inputs and outputs.

Extended Kalman Filter A method used in parameter and state estimation of a model.

Generalized linear models A class of models used to represent non-linear and bounded response variables.

Greybox models A class of physical models in which only a partial information is available to define the system.

Inverse modelling The process of identifying the parameter and unknown state information through a data-driven model's response.

Logistic regression A regression model predicting the probability of a binary response variable.

Markov random process Sequence of stochastic state transitions.

Maximum likelihood estimation A method of estimating the parameters of a statistical model.

Model-based predictive control A control method in which a model of the controlled system is used to meet setpoints.

Particle Filter A sequential Monte Carlo method for state and parameter estimation.

Survival models A statistical method used to analyze the expected duration of time until an event happens.

Whitebox models A class of physical models in which all the necessary information is available to define the system.

Chapter 1

Introduction

People spend 90% of their time indoors [1,2]. Thus, it is imperative to achieve comfortable, healthy, and productive indoor environments. Today, more than 35% of the primary energy use in North America is devoted to indoor climate control for thermal and visual comfort, and indoor air quality (IAQ) – split evenly between residential and commercial buildings [3,4]. Therefore, indoor climate control in buildings plays a significant role in our economic and environmental impact – responsible for 15-20% of the CO₂ emissions and a major driver for new energy infrastructure [2]. Although we can reduce this impact through better building designs, indoor climate control in existing buildings will remain as a substantial economic and environmental burden with the building floor area increase rate less than 2% annually [3]. In fact, anecdotal evidence suggests that improperly controlled building equipment wastes 30 to 50% of the energy used in commercial buildings [5–7]. Adoption of innovative indoor climate control strategies represents a low-cost and non-invasive means to address these inefficiencies in both new and existing buildings. In brief, these strategies mean better controls sequencing of heating, ventilation, and air-conditioning (HVAC) and lighting equipment. And, the research on this topic has been focused on developing transferable and scalable methodologies reaching near-optimal controls sequences for common building equipment. In particular, the new research idea in this thesis is

to exploit the untapped potential in the sensory data and distributed computational power in existing building automation systems (BAS) to create indoor climatic conditions tailored for the characteristics of the terminal equipment, zone occupancy and comfort expectations autonomously.

1.1 Background on indoor climate control

This section provides a background on the indoor climate control in commercial buildings. In a few places, terms related with the HVAC, comfort, and IAQ were used. When needed, the reader should refer to the ASHRAE Handbook of Fundamentals [8]. Depending on how fast the controlled variable responds to the building actuators, different control challenges were reported in the reviewed literature. Therefore, the discussion about these applications was separated into two groups: (1) control for thermal comfort and IAQ, and (2) control for visual comfort.

1.1.1 Control for thermal comfort and IAQ

Buildings are subject to environment and occupant-driven thermal loads. Environment-driven loads are the heat losses through a building's envelope due to conduction and outdoor air-infiltration, and the heat gains through exterior windows due to transmitted solar radiation. Occupant-driven loads are the metabolic heat emitted from occupants, plug-in equipment, and lighting. Occupants and some building components (e.g., paints, furniture) also emit pollutants, mainly CO₂ and volatile organic compounds. The HVAC equipment in buildings is controlled to maintain thermal comfort and IAQ in each occupied thermal zone at desired conditions.

The simple schematic shown in Figure 1.1 illustrates a generic variable air volume (VAV) air-handling unit (AHU) system – a common HVAC configuration in North

American commercial buildings [9]. An AHU is a system level HVAC equipment serving at least one thermal zone. It exhausts a portion of the air returned from thermal zones and mixes the rest of the returned air with some outdoor air. A portion of the sensible heat from the exhaust air can be used to temper the temperature of the outdoor air added to the ventilation air during the heating season – if the AHU has the heat recovery option. Normally, the proportion of the outdoor air in the supply air is selected based on the minimum ventilation requirements for IAQ. When the outdoor air is advantageous to partly or completely remove the cooling demand, the proportion of the outdoor air can be increased. This operating mode is known as the economizer mode – distinct from the heating and cooling modes. The supply air is first filtered, and then conditioned by passing through heating and cooling coils. In many institutional buildings, the heating and cooling are provided from central boiler or chiller plants. The heating plant remains operational only during a specified heating season, and the cooling plant remains operational only during a specified cooling season. Through ductwork, air is brought to individual VAV units. A VAV determines the amount of supply air rate to the thermal zone by modulating its damper. Different thermal zones served by the same AHU can have quite different heating and cooling load demands. To meet these interzonal sensible load differences, there can be reheat coils and radiant heating or cooling panels controlled by separate valves in each zone. The common control logic for an AHU is to meet the sensible heating/cooling load of the least demanding zone and to use the zone level HVAC equipment to provide the rest of the heating/cooling demand in each thermal zone.

The energy used by a VAV AHU system to control indoor temperature and IAQ depends on the amount of the heating and cooling produced by the plant and the plant efficiency, the electricity used by the fans and the pumps to circulate the air




Figure 1.1: A schematic of a simple VAV AHU system.

through the ductwork and the hot or chilled water through the pipes. The energy consumption will reduce as the HVAC loads reduce – depending on the equipment's part-load efficiency. Aside from the environmental and occupant-driven loads and physical properties of the building, the HVAC loads in a thermal zone depend on the operating setpoints and schedules. Commonly, they refer to heating and cooling temperature setpoints and setback schedules, and minimum and maximum ventilation rates in each thermal zone. They are determined by a controls technician or an operator in the BAS – occupants may or may not be given a temporary override option to change these setpoints through their thermostats. Determining appropriate setpoints and schedules for HVAC equipment is challenging due to three factors.

1. Diversity in a building's usage: Because more than a quarter of the North American workforce reported having flexible work hours [10–12], job specific restrictions and personal preferences translate into occupancy profiles. Consequently,

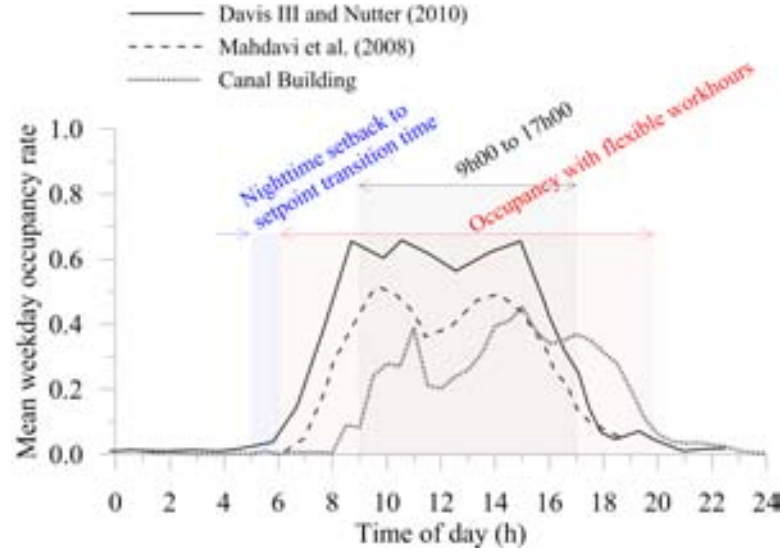


Figure 1.2: Weekday occupancy profiles in three different office buildings [13, 14].

The annotations occupancy with flexible work hours and nighttime setback-to-setpoint transition time do not emerge from neither of the three buildings. They illustrate the spread of operating hours due to flexible work hours and setback-to-setpoint transition periods.

in commercial buildings the occupancy spreads well before 9 am and well after 5 pm [13, 14] (see Figure 1.2). In some cases, occupancy may even extend to weekends and holidays [15]. The shift towards flexible work schedules from standard rigid work schedules makes it harder to predict a building's occupancy [16]. As a result, operators tend to choose conservatively long operating schedules in order to accommodate this diversity in building usage [7].

2. Diversity in occupants' temperature preferences: Thermal comfort depends on a large number indoor, outdoor, and contextual factors including but not limited to the clothing insulation [17], air speed, wall, window, ceiling, floor, and air temperatures, relative humidity, type of indoor activities undertaken, gender, age [8], adaptation to ambient climatic conditions [18], state of occupancy (e.g., just arrived, intermediate occupancy) [19]. However, in a thermal zone we typically control indoor

temperature using a single temperature sensor embedded inside a wall thermostat – which also inevitably has errors (e.g., calibration, drift, bias). In absence of this overwhelmingly long list of factors affecting occupants' thermal comfort, operators are motivated to choose a conservative set of temperature setpoints that minimizes the total number of complaints – with little consideration to save energy. For example, Federspiel et al. [19] investigated the number of unsolicited complaints in six office buildings and developed a model predicting the number of thermal comfort complaints (see Figure 1.3) – which was later adopted by ASHRAE Standard 55 [20]. The temperatures that minimize the total number of complaints were between 22 and 23°C. This does not imply that all occupants prefer temperatures in that range. In fact, a considerable number of complaints were observed even at this temperature range. This underlines the challenge in finding a single temperature setpoint value that works for everyone. In line with this, the pioneering work by Fanger [21] reported that there are individual differences in experiencing thermal environments and no thermal environment can satisfy everybody.

A simple solution to this problem stands out as the control of room thermostats by the occupants. Previous research pointed out that controllability of thermostats plays a vital role in achieving individual comfort expectations [22]. Humphreys and Nicol [23] observed that the same room temperature can be perceived more acceptable, if it is chosen – rather than being imposed. It was found that individual control of the indoor temperature improves productivity [24, 25] and employee satisfaction [26, 27]. Wyon [28] identified that $\pm 3^\circ\text{C}$ of individual control over the default temperature settings substantially improves occupant satisfaction.

Despite these benefits, the manual control of thermostats has been commonly avoided in controls implementations. One of the main reasons is that most occupants have very little knowledge of the HVAC systems and the thermal inertia of the

office building they work in [29]. Consequently, they tend to overcompensate when increasing or decreasing the setpoints – hoping to accelerate the heating/cooling process according to two thermostat use surveys [30, 31]. This can result in suboptimal control of the indoor temperature – both from the energy use and comfort perspectives. In addition, many occupants tend not to fine-tune their thermostat settings actively to exploit the seasonal variations in their clothing levels. For example, a survey-based study conducted by Karjalainen [30] identified that more than 60% of the office occupants reported interacting with their thermostats less frequently than monthly. They rather set their thermostats such that they feel thermally neutral in both heating and cooling seasons with minimum number of control actions – with little intention to save energy. In fact, even a thin cardigan or a sweater (~ 0.25 clo [20]) can change the optimal operative temperatures by about 1.5°C [20]. Schiavon and Lee [17]'s empirical dynamic clothing insulation model proposes that an average office occupant's clothing insulation level increases by about 0.6 clo seasonally – from a cooling (20 to 30°C) to a heating season (-10 to -20°C). However, this value can vary substantially depending on a large number of contextual factors – e.g., dress code, personal preferences, shared/private office environments.

3. Poor understanding of thermal dynamics and system's capacity: Heat transfer is a transient process. Particularly in buildings with large concrete slabs, warming or cooling the space may take hours. For example, Figure 1.4 illustrates the results from an experiment that was conducted in an empty perimeter office space during a heating season. The radiant panel heater was switched on when the indoor temperature was 16.5°C for 20 h – while keeping the VAV damper closed and the reheat coil off. In the

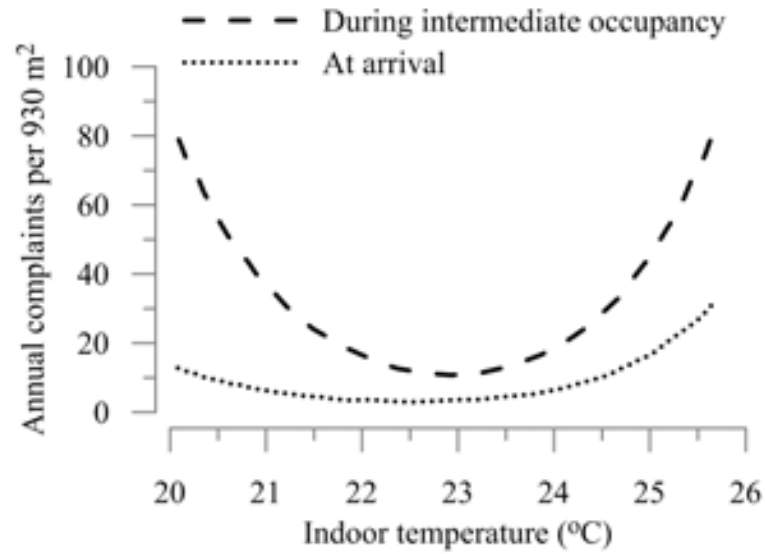


Figure 1.3: A statistical model predicting the number of complaints. (Taken from [19])

end of this period, the room temperature barely reached 19.0°C. When the radiant panel heater was switched off again, it took 12 h before the temperature return to 16.5°C.

Operators and controls technicians need to consider the length of the temperature transition period after a temperature setback period – so that the setpoints are met before occupants arrive. The temperature transition period depends on a number

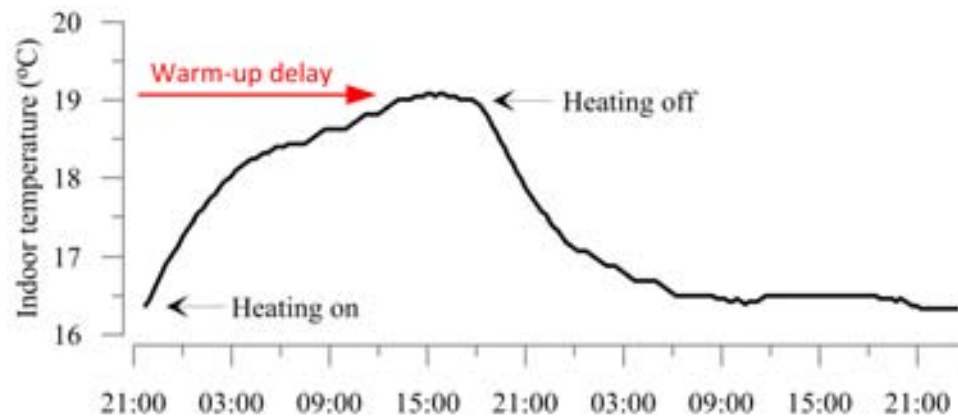


Figure 1.4: An illustration of the slow temperature response of buildings.

of factors: (a) terminal HVAC unit capacity, (b) thermal capacitance and its distribution (e.g., exposed vs. carpeted slab), (c) thermal insulation and airtightness of the building envelope, (d) ambient environmental conditions (e.g., solar radiation, temperature, wind), (e) lighting and plug-in appliance loads, and (f) shading position. These factors not only change in time (seasonally and daily) but they can also vary significantly between thermal zones. Although optimal start/stop algorithms – which dynamically determine the setpoint transition periods – have been commercially available, they often require manual tuning and setup (e.g., operators may need to guess the amount of thermal mass) and they often employ inaccurate simplistic methods [16, 32, 33]. Consequently, operators need to conservatively guess a fixed temperature transition period, so that all spaces reach to the setpoint temperature from the setback temperature before occupants start arriving [34].

1.1.2 Control for visual comfort

Window blinds and electric lighting are controlled to maintain a visually comfortable indoor environment. Traditionally, lighting and blinds are used manually by occupants. Simply put, they switch on electric lighting when it feels too dark or they close blinds when it feels too bright. Light switch on and blinds closing behaviours are associated with visual discomfort (e.g., light switch on in the absence of visual contrast or daylight, blinds closing upon glare). However, light switch off and blinds opening behaviours are not related with a physical discomfort. Occupants' light switch off behaviour can be explained with environmental or financial awareness [35, 36]. It is worth noting that occupants in commercial buildings are not directly affected by the cost of lighting electricity use. The primary reason for occupants to open their blinds is to increase their view and connection to outdoors. In absence of physical motivating factors, if lights are switched on, occupants almost never switch them off during

intermediate occupancy [37]. In fact, in a case study Pigg et al. [38] observed that the probability that an occupant leaves the lights on upon taking an intermediate break shorter than one hour is more than 0.75. This was about 0.50 for breaks between two to four hours. Similar to the infrequent use of lighting, if blinds are closed, occupants rarely open them until the next time they arrive [39, 40]. They almost never change their blind positions more than once a day [39, 41]. According to Bordass et al. [42], occupants – especially in shared offices – position their blinds to mitigate worst-case visual conditions. In brief, occupants' goal is to avoid visual discomfort at minimum number of interactions with lighting and blinds – with little consideration to exploit daylight to offset electric lighting.

In recognition of this, lighting and (although less commonly) window blinds have been automated. A recent review of the literature [43] indicates that many of the existing blinds and lighting automation systems fail to improve occupants' visual comfort and to reduce the lighting energy use – as intended. These automation problems result in temporary or permanent occupant overrides. For example, Heerwagen and Diamond [44] discuss how occupants taped over sensors to defeat automated lighting controls. In another case, Reinhart and Voss [45] reported that occupants overrode 88% of the attempts of an automation system to close window blinds. Other studies [46, 47] noted that automated window blind and lighting controls were sometimes deactivated due to complaints or had to be re-commissioned to tailor their behaviour to individual occupants. Perhaps somewhat surprisingly, manually controlled lighting fixtures which do not even meet the lighting standards can be perceived more satisfactory than the daylight-integrated automated lighting controls [22]. Gunay et al. [48] provided three explanations of occupants' discontent with automated lighting and blinds controls.

1. Desire for control: There is a strong relationship between occupants' perception

of control over their environment and productivity [49–51]. Automation systems that exclude occupants from the control-loop can infuriate the occupants [50–52]. Galasiu and Veitch [53] interpret this as the occupants' preference to have the capability to choose their environment rather than being obligated to accept the environment chosen for them. In line with this, Paciuk [54] reports that perception of control can increase comfort regardless of the physiological conditions.

2. Inappropriate vacancy detections triggering light switch off: Most automated lightswitches employ a passive-infrared and/or an ultrasonic sensor to detect occupants' presence – so that lights can be turned off in their absence [55]. Both PIR and ultrasonic sensors infer occupants' presence from their movements. When a movement is detected, the space is assumed occupied for some set delay period (typically between 5 and 60 min). An occupied instance may be incorrectly interpreted as absence and the lights can be turned off, when the occupant is immobile for more than the delay period. To eliminate occupant complaints, controls technicians or lightswitch manufacturers choose conservatively long delay periods. This causes lights to remain switched on for an extended period of time upon a departure – reducing the energy savings potential. Regardless of this conservativeness in selecting delay periods, inappropriate light switch off decisions cause frustration when all seated areas are not within the coverage of the sensors [56]. A walk-through survey conducted in 40 offices identified that about one fifth of the PIR sensors with 15 min delay were taped over to trip the occupancy-based lighting automation.

3. Desire for a customized indoor climate: Visual comfort depends on a large number of environmental and contextual factors. These factors include but not limited to the luminance of the light source, occupants' position and orientation relative to the source, background luminance, contrast in the field of view, colour of the light

source, reflectance of the furniture, individual preferences, and types of activities undertaken (drafting vs. typing) [57–59]. Despite the subtleness of the factors affecting visual comfort, it is impossible (or at least impractical) to foresee these factors when the control systems are being implemented. As a result, when daylight-integrated lighting and blinds automation systems were designed, the indoor illuminance measurements were taken by a photodiode (less commonly by a photoresistor or a phototransistor) sensor. The sensor's position also plays a crucial role on its readings. Even in a shallow perimeter office space, illuminance on the workplane can vary by a factor of ten or more [60]. For practical reasons, illuminance sensors are often positioned on the ceiling [61] or on the window frame measuring vertical illuminance on the view portion of the window surface [62]. However, research on occupants' visual comfort, and lighting and blinds use behaviour has been mostly focused on workplane conditions [22]. Given the diversity in environmental and contextual factors and the variability of sensors' position, illuminance measurements associated with occupants' visual comfort conditions vary substantially from one office to another. An observational study involving 45 office occupants revealed that preferred workplane light levels range from 91 to 770 lux [63]. Similarly in another study, researchers observed that preferred workplane light levels range from 230 and 1000 lux [53]. An investigation conducted on occupants' blinds use behaviour in 14 offices revealed that the workplane illuminance levels that trigger a blinds closing action at arrival vary between 3 and 9 klux among these offices [40]. It is challenging for controls technicians to account for this variability in selecting their lighting and blinds automation setpoints. They tend to select conservative setpoint values – e.g., setting the automated blinds to close when the illuminance on the view portion of the window exceeds 2 klux [62] or 1.8 klux on the workplane [64], setting lights to turn off above 500 lux [65]. Regardless, some occupants will dislike blinds opening when they feel it is still bright.

Some others will dislike blinds closing before it feels too bright, because they may want to preserve their view and connection to outdoors [45]. Some will dislike lights turning off automatically, when they feel that it is still dark.

1.2 Motivation and research objectives

As discussed earlier, occupants' visual and thermal comfort is governed by a diverse set of environmental and contextual factors. It is impractical (if not impossible) to account for these factors during controls implementation. The industry typically tries to maintain comfort through a few types of spatially limited sensors. Given the subtleness of the factors affecting occupants' comfort and sensor positioning, attempts to prescribe fixed operating setpoints that will work for everyone can be considered futile. The efforts to find a set of fixed setpoints that will make everyone happy merely force controls technicians and operators to be conservative – reducing the energy savings potential while still infuriating some of the occupants. Figure 1.5 illustrates two anecdotes where a controls technician made such assumptions in an academic office building. Likely to avoid complaints during the cooling season the indoor temperatures were maintained lower than 22°C, whereas during the heating season the temperatures were kept higher than 22°C. Given that occupant's clothing insulation levels tend to increase from a heating season to a cooling season [20], it is logical to expect that the indoor temperature preferences should increase from a heating season to a cooling season. Therefore, in the example shown in Figure 1.5, the seasonal changes in the indoor temperatures were against the expected seasonal variations for the comfort temperatures. Also, the VAV terminal unit serving the same four people zone was set to supply 250 L/s air for ventilation purposes including weekends (data shown in Figure 1.5 were from January 2015). Despite the

inefficiency in these temperature and ventilation scheduling choices, the work-order logs from this building indicate that too cold complaints in the cooling season and too hot complaints in the heating season were not uncommon. Similar observations were recently reported by Hilliard et al. [66]. The office building in which they have been studying advanced control strategies had no temperature setback, and the temperature setpoint was 22°C for the entire year. Perhaps even more surprisingly, the operators reported being unaware of these inefficiencies. The examples outside the published literature indicate that employing a constant 22°C in office buildings with afterhours temperature setbacks is very typical in North American facilities management practice. Therefore, we need to tailor operating setpoints in accordance with the preferences of the occupants in each controlled zone.

In a given building, occupancy patterns can be diverse since different occupants can have very different personal schedules. However, the usage patterns of individual zones can be predictable. Instead of trying to find a single operating schedule that works for all occupants, we need to tailor operating schedules in accordance with the recurring occupancy patterns in each controlled zone. In brief, if the occupancy and occupant comfort preferences in each office can be predicted at a reasonable accuracy, different operating schedules and setpoints for individual zones can be adapted – eliminating the need for wastefully conservative fixed operating schedules and setpoints.

Each thermal zone – given the uniqueness of its geometry and thermophysical properties, occupancy and occupant-driven load characteristics – has a unique temperature response subject to its HVAC equipment. If we filter out this signature temperature response for each zone, we can predict the length of the warm-up periods and help us choose better operating schedules.

Modern automation and control systems in commercial buildings provide access to

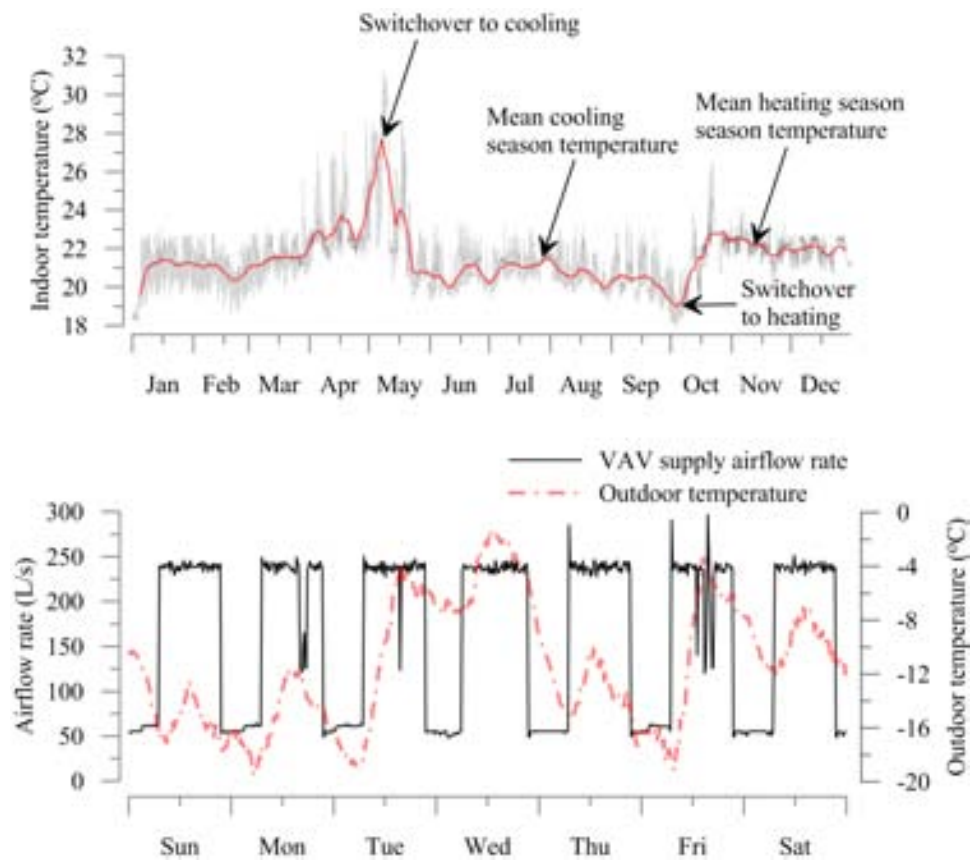


Figure 1.5: Two examples illustrating the conservativeness in technician choices in an academic office building.

real-time sensory data and distributed computational power. Most occupant actions can be registered through their interactions with control interfaces. These can include manual adjustments to the thermostat setpoints, lightswitches (if lighting is integrated to the BAS), and motorized blinds (if the blinds are part of the BAS). In addition, commercial building BASs are either already equipped with or they can be easily upgraded to have a range of sensors monitoring the indoor climate. Some of the common low-cost building sensor types include PIR motion detectors, CO₂ sensors, relative humidity sensors, photosensors, thermistors, and current sensors. Concurrent analyses of the occupants' control actions with the sensory measurements provide invaluable information about their comfort preferences. Simply put, we can learn the temperatures and light intensities disliked by the occupants in each zone through their interactions with thermostats and lightswitches. We can update operating setpoints in accordance with this information. The recurring occupancy patterns in each zone can be learned from PIR or other low-cost sensors; and this information can be employed to update operating schedules. Similarly, we can formulate inverse models by mapping the sensory information available inside a BAS onto simplified physical models. These inverse models can be employed in choosing the optimal start and stop times of the HVAC equipment.

The long term research objectives of the thesis are to improve energy efficiency and occupant comfort in office buildings through better control of the indoor climate. Within the scope of this thesis, a scalable and transferable indoor climate control tool was developed. The tool contained a set of novel algorithms that recursively learn from the occupancy patterns, adaptive occupant behaviours, and the temperature response in each control zone. The algorithms were designed to be embedded inside typical zone controllers (e.g., VAV controllers) and they were intended to undertake this learning process in real-time. The tool used the information derived from the

occupants and the building to autonomously choose better operating setpoints and schedules – tailored to exploit the nuances amongst subspaces in a building. This is also a major step to reduce the manual labour (e.g., continuous commissioning contracts with controls companies, operators) devoted to indoor climate control in commercial buildings. Due to the lack of standard controls infrastructure and quality data, residential buildings are not within the scope of this thesis.

1.3 Research methodology

The research was conducted in five stages:

1. Data collection and analyses for occupancy and occupant behaviour: The first stage in research is the statistical analyses of the human presence and behaviour patterns in office spaces using a longitudinal dataset gathered in an academic office building in Ottawa, Canada. In particular, occupants' lighting, blinds, and thermostat use patterns were studied. Different occupant model forms were introduced, and their appropriateness for controls-oriented applications was assessed. Upon this analysis, a model form for each behaviour type was selected. Algorithms for the recursive development of the selected models inside a controller were formulated.

2. Inverse modelling of the indoor temperature response: The second stage is the development of a data-driven greybox model predicting the temperature response of a thermal zone using the data gathered in the same academic office building. Different model forms at different complexities and sensory inputs were introduced, and their predictive accuracy and appropriateness for controls-orientated applications were assessed. An appropriate recursive parameter estimation methodology for the models to learn their unknown parameters in real-time was formulated. The simplest feasible model form that can mimic the physical response of the system was selected.

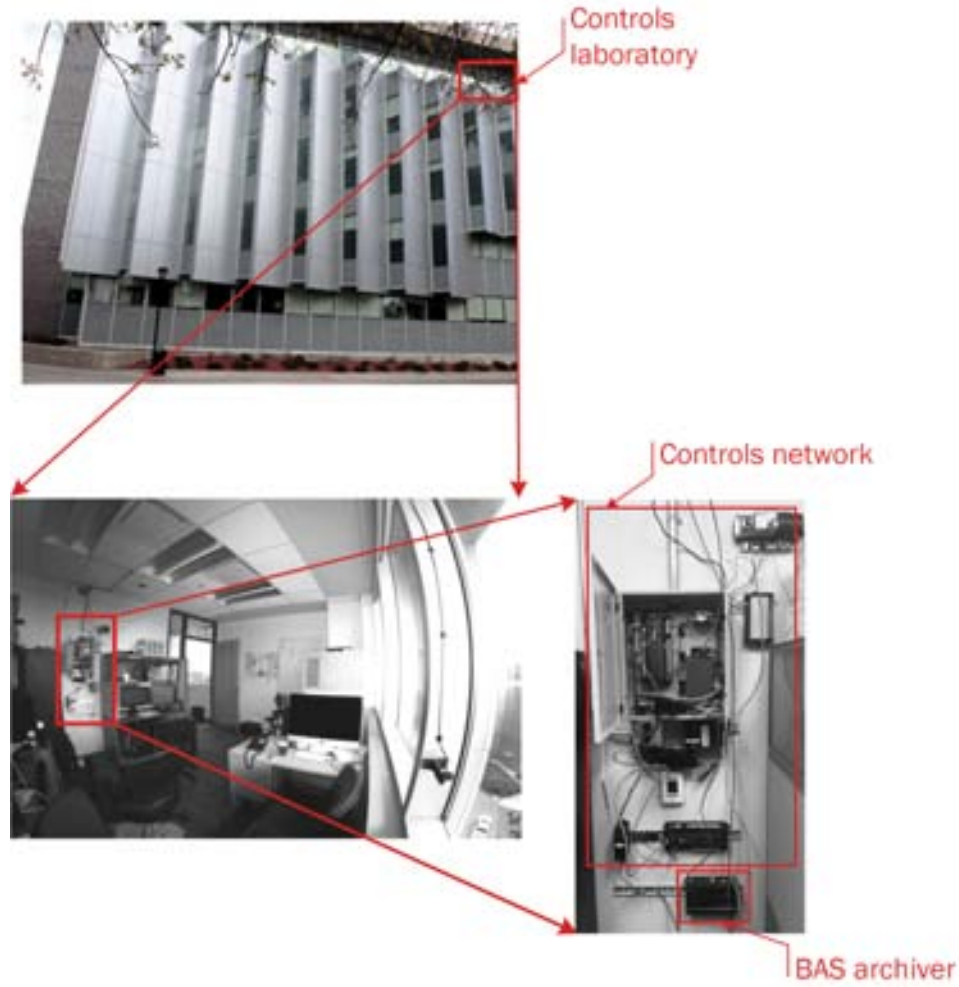


Figure 1.6: The controls laboratory.

Table 1.1 presents the characteristics of the dataset used in the first two stages (during the occupant model and data-driven greybox building model development process). In each analysis, a subset of this larger dataset was employed. In the following chapters of the thesis, the readers will be referred to this table, and further details will be provided as needed.

3. Laboratory implementation: The third stage is the laboratory implementation.

Table 1.1: An overview of the dataset used during the occupant model and data-driven greybox building model development process.

Data type	Number of offices	Acquisition period	Sensor type or monitoring method	Sampling frequency
Occupancy	16	Jan 2014 - May 2016	PIR (from BAS)	Event-based
Lighting	11	Jan 2014 - May 2016	Light switch (from BAS)	Event-based
Thermostat	11	Jan 2014 - May 2016	Keypress (from BAS)	Event-based
Thermostat	9	Jan 2016 - May 2016	Keypress (from BAS)	Event-based
Blinds	8	Feb 2014 - Oct 2014	Time-lapse photography	30 min
Blinds	1	Jan 2014 - May 2016	Keypress (from BAS)	15 min
Solar irradiance on facade	—	Jan 2014 - May 2016	Southwest and West-facing pyranometers (from local weather station)	15 min
Indoor illuminance on ceiling	11	Mar 2015 - May 2016	Photodiode (from BAS)	15 min
Indoor temperature	20	Jan 2014 - May 2016	Thermistor (from BAS)	15 min
Outdoor temperature	—	Jan 2014 - May 2016	Thermistor (from BAS)	10 min
VAV airflow rate	17	Jan 2014 - May 2016	Pressure sensor (from BAS)	10 min
VAV discharge air temperature	17	Jan 2014 - May 2016	Thermistor (from BAS)	10 min
Radiant heater state	17	Jan 2014 - May 2016	Feedback (from BAS)	10 min
Door position	1	June 2015 - May 2016	Door contact sensor (from BAS)	Event-based
CO ₂	1	June 2015 - May 2016	Non-dispersive infrared (from BAS)	1 min
Transmitted solar irradiance	1	June 2015 - May 2016	Pyranometer (from BAS)	1 min

The algorithms of the indoor climate control tool were implemented inside a controls laboratory for testing and further analysis. They learned from the recurring occupancy patterns, occupants' comfort preferences, and the temperature response of the space. The recursively derived information was used to adapt setpoints and schedules. The controls laboratory was a Southwest-facing shared-office space intermittently used by four graduate students (see Figure 1.6). It had a standalone controls network that actuates the valve of a perimeter radiant heater, the damper of a VAV unit, recessed fluorescent lighting and automated roller blinds. The laboratory was redundantly equipped with a variety of sensors (see Figure 1.7). It had a BAS archiver to permanently store the data. Through this laboratory investigation, the implementation challenges were identified, and the algorithms were verified and improved.

4. Field implementation: The fourth stage is the field implementation. Upon laboratory testing, in the same academic office building, the algorithms were implemented inside controllers serving eight private office spaces for the indoor temperature control and five private offices for the lighting control (see Figure 1.8). As shown in Figure 1.9, the offices had two VAVs with a reheat option serving them (a VAV for four offices). Each office had an individually controlled radiant perimeter heater. They had recessed fluorescent lighting integrated to the automation system. Each room was equipped with a PIR motion sensor, a ceiling-mounted photodiode sensor, a thermistor sensor, and a CO₂ sensor. Outdoor air temperature measured at the AHU intake was a common database object for the entire network. Offline plug-in equipment load measurements were also taken from these offices. Blind positions in these offices were also monitored through time-lapse photography. This part of the study rendered the potential to analyze the energy and comfort related implications




Figure 1.7: The sensors and actuators in the controls laboratory.

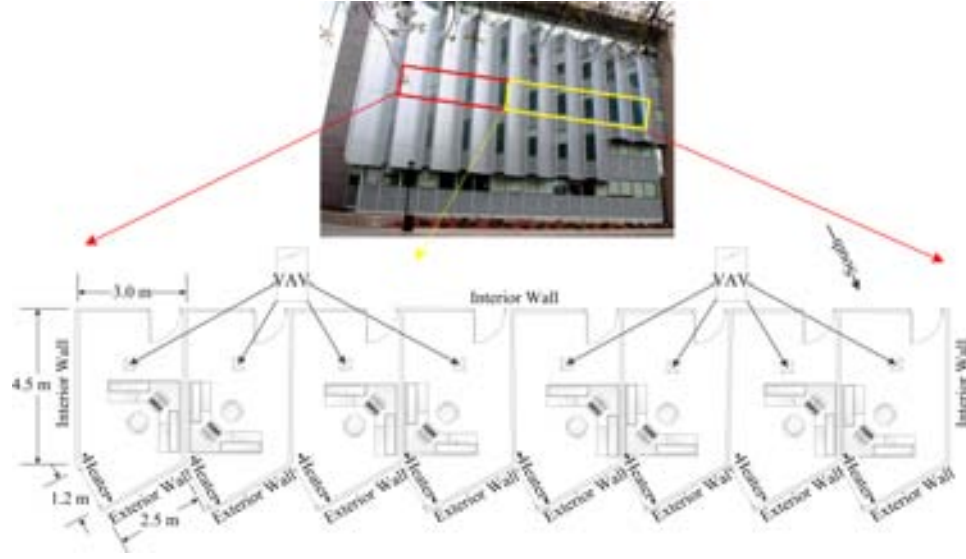


Figure 1.8: The offices employed in the field implementation. The red lines represent the eight rooms used in the implementation for the temperature controls and the yellow dashed lines represent the five rooms used in the implementation for the lighting controls.

of the adaptive control algorithms developed in this thesis. A control office group was selected from the same building a storey above the eight offices. They had identical geometry, construction, and occupancy characteristics with the exception that they operate with the existing controls programs – not with the adaptive controls algorithms developed in this thesis.

5. Simulation-based investigation: The fifth stage is the simulation-based investigation of the adaptive indoor climate control tool's functionalities. For temperature controls, a building performance simulation (BPS) model of the eight offices, which were employed in the field implementation, was built in EnergyPlus [67]. The algorithms were implemented inside the Energy Management System (EMS) application of EnergyPlus to actuate the heating and cooling equipment of the simulated offices. For lighting and blinds controls, a daylighting model of the offices, which were employed in the field implementation, was built in Daysim [68]. Virtual experiments




Figure 1.9: The sensors and actuators in offices whereby the algorithms were implemented.

were conducted under ambient conditions defined in the CWEC [69] for Ottawa. The objectives of these simulation-based investigations were (1) to estimate the energy-savings potential of the indoor climate control tool – as the resolution of utility metering in the building did not permit us to quantify this experimentally, and (2) to study the sensitivity of the indoor climate control algorithms' performance by varying the operational parameters that could be selected by different controls technicians.

1.4 Document Structure

The thesis is structured to present the data analyses to develop the occupant and the building temperature response learning algorithms first.

Chapter 2 presents the development of an adaptive presence detection algorithm from sensory data gathered in the laboratory. In addition, the analyses of the recurring occupancy patterns in an academic office building are presented. Upon these analyses, an algorithm that can recursively learn from these occupancy patterns is developed.

Chapter 3 presents the analyses of the thermostat keypress actions with concurrent occupancy and temperature data, and a recursive algorithm is formulated to develop thermostat use models inside zone controllers. In a similar fashion, the analyses of the occupants' lighting use patterns are presented, and a recursive algorithm is developed to learn from occupants' light switch behaviour inside building controllers.

Chapter 4 presents the inverse modelling of the temperature response inside individual offices. The purpose of this exercise is to filter out the warming/cooling capacity of the terminal HVAC units subject to environment and occupant-driven loads. This information is needed in calculating the near-optimal temperature setback-to-setpoint transition time in each zone prior to the occupants' arrival.

These chapters are followed by two others to introduce the use of the algorithms in temperature, and lighting and blinds control.

Chapter 5 presents an adaptive indoor temperature control algorithm. It employs the recursively derived information from the learning algorithms presented in Chapters 2 to 4. Laboratory and field implementation results and the comfort implications were presented. Energy-savings potential was further investigated through EnergyPlus simulations.

Chapter 6 presents an adaptive indoor illuminance control algorithm. It employs the recursively derived information from the learning algorithms presented in Chapters 2 and 3. The laboratory and field implementation results and the comfort implications were presented. Energy-savings potential was investigated through a daylight simulation study. Considering the breadth of the topics covered within this thesis, the relevant literature is presented in the beginning of each chapter.

Chapter 2

Learning from occupants' presence

The first section of this chapter presents a case study conducted in the controls laboratory to examine different methods for detecting human presence in offices. Upon the results of this case study, an adaptive presence detection algorithm that learns from the frequency of the occupants' detectable movements was formulated. In the second section of this chapter, the recurring occupancy patterns in 16 private offices in an academic building were analyzed. An algorithm that can recursively learn from these occupancy patterns was formulated.

The chapter covers the development of the algorithms. The laboratory and field implementation of the algorithms for verification and further testing are presented in Chapter 5.

2.1 Detecting presence

Detecting occupants' presence in office spaces is central for many lighting and heating, ventilation, and air-conditioning (HVAC) control applications – e.g., demand controlled ventilation, motion sensor-based lighting control [70]. Particularly, incorrect vacancy detections in lighting controls can lead to occupant annoyance, and in



Figure 2.1: An example where the occupant taped over the motion detectors (credit: Sara Gilani).

many cases occupants place an opaque material to cover the surface of motion detectors [43]. Figure 2.1 presents an example from a private office space in an academic building in Carleton University. In addition, there has been a surge of research on innovative lighting and HVAC controls which simply takes the detection of human presence in office spaces for granted – e.g., Roisin et al. [71]'s lighting control algorithm. The effectiveness of such control strategies is dependent on the quality of presence detections.

2.1.1 Literature review

PIR sensors are the most common sensor type used in office buildings to detect human presence. They detect movements from changes in the infrared-radiation impinging on them [72]. Given that movements are discrete-events, in practice a delay value (e.g., 15 to 60 min) is heuristically selected to avoid incorrect vacancy detections during immobility. After each movement detection, the space is assumed occupied for this delay period. The uncertainties in occupants' activeness (frequency of detectable movements), office layouts and sensor positioning play a nebulous role over the reliability of PIR sensors. These uncertainties force controls technicians and programmers to select conservatively long delay values – which diminish the energy

savings potential. Nagy et al. [73] demonstrated that optimal delay values in different space types can differ by a factor of five (4 to 20 min).

Within a building controls network, PIR sensors are often built-in inside control interfaces such as thermostats or lightswitches and mounted on vertical surfaces, or they are mounted on ceilings as standalone sensors. When they are built-in inside control interfaces, the design purpose of the control interface dictates where the PIRs are placed. For example, the lightswitches in commercial buildings are typically placed about 100 cm above the floor level. For thermostats, this value is about 150 cm. Thus, the same motion sensor placed inside a wall thermostat rather than a light switch would tend to have a lesser coverage of a seated occupant [74]. Within a wall thermostat (which contains a thermistor sensor), a PIR sensor needs to be placed away from terminal HVAC units and exterior walls.

Ultrasonic sensors – similar to PIRs – detect human movements, and thus suffer from the same types of limitations as the PIR sensors with the exception that they were reported to be more sensitive to the motion of inanimate objects (e.g., blowing curtains) [75]. An early study Maniccia [76] reviewed 23 commercially available ultrasonic and PIR sensors. Of them, 18 failed to detect movements within the manufacturer's claimed coverage range. Almost two decades after this study, there is no evidence whether this remains as a valid issue.

The building controls industry has begun to integrate CO₂ sensors inside wall thermostats. The CO₂ generation rate of a sedentary occupant is about 0.3 L/min [77]. Thus, a single occupant in a hypothetical perfectly airtight and unventilated 15 m² office space can increase the CO₂ concentration by about 350 ppm/hour. The CO₂ sensors render the potential to infer occupants' presence by filtering out occupants' influence on the CO₂ concentration of the office spaces. In the reviewed literature, this was done via simple physical models assuming perfectly mixed indoor air [78–80].

Although the early attempts to estimate the whole-building occupancy by taking CO₂ readings at the AHU's supply and return air were successful [79], estimating the spatial distribution of occupancy at the zone level were found challenging [80]. Transport of CO₂ emitted by occupants is a transient process. Depending on the office furniture layout and the distance between occupants and the CO₂ sensor, the influence of occupants on the sensor's readings will lag. Although room specific anecdotal observations were reported for the time lag values in the literature (e.g., Arora et al. [80] 30 min or Dong and Andrews [81] 20 min), the variables such as ventilation rate and style (e.g., displacement or mixing ventilation), room size, window/door positions can affect the CO₂ transport rate (by changing the process from diffusion to advection). Furthermore, estimating occupants' influence on CO₂ sensors' readings may require auxiliary information from VAV units' ventilation airflow rate and CO₂ concentration, door position and corridor CO₂ concentration, and air permeability of the envelope and outdoor CO₂ concentration. As a result, CO₂ sensors have been treated as a secondary source of information after the PIR or ultrasonic sensors [81–83]. Particularly, in shared office spaces – where PIR sensors typically do not have a direct line-of-sight to the motion for all occupants – inclusion of CO₂ sensors may improve the accuracy of presence detections [82].

Acoustics sensors have also been used to detect occupants' presence in office spaces [84, 85]. Challenges in distinguishing sounds from inanimate sources (traffic from outside), acoustic attenuation between occupants and the sensor, and lack of acoustic sensors in typical building controls networks limit the applicability of the acoustic sensors beyond research.

Advances in computer vision provide opportunities to generate high-resolution occupancy information [86–88]. Given the computational limitations to train image-processing algorithms inside BASs, two different approaches were developed: (1) a

standalone camera mounted on a low-cost microcomputer board that executes the image processing locally and provides simple occupied/unoccupied information to the BAS [87] or (2) images coming from a camera network (e.g., surveillance) were processed inside a central computer [86, 88]. Brackney [87]'s image processing occupancy sensor adopted the first approach and it can be employed within typical BASs. Currently, an image recognition-based occupancy sensor is not commercially available.

Decision trees and Bayesian statistical inferences were commonly employed to blend the information gain through various sensors [72, 80, 89]. These studies ranked the information gain through different sensor types and they typically agreed that PIR sensors provide the largest information gain followed by plug load, door contact, CO₂, and acoustics sensors.

With the innovations in advanced low-cost sensing technologies, the focus in occupancy detection research has diverted to tracking occupants' position and activities from monitoring spaces (to infer anonymous occupancy). For example, Li and Becerik-Gerber [90] and Philipose et al. [91] employed RFID tags, Milenkovic and Amft [92] used ultrasound range finders mounted on computer monitors, and Nguyen and Aiello [84] monitored occupancy through pressure sensors placed on seats. Tracking individuals – in lieu of monitoring spaces – bring about privacy concerns. And, arguably it will provide very limited benefit for the operation of existing buildings because of (1) the capital cost to enhance sensing capabilities and (2) the limited spatial resolution of the BAS-integrated building actuators (e.g., in existing buildings terminal HVAC units and lighting cannot be controlled to maintain cubicle level thermal and visual comfort).

2.1.2 Methodology

Due to ethical constraints in collecting ground-truth occupancy data through a camera, this study was conducted inside the controls laboratory only. Recall that the laboratory is a 17 m² office space intermittently used by four graduate students – whose consents were obtained. However, with occasional visitors the number reached five in a few instances.

During these experiments, the data from three PIR sensors were used. Of them, two were built-in inside commercial thermostats mounted on the walls (at 1.2 and 1.4 m heights above floor). The same wall thermostats contained CO₂ sensors as well. The voltage output of the CO₂ sensors was converted to CO₂ concentration readings between 0 and 2000 ppm using the analog input configuration details provided by the thermostat manufacturer. The calibration of the sensors was tested against a different CO₂ sensor. It is observed that the readings of the individual CO₂ sensors were biased (-40 and -80 ppm for the two CO₂ sensors). However, because the algorithm needs to make presence detections in presence of the uncertainties typical of commercial BAS instrumentation, the instrumentation was not upgraded or calibrated specifically for this study – i.e., the instrumentation was assumed as-is. Instead, the viability of filtering out occupants' presence from the response of a noisy CO₂ sensor was examined.

Figure 2.2 presents the PIR sensor locations and their coverage ranges. Note that the line-of-sight of PIR 1 was blocked by the adjacent cubicle; whereas the adjacent wall blocked PIR 2's coverage. These are typical conditions in open-space offices. The third PIR sensor was a standalone wireless sensor communicating with the network through an EnOcean-to-BACnet gateway, and it was mounted on the ceiling. The coverage cone had a radius of 1.6 m at the workplane. The PIR sensors generate an instantaneous digital signal upon movement detection. Each movement detection




Figure 2.2: Sensor locations and their coverage ranges.

was stored inside the BAS archiver. A reed switch sensor monitored the door state (open or closed). A differential pressure sensor was readily available in the network to control the VAV unit's damper position. The CO₂ concentration of the AHU's supply air was also part of the network database. A fisheye lens camera was set to take low-resolution photos at 1 min intervals. Manual interpretation of these images represents the ground-truth occupancy information. Table 2.1 provides a summary of the characteristics of the sensory data used in this part of the study.

2.1.3 Results and discussion

When a PIR sensor detects a movement, the probability of presence becomes one. As the time elapses in absence of a movement detection, the likelihood of presence decreases. Commonly, we introduce a delay period based on our prior experience

Table 2.1: Characteristics of the sensory data used in occupancy detection study.

Sensors	Properties	Sampling
PIR 1 & 2	5 m range, 100° horizontal and 80° vertical coverage	Event-based sampling
PIR 3	Detection cone radius 1.6 m for workplane	
CO ₂ 1, CO ₂ 2, CO ₂ of supply air	0-2000 ppm non-dispersive infrared CO ₂ sensor	1 min per sample
Door	Reed switch sensor	
Pressure	Differential pressure sensor	
Camera	180° vision 320 x 240 pixels	

such that if a movement is detected, the room is assumed occupied over this delay period. Figure 2.3 presents the individual movement detections over an illustrative 12 h period by the three PIR sensors. With the 10 min delay, the PIR sensors were converted to occupancy time-series. In reality, there were seven intermediate vacancy periods. Except the 3.2 h break in the morning and the 0.9 h break in the afternoon, the breaks were usually less than 10 min. Detections made using PIR 1 was only able to capture four of the intermediate breaks; and we could detect three and four intermediate breaks with the PIR facing the door (PIR 2) and the PIR on the ceiling (PIR 3), respectively. Perhaps even more importantly, the PIR 1 failed to detect the first arrival of the day. This is likely because its vision was obstructed by the edge of the cubicle near the door (see Figure 2.2). This underlines two major limitations: (1) PIR sensors alone can fail to capture short intermediate breaks, as they cannot distinguish them from mere immobility. (2) A PIR sensor should be carefully positioned so that all seated areas can be within its coverage. However, this may not be always possible in an open-space office space. In cases where a single PIR has to be used and it cannot have a direct line-of-sight to all cubicles, placing a sensor facing the entrance would at least guarantee to capture each occupants'

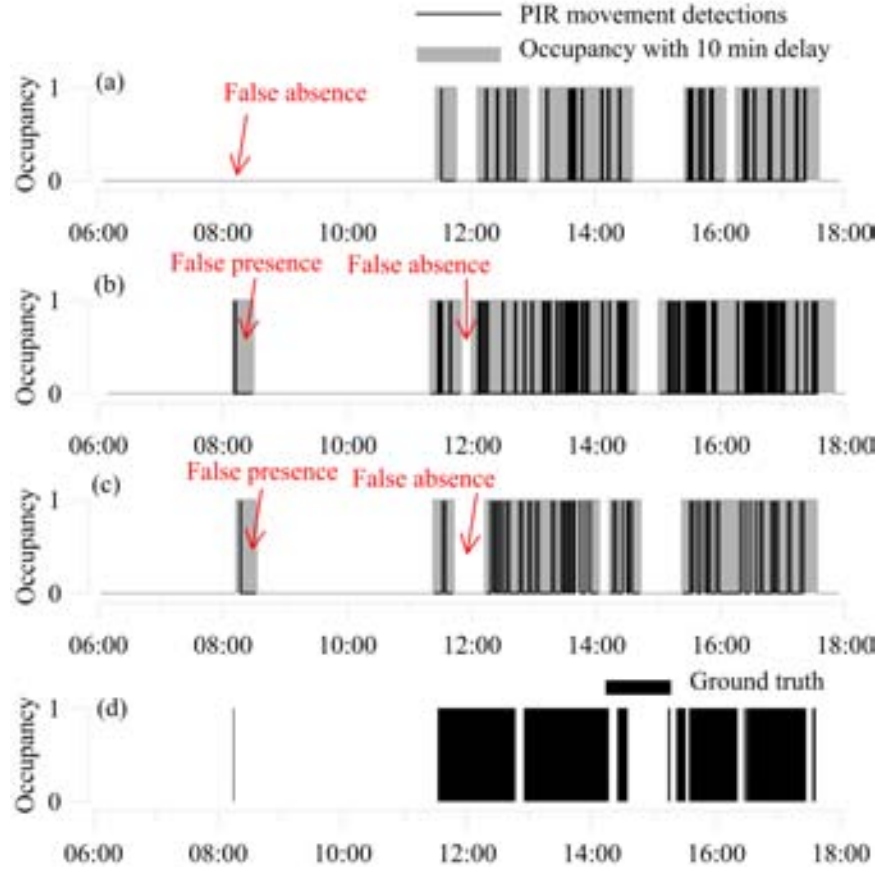


Figure 2.3: Movement detections over an illustrative 12 h period by (a) PIR 1, (b) PIR 2, (c) PIR 3; and (d) the ground-truth occupancy for the same period.

arrivals and departures. Not having PIR coverage to certain occupants would create a disproportionate annoyance – in contrast to those who were within the PIR coverage. The second most important choice that one needs to make regarding PIR sensors – after deciding where to place them – is the delay period. In Figure 2.3, we illustrated the conversion of movement detections to occupancy time-series using an arbitrary 10 min delay. For the full dataset (not the illustrative 12 h shown in Figure 2.3), Figure 2.4 presents the accuracy of presence and absence predictions at different delay values. The accuracy of presence predictions equals the ratio of correct presence detections, and the accuracy of absence predictions equals to the ratio of correct

absence detections. They were computed by contrasting the detected occupancy time-series to the ground-truth. As the length of the delay period increases, the accuracy of presence detections increases at a reducing rate. Note that high accuracy presence detections are essential for occupant comfort – e.g., ensuring that lights do not turn off when the occupants are present but immobile. In contrast, the accuracy of absence detections decreases – meaning that we start interpreting absence periods as presence periods – as the length of the delay period increases. From controls point of view, this causes a reduction in the potential energy savings. According to Nagy et al. [73], delay values for lighting controls should be selected to achieve a minimum presence detection accuracy of 95% for occupant comfort. Depending on the purpose of the occupancy detections, larger presence prediction errors can be acceptable. For example, incorrect detection of the intermediate vacancy periods, despite being a major source of occupant annoyance for lighting controls, may not cause a noticeable thermal discomfort given that buildings respond to thermal loads slowly.

The results of the different scenarios examined include:

1. Individual PIR sensors with an adaptive delay algorithm

Even when an acceptable detection inaccuracy is decided, it is challenging to choose the delay value for a PIR sensor achieving it. The PIR delay value achieving 95% accuracy in detecting presence ranged from 14 to 20 min (see Figure 2.4), while there is no guarantee that a fourth sensor will be in this range. Naturally, the ground-truth occupancy information for the calibration of PIR sensors' delay are not readily available in practice. This naturally forces controls technicians to select conservatively long delay values. Note that selecting a 60 min time delay can incorrectly interpret 22% of the unoccupied periods as occupied. This value is 14% for a 30 min time delay.




Figure 2.4: Accuracy of presence and absence detections as a function of the delay value using (a) PIR 1, (b) PIR 2, (c) PIR 3, (d) all three PIRs simultaneously.

Following the optimal time delay concept introduced by Nagy et al. [73], a practical way of estimating the optimal delay value in operation was proposed. Recall that probability of presence $p(\mathbf{p})$ becomes one as soon as a movement is detected and then it starts decreasing as the time elapses without a new movement detection – $p(\mathbf{p}|\Delta t)$. Here, we will denote the time between two consecutive movement detections as Δt . Nagy et al. [73] examined two different PIR delay values achieving 90 and 95% accuracy in predicting presence periods for lighting automation. It was observed that the former generates frequent user complaints while the latter does not. In line with Nagy et al. [73], the optimal delay value is taken as the minimum delay achieving 95% accuracy in predicting presence periods. We wish to compute the Δt value that would provide a 5% chance of presence – so that we can select it as the optimal delay value. Bayesian inference suggests $p(\mathbf{p}|\Delta t) \propto p(\Delta t|\mathbf{p}) p(\mathbf{p})$ where $p(\Delta t|\mathbf{p})$ is the empirical likelihood distribution for Δt during true occupied instances. Because occupancy cannot be truly observed, it is assumed that $p(\Delta t|\mathbf{p}) \propto p(\Delta t)$. The handicap of this assumption is that the PIR sensor must cover all occupants – or at least their arrivals and departures – in order to be effective. For example, PIR 1 does not have a clear line-of-sight to two of the cubicles and the door. As shown in Figure 2.5, the movement detections by PIR 1 tend to be separated by a small Δt (i.e., frequent movement detections). However, the distribution of Δt will chronically fail to represent true optimal time delay, because it fails to monitor two of the occupants. Although more than 95% of the movement detections by PIR 1 were separated by less than 9 min (see Figure 2.5.a), the choice of 9 min delay would result in incorrect interpretation of an occupied office as unoccupied at 15% of the occupied times (see Figure 2.4.a). The assumption $p(\Delta t|\mathbf{p}) \propto p(\Delta t)$ appears to be a reasonable approximation for PIRs 2 and 3 (see Figure 2.5.b and c). Therefore, a self-adaptive time-delay learning algorithm can be developed and implemented for PIRs

that can monitor all cubicles and/or face the door. However, learning the distribution of movement detections to estimate the optimal delay can be problematic in cases where a PIR sensor fails to monitor some of the cubicles.

2. A network of PIR sensors

When all three PIRs were used at the same time, the optimal delay became 5 min (Figure 2.4.d). This underlines the potential of using PIR sensors as a network to improve the accuracy of presence detections. Recall that the minimum delay value that can achieve 95% accuracy in predicting presence by using individual PIR sensors were between 14 and 20 min. In addition, 95% of the time intervals between movement detections were less than 5 min (Figure 2.5.d). Therefore, optimal delay estimates by learning the distribution of the time intervals between movement detections appear to be more appropriate with a PIR sensor network than individual PIR sensors. When available, more than one PIR sensor should be employed in open-space offices in occupancy detection and the optimal delay time can be approximated through learning the distribution of the time intervals between movement detections recursively.

3. A PIR and a door contact sensor

In commercial buildings, integrated access control can provide information about the door position in private or shared-office spaces. Figure 2.6 illustrates the relationship between the manually controlled door position and occupancy. When a manually controlled door is in closed or open state, occupants may or may not be inside the office. However, the door closing actions can provide useful information regarding the direction of the occupants' movement – if processed with a PIR sensor facing the door. Although a door closing action does not guarantee that the space is completely vacated, it permits us to make slightly more aggressive delay period choices upon a departure event. After a door closing action if a new movement has not been detected

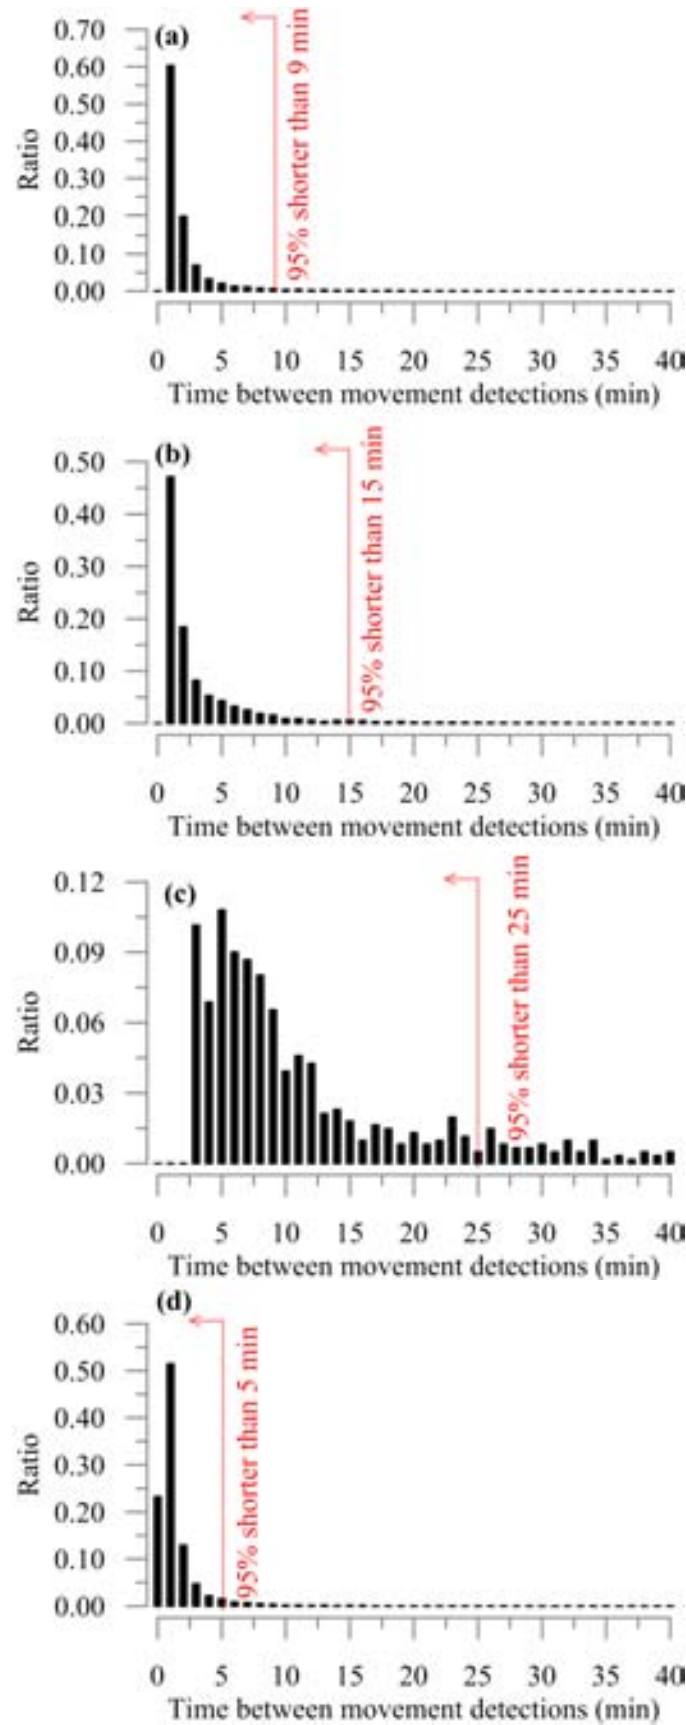


Figure 2.5: Frequency of movement detections during intermediate occupancy by using (a) PIR 1, (b) PIR 2, (c) PIR 3, and (d) all three PIRs simultaneously.

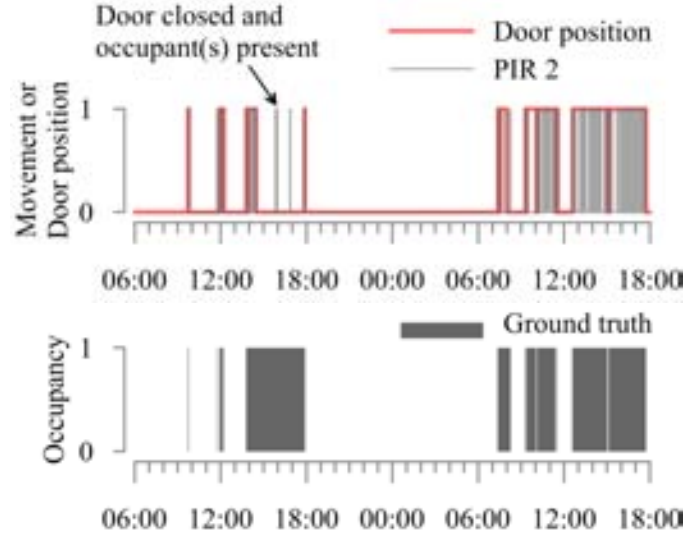


Figure 2.6: Door position, the movement detections by PIR 2, and the ground truth.

until the next time the door is opened, it is assumed that the space was vacated with the last door closing action. The ratio of these events to the total number of detected departures can be learned in recursion inside a building controller. For the data gathered during these experiments, more than 90% of the door closing actions resulted in an unoccupied office. In other words, for this office there is less than 10% chance of presence after a door closing event. In addition, more than 50% of the movements detected by PIR 2 were separated with a time interval less than 3 min. Therefore, 3 minutes after a door closing event, if a movement has not been detected, the space was most likely ($\sim 95\%$) vacated. Figure 2.7 illustrates the influence of including the information from a door sensor on the accuracy of PIR 2. Normally, the PIR 2 can still be used with a 15 min delay as estimated in Figure 2.5.b. However, after a door closing action if the PIR 2 does not detect any movements for 3 min, it is assumed that the office is unoccupied. This assumption did not affect the accuracy of the presence detections – i.e., it remains about 95% for a 15 min delay during intermediate occupancy. However, this resulted in increasing the accuracy of absence detections

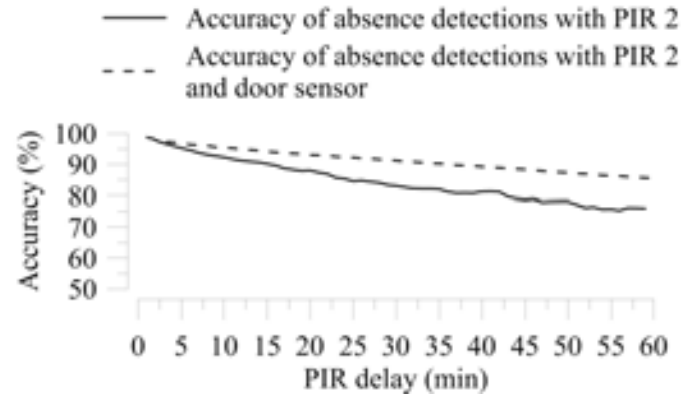


Figure 2.7: Effect of door sensor on the accuracy of presence predictions at different PIR delay values.

by 4%. In other words, we were able to truly interpret 96% of the unoccupied periods in lieu of 92% with the addition of a door sensor. This is equivalent to being able to employ occupancy-based control strategies for an extra 240 h/year.

4. A PIR and a CO₂ sensor

In general, commercial buildings are slightly pressurized with ventilation air, so that part of the room air is forced into the corridor, the return air grille, and exfiltrate into the outdoor air. Depending on the magnitude and orientation of wind and door position, the pressurization can reverse and the corridor and outdoor air can infiltrate into the room air. Due to differences in CO₂ concentration of ventilation, ambient, and corridor air, the room CO₂ concentration varies in time. Figure 2.8 schematically illustrates these processes in a perimeter office space. Here, the goal is to filter out occupants' impact on CO₂ concentration from the CO₂ response of the space. This information is used to complement a PIR sensor. Figure 2.9 presents the occupancy, door position, indoor and supply air CO₂ concentrations over a five day period. During this period, the VAV unit's supply air pressure was 10 Pa. Using the sensory data, a data-driven dynamic model that can characterize the CO₂ concentration response




Figure 2.8: Mass transfer processes affecting the CO₂ concentration in a perimeter office space.

of the office was formulated as follows:

$$\begin{aligned} \frac{dC_{in}}{dt} &= (C_{sa} - C_{in}) \sqrt{p_{sa}} x_1 + (C_{out} - C_{in}) x_2 + N_{occ} x_3 \\ C_m &= C_{in} + x_4 \end{aligned} \quad (2.1)$$

where C_{in} (ppm) is the indoor CO₂ concentration, C_{sa} (ppm) is the concentration of the supply air, p_{sa} (Pa) is the pressure sensor reading across the VAV unit, C_{out} (ppm) is the outdoor CO₂ concentration taken as a constant 400 ppm, N_{occ} is the number of occupants, and C_m (ppm) is the measured CO₂ concentration from sensors CO₂ 1 or CO₂ 2. The non-physical parameters x_1 to x_3 are weighing factors mapping model inputs onto measured output. The first component $(C_{sa} - C_{in}) \sqrt{p_{sa}} x_1$ acts as a proxy for the role of ventilation air on the rate of change in the indoor CO₂ concentration. Note that pressure sensors in VAV units are used to control airflow rate (i.e., square root of pressure drop is proportional to the airflow rate). The second component $(C_{out} - C_{in}) x_2$ represents the role of outdoor air infiltration on the rate of change in the indoor CO₂. The third component $N_{occ} x_3$ represents the role of the occupants on the rate of change in the indoor CO₂. This implicitly takes into account

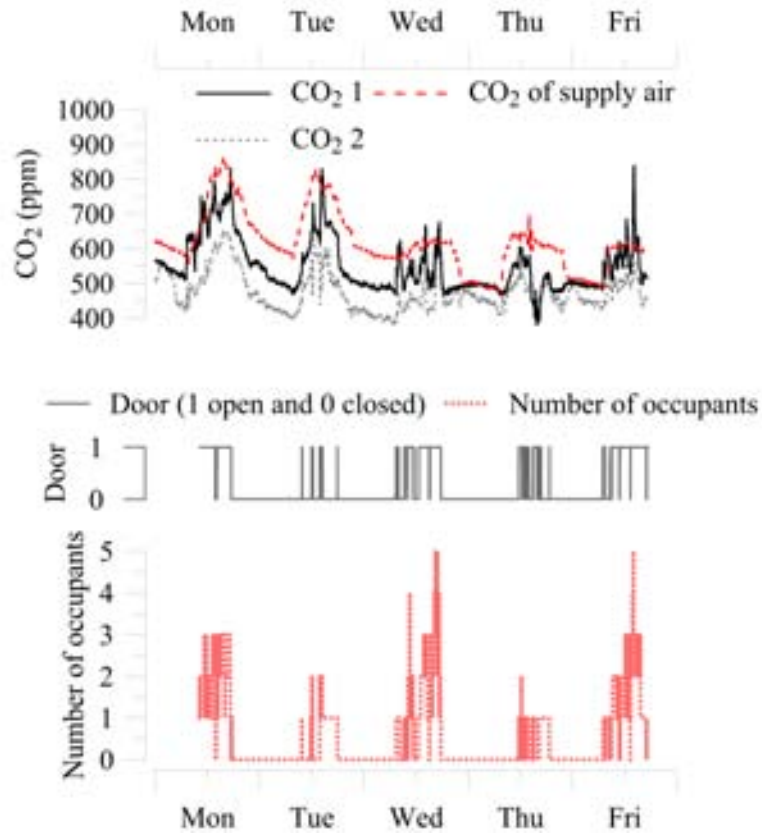


Figure 2.9: Indoor and supply air CO₂ concentrations, door position, and number of occupants on weekdays.

occupants' habit to keep the door open. The door position was originally included as part of the model; then it was left out as it provided only a marginal improvement in the predictive accuracy and to ensure the transferability of the method to other offices.

Given CO₂ emissions from the occupants, it is logical to expect that the indoor CO₂ concentration becomes higher than the CO₂ concentration of the supply air. However, readings of CO₂ 2 were always lower than the CO₂ of the supply air (see Figure 2.9). This can be explained by the aforementioned sensor bias.

Table 2.2: The parameters of the model shown in Eqn. 2.1. The parameters are shown in mean \pm standard deviation format.

Sensor	\mathbf{x}_1 ($1/h\sqrt{Pa}$)	\mathbf{x}_2 ($1/h$)	\mathbf{x}_3 ($ppm/h - person$)	\mathbf{x}_4 (ppm)
CO ₂ 1	3.05 \pm 0.04	0.38 \pm 0.09	35.7 \pm 3.7	-36.0 \pm 4.0
CO ₂ 2	1.11 \pm 0.01	0.30 \pm 0.03	24.6 \pm 0.8	-86.0 \pm 0.5

To account for this, an unknown parameter representing the sensor bias x_4 is introduced. The model parameters $x_{1 to 4}$ were estimated recursively using a particle filter (see Table 2.2). The model made predictions at one minute timesteps using a prior estimate of the parameters. The prior estimates of the parameters were rough guesses to initialize the filter. At the end of each minute, a new set of measurements were made available to the filter. Each time a new measurement was taken, the measured CO₂ concentration was compared with the predictions made by using the prior parameter estimates. Based on this comparison, the parameter estimates were updated. This procedure was repeated iteratively. Given the particle filter is recursive and the model shown in Eqn. 2.1 is simple, one can easily embed this as a self-calibrating model inside an existing building controller. Further information on recursive parameter estimation and filtering is provided in Chapter 4. The model together with these parameter estimates were employed to make predictions over a 5 day prediction horizon. Results shown in Figure 2.10 indicate that over this time horizon the model was able to predict the CO₂ response with a mean absolute error of ~ 25 ppm with CO₂ 1 and ~ 40 ppm with CO₂ 2.

The PIR and CO₂-based occupancy detection method was formulated as follows:

- (1) if a movement is detected, the room is assumed occupied over a delay period,
- (2) if a movement was detected in the previous delay period but a new movement was

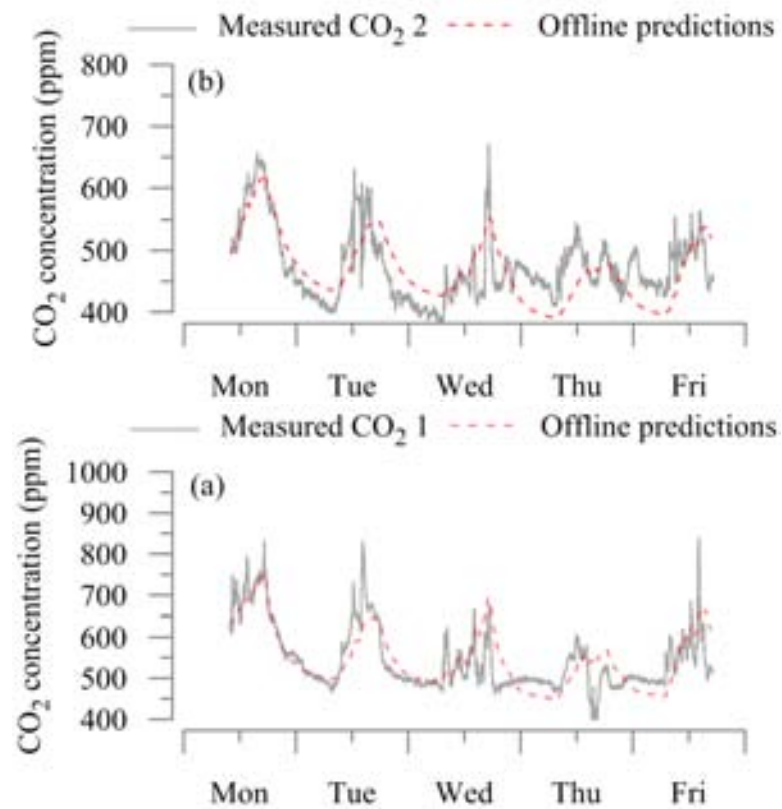


Figure 2.10: The predictions made by the model shown in 2.1 using the parameter estimates shown in Table 2.2 for (a) CO₂ 1 and (b) CO₂ 2 sensors.

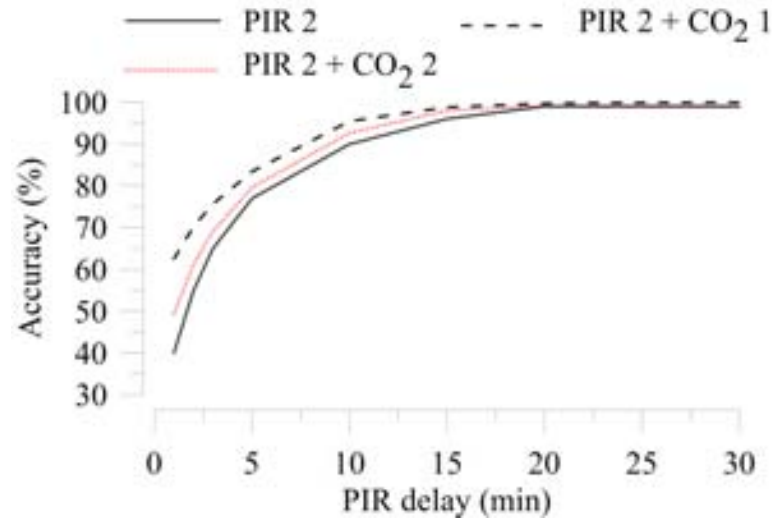


Figure 2.11: The influence of using CO₂-based occupancy detections together with PIR 2 over the absence detection accuracy.

not detected, call the CO₂ model (see Eqn. 2.1) to determine whether the space is still occupied. The results shown in Figure 2.11 indicate that the use of the CO₂ sensors slightly improved the accuracy of presence detections particularly when the PIR delay period was less than 15 min. Together with the interior CO₂ sensor (CO₂ 1), the PIR 2 delay time achieving 95% accuracy in detecting presence reduced from 14 to 10 min. This reduction was only 2 min (from 14 to 12 min) with the use of CO₂ sensor facing the door. These reductions in PIR delay times account for some minor improvements in extending the time periods correctly interpreted as absence. In addition, they render slightly better representation of the brief intermediate vacancy periods. Note that CO₂ may provide information regarding the number of occupants, however counting the number of occupants was not within the scope of this thesis.

In brief, presence detection accuracy of a PIR sensor can be improved marginally with the inclusion of door, CO₂ or other PIR sensors. However, addition of new sensors would limit the transferability of the algorithm to existing buildings. Therefore, using individual PIR sensors with an adaptive delay algorithm was found acceptable

and it was used in the rest of this thesis. A recursive controls algorithm that learns from the frequency of movement detections of a PIR sensor to adapt its delay value autonomously was formulated as a standard BACnet program object – in the scripting language GCL+ [93]. The code was implemented and verified inside the laboratory's controls network. Based on the comparison between the ground-truth and the three PIR data records, a false presence detection due to an inanimate movement (e.g., roller blinds moving due to supply air from a nearby duct) was not observed. However, false presence detections may become a concern in office spaces with operable windows. Transferability of the methods used in this section to offices with operable windows should be studied.

2.2 Predicting presence

In recognition of the transient nature of the heat transfer within building fabric, real-time occupancy detections are not adequate to control heating and cooling systems. In order to adapt the operation of heating and cooling systems, the likelihood of occupancy needs to be predicted well before the space is occupied.

2.2.1 Literature review

The challenge in predicting the recurring occupancy patterns in offices is in representing the randomness inherent in their daily schedules in a statistical model form. In the context of this thesis, the occupancy model form should help us retrieve useful information to improve operation.

Randomness in occupancy patterns can be classified in two groups: (1) day-to-day variations of an occupant's schedule and (2) the diversity amongst different occupants' schedules inside a building. An occupant's schedule, despite changing from time

to time, is individual but not arbitrary suggesting that occupants consciously and consistently choose their daily schedules [34]. On the other hand, different occupants' schedules in the same building can be vastly different [94]. As reported by Bordass et al. [95], many office occupants today tend to have flexible work hours. For example, 27% of the workforce in the United States were reported having flexible work hours in 2004 – up from 15% in 1995 [10, 11, 96]. Only two in five of the workforce in Canada has a standard work schedule [12]. As a result, personal preferences and job restrictions translate into office occupancy patterns. Some occupants prefer to arrive as early as 6 am and others as late as 12 pm [97]. In some cases, some occupants may even extend their work hours over weekends and holidays [15]. As a result, in office buildings the occupied hours can extend from 6 am to after 8 pm on weekdays, despite the fact that individuals' work hours are only about half this time [98]. Moreover, peak occupancy levels rarely exceed 50% [13, 14, 99, 100]. For example, Mahdavi et al. [13] monitored 48 offices in different types of buildings: a university building, a large office complex, and a government building. In all office types, it was found that the workstations were unoccupied at least half of the time, and the occupancy differed significantly from one office space to another. Therefore, we need to model occupancy at the individual thermal zone level in order to exploit the diversity amongst different occupants' schedules in the operating schedules of the HVAC equipment.

In the past two decades, researchers have instilled the basics of modelling human presence in office buildings [101]. In presence modelling, three different methods have been used. The most common method is building weekly occupancy schedules – presenting the likelihood of presence as a function of the time of day and the day of week [94, 102]. This model form is easy to interpret by building operators and controls technicians. Building specific occupancy schedules provide valuable insights that can help operators choose operating schedules.

The second method used in occupancy modelling is the Markov chains [103, 104]. The model predicts the likelihood of an arrival when occupants are absent and it predicts the likelihood of a departure when occupants are present. The strength of this approach is that, unlike the traditional schedule-based models, the likelihood of observing an arrival or departure for the rest of the day can be estimated – given current time and current state of presence. This may help making midday control decisions such as temperature setbacks when the likelihood of observing an arrival is very small for the rest of day [34]. A weakness of the Markov occupancy models is that they treat arrival and departure events independently. In reality, occupants may depart early when they arrive early or they may depart late when they arrive late [104]. Survival models (the third method) appear to be a promising alternative to tackle this limitation [105]. Survival models can predict the duration of an intermediate vacancy period following a departure or they can predict the duration of an intermediate occupancy period upon an arrival [106].

The research on modelling occupancy in buildings has been limited to realistically representing human presence inside BPS tools (e.g., EnergyPlus). The ability of different occupancy models in providing information to improve the control of HVAC equipment has not been actively explored. Only in a few cases, the occupancy models derived from a specific building's dataset were used to employ temperature setback in vacant or intermittently used thermal zones [70, 98, 107, 108]. Because these occupancy models have been derived from a building's dataset retrieved at a certain point during operation, they may become unrepresentative if the occupancy characteristics change during the operation. In fact, the evidence suggests that occupancy data in office buildings is non-stationary meaning that its characteristics tend to change slowly in time [104] and space [100]. In addition, the development of a different offline

occupancy model for each thermal zone in a large office building is a cumbersome and intrinsically labour-intensive process.

2.2.2 Model forms

Two years' worth of motion sensor data gathered in 16 private offices in an academic office building were analyzed to formulate a recursive algorithm that predicts the presence of occupants in a zone. The motion sensor data were converted to occupancy time-series data using the adaptive delay algorithm for PIR sensors. In the monitored offices, the office furniture did not obstruct the PIR motion sensors. All of the offices belong to a full-time faculty member. The occupants were typically engaged in sedentary office activities – typing, reading, and writing. The occupants had flexible work schedules interrupted with irregular absences as they were also engaged in work outside their offices (e.g., meetings, conferences, teaching, and work from home).

Figure 2.12 presents the mean weekly occupancy schedule in the monitored offices over the two year period. This type of analysis could reveal that the building occupancy was low – yet it occasionally extended over the weekends. The weekday occupancy spreads from 7 am to 8 pm, and the total number of occupants present in these 16 offices rarely exceeds 5 at any given moment. We can use this model form to identify the earliest expected arrival time and the latest expected departure time on different weekdays in each thermal zone. Evidently, these parameters can be used in choosing the temperature setback periods for each thermal zone.

Figure 2.13 presents a discrete-time Markov occupancy model predicting the likelihood of observing the first arrival of a weekday in the next one hour at different instances. This analysis reveals that the first arrival times in individual zones can be a multimodal distribution. Some of the occupants – even if they did not arrive in the




Figure 2.12: Weekly occupancy schedules in the studied offices.

morning – can start using their offices in the afternoon. However, the relative likelihood of a first arrival in the morning is about four times larger than a first arrival in the afternoon. From this model form, we can retrieve two important pieces of information to improve the operation: the earliest and the latest expected arrival times in each thermal zone. The latest expected arrival time can be used in executing midday temperature setback decisions. For example, if the probability of observing an arrival for rest of the day is negligible and the occupant has not arrived yet, the temperature setback can be reinstated. Note that the areas under the models do not equal one. This is because there are absent weekdays (e.g., work from home days, travels for work and leisure). Therefore, the model provides another useful information that we can use to improve the operation: the ratio of absent weekdays in an office.

Figure 2.14 presents a discrete-time Markov occupancy model predicting the likelihood of observing the last departure of a weekday in the next one hour at different instances. Similar to the distribution of the first arrival times, the last arrival times in

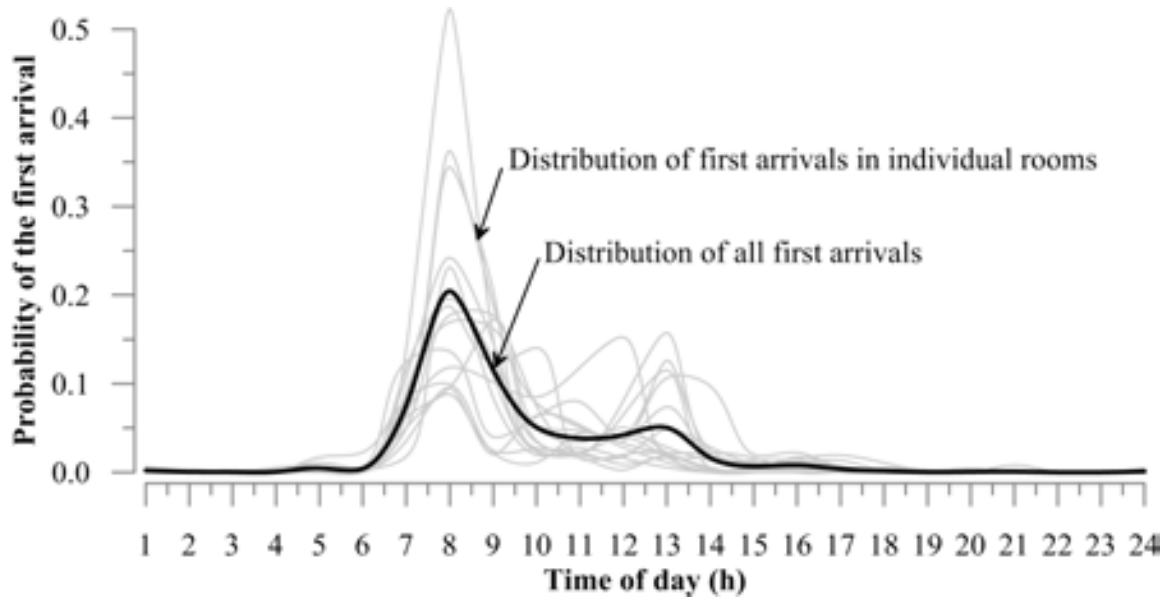


Figure 2.13: Discrete-time Markov model predicting the likelihood of a first arrival on a weekday in the next one hour.

individual zones can be multimodal distributions. An occupant either leaves his/her office early (e.g., before taking a lunch break) or the office remains occupied for a long period of time. Although the definition of the terms early and long period of time vary from individual to individual, this trend was observed in most of the monitored offices. However, the peak likelihood of the departures before noon was significantly less pronounced than the peak likelihood of the departures after 5 pm. This model form can help us retrieve another important piece of information to improve the operation: the latest expected time of departure.

Figure 2.15 presents a survival occupancy model predicting the duration of an uninterrupted intermediate occupancy period. The model form provides an information that can potentially help us extend the temperature setback periods more confidently: the longest expected duration of occupancy. Similarly, Figure 2.16 presents a survival occupancy model predicting the duration of an intermediate vacancy period. From

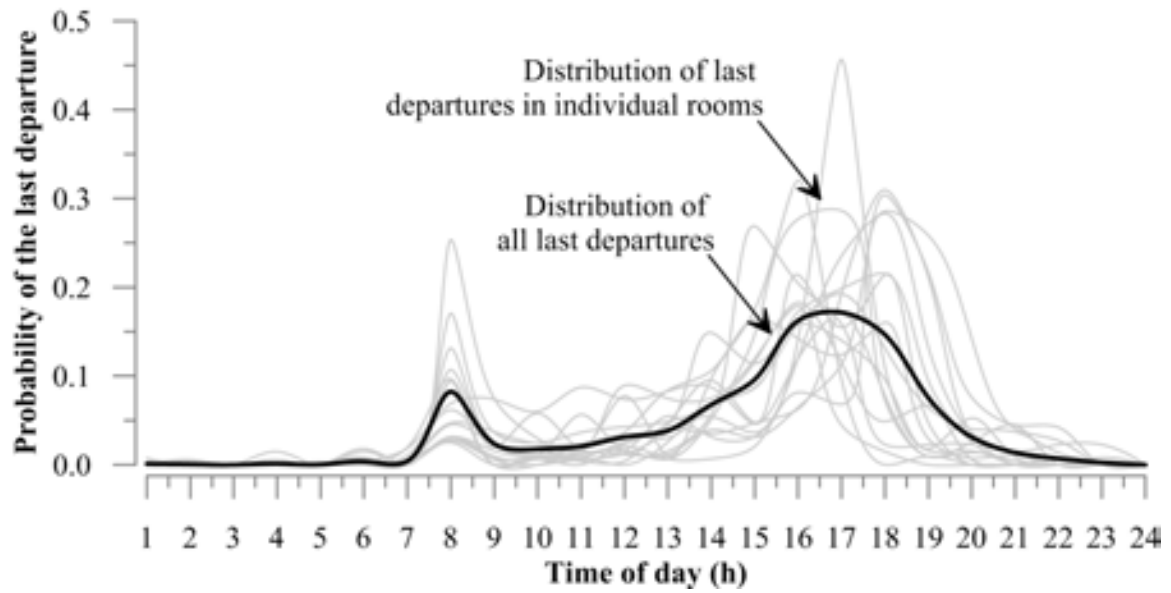


Figure 2.14: Discrete-time Markov model predicting the likelihood of the last departures in a weekday.

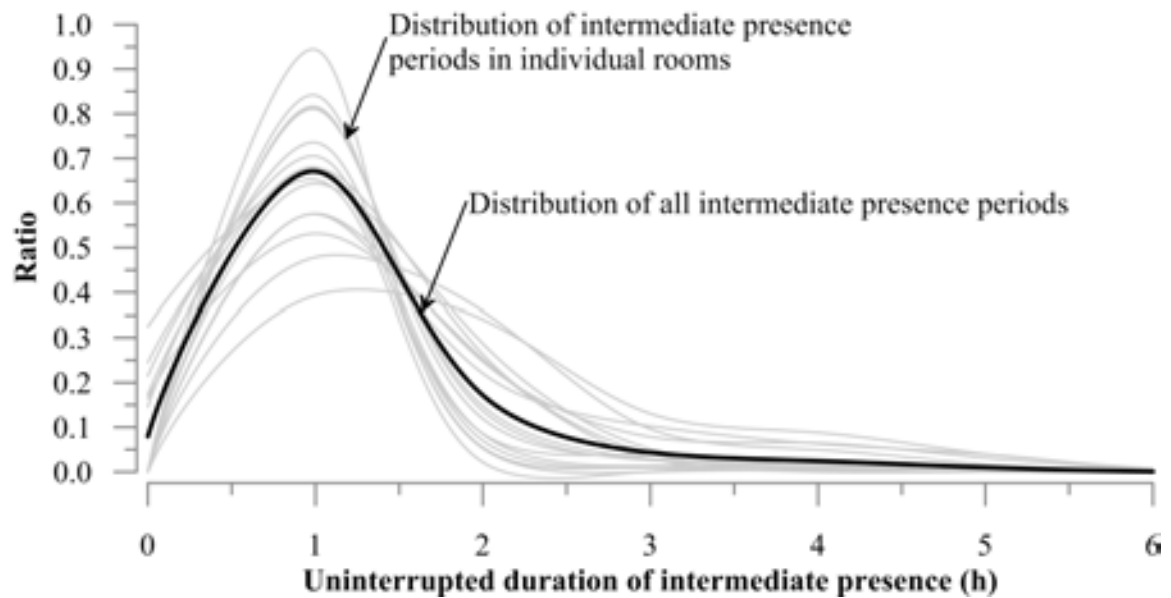


Figure 2.15: Survival model predicting the duration of an uninterrupted intermediate occupancy period.




Figure 2.16: Survival model predicting the duration of an intermediate vacancy period.

this model, we can retrieve another useful information to improve the operation: the longest expected duration of intermediate absence. This can be used to execute mid-day temperature setback decisions as well. For example, if an occupant arrived in the morning and then left the office longer than the longest expected duration of intermediate absence, the temperature setback can be restored – without waiting for the latest expected time of departure in the zone.

2.2.3 Diversity aspects

The occupancy schedules for individual offices are shown in Figure 2.17. When occupancy profiles were decomposed into individual rooms, it became clearer that the occupancy periods were very diverse. In rooms used by different occupants, the peak weekday occupancy ranged from 75% to 10%. In only one of the 16 offices, the

occupancy regularly extended over the weekend. In two of the 16 offices, the occupancy profiles appeared to be somewhat different on different weekdays. In general, occupancy schedules were similar on different weekdays.

This was further investigated by building autocorrelation plots (see Figure 2.18). An autocorrelation plot presents how well a data record correlates to itself, if the data record is shifted by a time lag. After shifting the data record by a time lag, if the data record still resembles the original data record, it would yield non-zero values [109]. The time lag values which yield a local maximum implies that the occupancy patterns tend to recur at that time interval. For example, the autocorrelation plots shown in Figure 2.18 underline the 24 h periodicity in all 16 rooms – i.e., when the data records are shifted by 24 hours, it still resembles the original data record. This is clearly not too surprising because occupants typically have daily schedules. But, the question that we wish to answer by looking at these autocorrelation plots is that: *Would shifting the individual data records by 168 hours in lieu of 24 hours result in better correlation to the original data set?* This result was crucial because based on this, we could identify whether or not the occupancy-learning algorithm should learn the recurring occupancy patterns for each weekday separately.

The results shown in Figure 2.18 indicate that in two of the 16 rooms the periodicity was significantly more pronounced at 168 hours than it was at 24 hours. This implies that occupancy on a given day of the week is a better indicator than occupancy from all weekdays. This might be expected for the subject offices, which are occupied by professors with weekly teaching schedules. It is also important to note that an algorithm to learn the occupancy patterns on each day separately would require a much longer training period than an algorithm to learn the occupancy patterns on




Figure 2.17: Occupancy schedules in individual offices. Each plot represents one of the 16 occupants. The week starts on Monday.




Figure 2.18: Autocorrelation of the occupancy time-series in individual offices.

weekdays in general – due to the reduction of the observation frequency (e.g., the algorithm would need to wait till the next Monday to update the parameters describing the occupancy patterns on Mondays).

Six occupancy parameters were introduced based on the information that one can retrieve from the occupancy model forms: (1) the earliest expected time for the first arrivals ($t_{arr,early}$), (2) the latest expected time for the first arrivals ($t_{arr,late}$), (3) the latest expected time for the last departures ($t_{dpt,early}$), (4) the ratio of absent weekdays (p_{abs}), (5) the longest expected duration of intermediate presence (D_{int}), and (6) the longest expected duration of intermediate absence (D_{abs}). In this study, the terms earliest/latest expected time or the longest expected duration stand for an occupancy event with less than 10% chance of occurrence. For example, the earliest expected arrival time means that only 10% of the first arrivals occur before that time of the day. These parameters were tabulated in Table 2.3 for each of the 16 occupants. For about half of the occupants, 7 am is the earliest arrival time, and it is 8 am for rest of

Table 2.3: The occupancy parameters in the monitored offices.

Occupant	$t_{arr,early}$ (h)	$t_{arr,late}$ (h)	$t_{dpt,late}$ (h)	p_{abs} (h)	D_{abs} (h)	D_{int} (h)
1	8	12	18	0.41	2	3
2	8	9	19	0.28	3	4
3	7	9	19	0.11	2	5
4	8	13	18	0.64	4	2
5	7	10	18	0.38	3	3
6	8	13	20	0.24	3	3
7	8	13	18	0.30	2	4
8	8	14	19	0.44	2	1
9	8	10	18	0.53	4	3
10	8	13	20	0.38	3	1
11	7	13	17	0.41	3	3
12	7	12	16	0.64	4	3
13	8	12	18	0.29	2	3
14	7	11	18	0.37	2	4
15	7	10	19	0.35	2	4
16	8	11	20	0.19	2	3

the occupants. In contrast to this limited individual variability in $t_{arr,early}$, the latest expected arrival time for different individuals vary between 9 am and 2 pm. Similarly, the latest expected departure times for different individuals range between 4 pm and 8 pm. The ratio of absent weekdays varies by a factor of six amongst monitored occupants. The longest expected intermediate break and presence durations also vary substantially from one occupant to another.

The monitored offices have individually controlled perimeter heaters. However, four private offices share a VAV unit for cooling. Therefore, each office – despite being an individual thermal zone for heating – was only a portion of a thermal zone for cooling. The uniqueness of the six parameters reduce when the terminal units serve more than one person. For example, the earliest expected time of the first arrival can be determined by the occupant that has the earliest arrival habits amongst people

using the thermal zone. Similarly, the longest expected intermediate break can be determined by the occupant that has the tendency to take the longest intermediate breaks amongst occupants using the thermal zone. The effect of diversity was studied through 16 different hypothetical zone sizes (from 1 occupant zones to 16 occupant zones). The occupants in each hypothetical zone were randomly sampled amongst 16 monitored offices. As the number of people in the zone increases, the likelihood that the zone is used earlier in the day and later in the evening increases. The latest expected time of first arrival converges to that of the occupant with the latest arrival habits (see Figure 2.19). The likelihood of observing an absent weekday reduces, and the chances of one of the occupants arrive after a long break increase (see Figure 2.20). Therefore, zone level occupancy learning can be most effectively executed in buildings with high granularity in terminal HVAC equipment. However, it is worth noting that as the zones get larger the HVAC equipment cannot account for the differences envelope and casual gains, and this can lead to discomfort problems.

2.2.4 Occupancy-learning algorithm

An algorithm that iteratively learns the recurring occupancy patterns inside an office from a PIR sensor's movement detections was developed. The insights acquired by analyzing the observational data were synthesized to shape the characteristics of this algorithm. The algorithm is intended to retrieve the six parameters $t_{\text{arr,early}}$, $t_{\text{arr,late}}$, $t_{\text{dpt,late}}$, p_{abs} , D_{abs} , D_{int} parsimoniously. It is intended to be embedded inside local building controllers. The recursively learned parameters are employed by the temperature control algorithm (see Chapter 5) in selecting the temperature setback schedules.

The first movement detection of the day is an arrival observation. Upon this

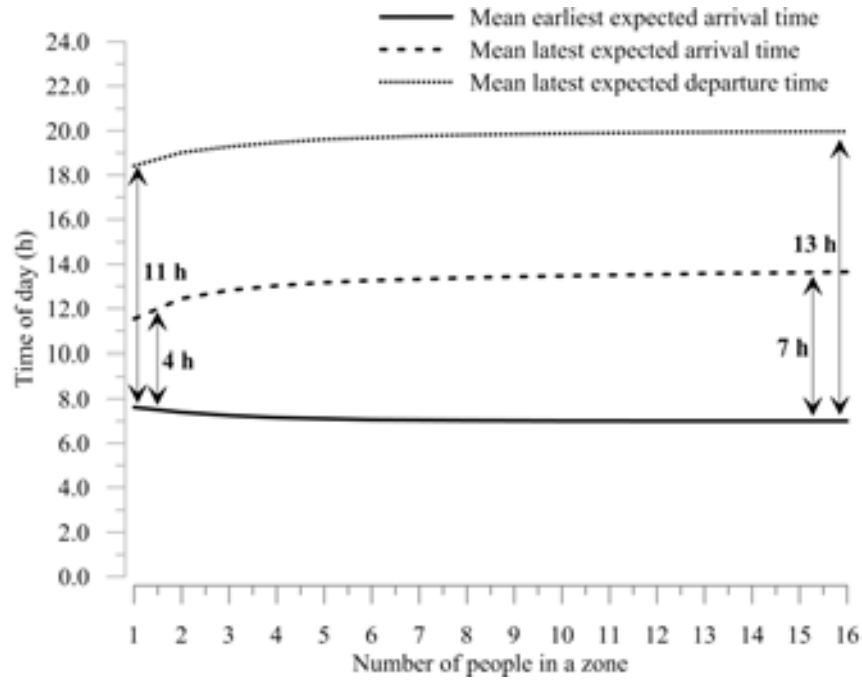


Figure 2.19: Effect of zone granularity on the occupancy parameters $t_{\text{arr,early}}$, $t_{\text{arr,late}}$, and $t_{\text{dpt,late}}$.

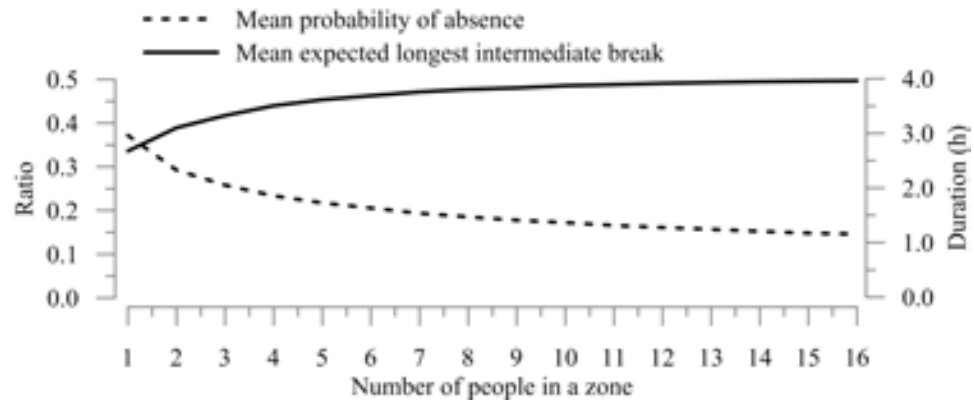


Figure 2.20: Effect of zone granularity on the occupancy parameters p_{abs} (h) and D_{abs} (h).

observation, the mean and the variance of the first arrivals are updated as follows:

$$\begin{aligned}\mu_{arr,i+1} &= \mu_{arr,i} \frac{i-1}{i} + (t_{arr} - \mu_{arr,i}) \frac{1}{i} \\ \sigma_{arr,i+1}^2 &= \sigma_{arr,i}^2 \frac{i-1}{i} + (t_{arr} - \mu_{arr,i+1})^2 \frac{1}{i}\end{aligned}\tag{2.2}$$

where μ_{arr} (h) and σ_{arr} (h) are the mean and the standard deviation of the first arrival times, respectively; t_{arr} (h) is the time of the day when the new first arrival is observed, and i is the index number of the new arrival observation.

The last movement detection of the day is a departure observation. In a similar fashion, after the last departure observation of a weekday, the mean and the variance of the last departures are updated as follows:

$$\begin{aligned}\mu_{dpt,i+1} &= \mu_{dpt,i} \frac{i-1}{i} + (t_{dpt} - \mu_{dpt,i}) \frac{1}{i} \\ \sigma_{dpt,i+1}^2 &= \sigma_{dpt,i}^2 \frac{i-1}{i} + (t_{dpt} - \mu_{dpt,i+1})^2 \frac{1}{i}\end{aligned}\tag{2.3}$$

where μ_{dpt} (h) and σ_{dpt} (h) are the mean and the standard deviation of the last departure times, respectively; and t_{dpt} (h) is the time of the day when the new last departure is observed.

If no motion sensor detections are observed on a weekday, this is an event of absence. The p_{abs} is simply $(1 - i) / j$ where j is the index number for the observed business days.

For controls-oriented purposes, the parameters $t_{arr,early}$, $t_{arr,late}$ and $t_{dpt,late}$ were approximated as follows:

$$\begin{aligned}t_{arr,early} &= \mu_{arr,i} - 1.3\sigma_{arr,i}(1 - p_{abs}) \\ t_{arr,late} &= \mu_{arr,i} + 1.3\sigma_{arr,i}(1 - p_{abs}) \\ t_{dpt,late} &= \mu_{dpt,i} + 1.3\sigma_{dpt,i}\end{aligned}\tag{2.4}$$

The multiplier 1.3 is intended to compute the 90th percentile for the three parameters. The underlying assumption in this computation is that the first arrival and the last departure times can be approximated as Gaussian distributions. As illustrated in Figures 2.13 and 2.14, this is not always true. The first arrival and the last departure distributions are bimodal distributions – yet for most occupants the major mode appears to be much more pronounced than the secondary mode. Arguably, the bimodality in these faculty offices can be explained as the first arrivals after the lunch break and the last departures prior to taking a lunch break. Alternatively, discrete likelihoods could be computed for observing an arrival or a departure at a certain time during the day. If the discrete likelihoods were set at a high temporal resolution, the characteristics of the arrival and departure patterns would be better preserved, albeit at an increased computational burden. As reported in the GCL+ programming guide, a single BACnet program object inside a local controller is permitted to only 5500 bytes, and a variable array of 100 elements uses about 4000 bytes of memory. The next generation of controllers might overcome these limitations, but these limitations had to be acknowledged to ensure the methodologies developed in this thesis are useful for the existing buildings. An advantage of approximating the first arrival and the last departure distributions as Gaussian distributions is regularizing the likelihood weights as continuous distributions.

If an occupant arrives in the morning and then departs, this value is temporarily held as a candidate for the last departure time of the day until midnight. If an occupant returns after an intermediate break, the duration of intermediate absence value is computed by subtracting the current time from this temporarily held departure value. Then, it is rounded to the nearest hour (e.g., an intermediate break of 2h20 is rounded down to 2h). The distribution of the duration of intermediate absence periods was represented in terms of seven discrete likelihood values for intermediate

breaks between 0 and 6 h – in lieu of a standard continuous probability distribution. This is because the intermediate vacancy period distributions shown in Figure 2.16 had very different shapes for different occupants. With each new intermediate vacancy observation, the discrete likelihood weights are updated. These recursively updated weights are used to select the longest expected vacancy period D_{abs} . In a similar fashion, the longest expected intermediate presence period D_{int} was also computed.

The occupancy-learning algorithm is formulated as a standard BACnet program object – in scripting language GCL+ [93]. The algorithm is intended to be executed recursively. It is called to update the parameters only when an occupancy event takes place (e.g., the beginning and the end of an occupancy period). The code was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices. The details of the implementation results are presented in Chapter 5.

2.3 Summary

This section presents the development of an adaptive algorithm for PIR sensor-based presence detections and an adaptive algorithm for zone level presence predictions.

For the development of the presence detection method, data from three PIR, a door contact, and two CO₂ sensors and a fisheye lens camera gathered in a shared-office space. The analyses of the data and the ground-truth occupancy information provided following insights with regard to the use of PIR sensors in presence detection algorithms in shared-office spaces:

- When it is possible, a PIR sensor should be carefully positioned so that all seated areas remain within its coverage. In cases where a single PIR has to be

used and it cannot have a direct line-of-sight to all cubicles, placing a sensor facing the entrance would at least guarantee to detect each occupants' arrivals and departures. In this study, a single PIR sensor facing the door correctly interpreted 95% of the occupied periods and 93% of the unoccupied periods with a 14 min delay. If there are multiple PIR sensors available inside one room, the delay period can be reduced and our ability to detect absence can be increased. In this study, the PIR delay that can achieve 95% accuracy during occupied periods could be reduced to 5 min by using the three PIR sensors together.

- Departures by blending the door closing actions as a secondary source of information with a PIR sensor can be detected. Upon the detection of a departure with a door closing action, the PIR sensor's delay periods can be shortened. In this investigation, this did not affect the accuracy of the presence detections. The accuracy of absence detections was able to increase by 4% of the unoccupied times. This was equivalent to being able to employ occupancy-based control strategies for an extra 240 h/year.
- Occupants' impact over the CO₂ concentration response of the office can be filtered out by using a simple data-driven model trained with a particle filter. The model-based information about occupants' presence can be used in tandem with a PIR sensor. The results indicate that the use of CO₂ sensors slightly improve the accuracy of presence detections – particularly when the PIR delay period was less than 15 min. The improvement was more evident with a CO₂ sensor located away from the door.
- The distribution of the time between consecutive movement detections can be used to approximate the optimal PIR delay period when PIRs are positioned

to monitor all cubicles and/or face the door. For example, if 95% of the consecutive movement detections were separated by a time interval less than 10 min, choosing a 10 min PIR delay would correctly interpret about 95% of the occupied duration as presence periods. This approximation renders the potential of developing self-calibrating PIR sensors that can autonomously determine their optimal delay periods. To this end, a recursive controls algorithm that learns from the frequency of movement detections of a PIR sensor to adapt its delay value autonomously was formulated as a standard BACnet program object.

For the development of the zone level presence prediction algorithm, two year's worth of occupancy data from 16 private offices in an academic office building were analyzed. In this part of the study, following research steps were taken:

- Using the occupancy data, three different types of statistical model forms were built: (1) weekly occupancy schedules predicting the likelihood of presence as a function of the time of day and the day of week, (2) a discrete-time Markov model predicting likelihood of observing an arrival or a departure event in the next one hour, and (3) a survival model predicting the duration of an intermediate occupancy and a vacancy period.
- From these models, six key occupancy parameters were introduced: (1) the earliest expected time for the first arrivals ($t_{arr,early}$), (2) the latest expected time for the first arrivals ($t_{arr,late}$), (3) the latest expected time for the last departures ($t_{dpt,late}$), (4) the ratio of absent weekdays (p_{abs}), (5) the longest expected duration of intermediate presence (D_{int}), and (6) the longest expected duration of intermediate absence (D_{abs}). The diversity of these parameters between the 16 monitored offices was presented, and the effect of the number of occupants inside a zone on the uniqueness of these parameters was discussed.

- Further simplifications were made in the model forms to retrieve these six parameters recursively through an occupancy-learning algorithm embedded inside a local controller. The algorithm iteratively learns the recurring presence patterns inside an office from a PIR sensor's movement detections. The recursively learned parameters were employed by the temperature control algorithm (see Chapter 5) in selecting the temperature setback schedules. The code was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices. The details of the implementation results are presented in Chapter 5.

Now that we have the algorithms to detect and predict zone level presence, in the next chapter we will look at developing algorithms to learn from occupant's adaptive comfort.

Chapter 3

Learning from adaptive behaviours

This chapter first presents a review of the literature on the adaptive occupant behaviour modelling. Subsequently, the thermostat keypress actions with concurrent occupancy and temperature data in twenty office spaces were analyzed, and a recursive algorithm was formulated to develop thermostat use models inside local building controllers. In a similar fashion, the light switch and blinds use behaviours in ten private offices were analyzed, and a recursive algorithm was developed to learn from occupants' light switch and blinds use behaviour inside building controllers.

The chapter covers the development of the algorithms. The laboratory and field implementation of the thermostat learning algorithm for verification and further testing are presented in Chapter 5. The implementation results for the lighting and blinds learning algorithm are presented in Chapter 6.

3.1 Literature review: adaptive behaviour modelling

Occupants are not passive recipients of the indoor climates selected for them. They undertake adaptive behaviours to restore their comfort, when they feel uncomfortable [110]. These actions often involve adjusting their indoor environment through

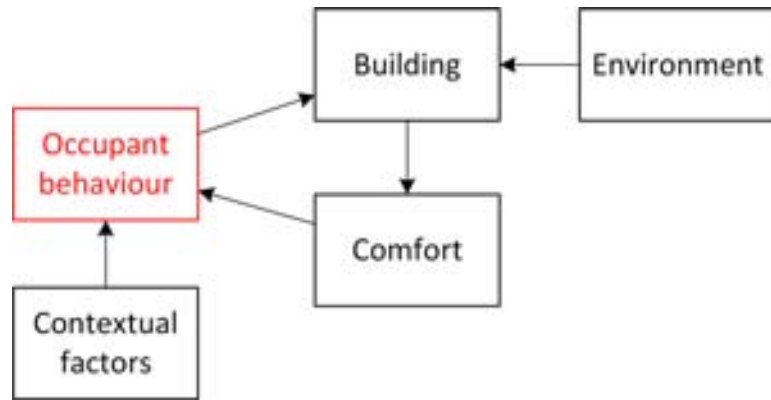


Figure 3.1: Interactions between a building and its occupants.

interactions with blinds, lighting, windows, and thermostats. The way these building components are used accounts for great uncertainty over a building's energy use and occupants' comfort [22,101]. Figure 3.1 schematically illustrates the dynamic interactions between a building and its occupants. The construct of comfort affects the way occupants behave, and occupants' behaviours affect the energy use of the building. The contextual factors in Figure 3.1 are occupant habits, attitudes, and beliefs that influence the way occupants undertake adaptive behaviours [43]. A common example of this is that occupants undertake adaptive behaviours less frequently in shared offices than in private offices [40]. Adaptive occupant behaviours have been widely studied through field surveys; and upon these surveys data-driven adaptive behaviour models have been developed [22]. The primary purpose of occupant behaviour modelling has been to better understand the occupants' influence on buildings' energy use. Only a few studies attempted to exploit the potential of adaptive occupant behaviour models to retrieve unsolicited information about occupants' adaptive comfort [48,111].

Model formalisms:

Although the model formalisms that mimic occupants' blinds, lighting, and window use behaviours are established [37,112,113], the research on modelling thermostat

use behaviour remains anecdotal [105, 110, 114, 115]. Occupants' thermostat use behaviour in offices has not been analyzed with concurrent occupancy and temperature data. Therefore, the literature on adaptive behaviour modelling is briefly presented as a whole.

In the reviewed literature, four different adaptive behaviour modelling formalisms were found: (1) schedules, (2) Bernoulli models, (3) discrete-time Markov models, and (4) discrete-event Markov models. The formalisms classify whether the models predict the occupants' adaptive actions or the state of the building components with which occupants interact.

The traditional way of modelling adaptive behaviours is building schedules – e.g., presenting the ratio of the lights switched on or the mean blind occlusion rate averaged over a week or a month [116]. The underlying assumption of this method is steady-periodicity. In other words, the model form predicates that the time of week or the month of year alone is adequate to make predictions for an adaptive occupant behaviour. This assumption arises from the fact that indoor and outdoor environmental factors that influence adaptive behaviours tend to recur in daily or seasonal cycles.

The second method used in adaptive behaviour modelling is the Bernoulli random processes [117]. The Bernoulli behaviour models predict the likelihood that an occupant leaves a building component at a certain state (e.g., probability of the lights being on) as a function of the predictor variables (e.g., ambient illuminance). These two modelling approaches cannot be used for modelling thermostat use behaviour because office occupants do not turn on or off the HVAC equipment directly; instead they determine the temperature setpoints. In addition, these two model forms attempt to predict the state of the building components with which occupants interact in lieu of the occupants' adaptive actions. Consequently, despite being appropriate

when the occupant models' purpose is to better represent a building's energy use, they do not provide any information about the occupants' adaptive comfort [118]. Although Bernoulli occupant models were developed with both indoor and outdoor explanatory variables in the literature [110, 117], Gunay et al. [119] discussed that they are more appropriate to be used with outdoor variables. This is because the adaptive behaviours trigger changes in the indoor environment which contains both the explanatory and the response variables.

The third method used in modelling adaptive behaviours is the discrete-time Markov chains [112]. The discrete-time Markov models predict the likelihood of undertaking an adaptive behaviour in the next timestep. They can be developed by both indoor and outdoor environmental variables, because they are derived upon conditions just before occupants undertake the behaviour. The Markov models treat adaptive actions and their reversals independently and are capable of predicting behaviour patterns realistically. Note that realistic representation of the frequency and timing of adaptive actions can act as a proxy to occupants' comfort. From thermostat overrides one can infer thermal discomfort conditions and from overrides to lighting and blinds automation systems one can infer visual discomfort conditions [16]. A common issue regarding the discrete-time Markov models is their dependency on fixed timesteps [120]. They only provide the likelihood of an event in the next timestep. Fixing the timestep implies the frequency of an occupant's instances of decision-making remains constant; it is logical that these instances increase in frequency during periods in which environmental conditions are rapidly changing (e.g., at arrival) [120].

Discrete-event Markov models (fourth method) link an occupant action model to an external event [37, 121]. For example, in Reinhart [37]'s light switch model, simulated occupants are modelled to turn on their lights at arrivals (event). In Rijal

et al. [121] occupants were modelled to consider window opening and closing upon a change in the predicted mean vote (event) [20]. In a similar fashion, Gunay et al. [118, 120] treated discrete-events for the light switch-on behaviour as a change larger than 100 lux in the workplane illuminance levels. When appropriate events triggering the behavioural adaptation of the occupants can be identified, the models' predictive accuracy are shown to improve in contrast to discrete-time Markov models [118]. However, discrete-event Markov modelling is challenged by finding an appropriate event definition to replace the timestep concept. Another limitation of this approach is that its predictive performance relies on the accuracy of the external events' predictions. For example, the predictive performance of the discrete-event Markov light switch model for arrival is subject to our ability to detect intermediate arrival and departure events accurately.

Regression methods:

In the reviewed literature on adaptive behaviour modelling, two regression methods were found: (1) linear regression (e.g., linear or polynomial regression) [122–124] and (2) generalized linear regression (e.g., logistic, probit regression) [40, 112, 117, 125–130, 130]

The advantage of the first method is that the regressed parameters have a closed-form solution – as the least-squares method can be used. The recursive least-squares method could be easily formulated inside controllers – in which computational and analytical capabilities are strictly limited. The handicap of this method is that linear regression is not appropriate for probabilistic models where the response variables are bound between 0 and 1.

The generalized linear regression is the de-facto standard in occupant behaviour modelling [113]. It employs a non-linear link function (e.g., probit or logit) to map the explanatory variables (e.g., indoor temperature) onto bounded response variables

(e.g., probability of observing a thermostat override). The maximum likelihood estimation (MLE) method is employed to solve for the parameter estimates of a generalized linear regression model. The challenge in conducting generalized linear regression inside a controller is that the MLE method does not provide a closed-form solution for the parameter estimates, meaning that solution needs to be converged iteratively. Because the literature has been focussing on developing adaptive behaviour models offline using batch data with access to established statistical tools and computational power, recursive formulation of a parsimonious algorithm to develop generalized linear models of adaptive behaviours inside building controllers remained an open research question.

3.2 Development of the thermostat learning algorithm

In this section, the thermostat keypress data from twenty offices were analyzed with concurrent indoor temperature and occupancy data. A recursive algorithm that learns from user interactions with thermostats to adapt the operating temperature setpoints was formulated as a standard BACnet program object – in the scripting language GCL+.

3.2.1 Analysis of the observational data

Concurrent indoor temperature, occupancy, and thermostat keypress data were gathered from twenty offices from two academic office buildings in Ottawa, Canada. The data in building 1 were collected only for five months. The data in building 2 were collected for 14 months. For both buildings, the heating season was the period during

which the central heating plant was operational – between October 5 and May 6. The cooling season was the period during which the chillers serving these buildings were operational between May 6 and October 5. Table 3.1 provides an overview of the monitored offices.

Each room had a commercial thermostat integrated to a BAS. Each thermostat contained a 10 k Ω thermistor sensor and a PIR motion detector. The accuracy of the thermistor-based temperature sensors was verified against a high accuracy reference thermometer (with a resistance temperature detector probe) by taking measurements from ten identical thermostats at four random instances. The mean and the standard deviation of the mean absolute errors were 0.3 and 0.2°C, respectively. The temperature measurements were taken at 15 min intervals and stored inside a local commercial BAS archiver. The PIR motion sensors' movement detections and the users' thermostat keypress actions were collected per event basis and stored in the same BAS archiver. The occupancy data records in each room were generated from the movement detections using the adaptive time delay algorithm (see Chapter 2).

The cooling and ventilation in each room were provided through a VAV terminal unit; and the heating was provided through ceiling-mounted radiant perimeter heaters and/or the VAV units' reheat coils. Through these HVAC units, the occupants were permitted to change their temperature setpoints up to $\pm 3^\circ\text{C}$ temporarily for the rest of the day. Once an override is ended, the setpoints were reset to the default setpoint value – which was 22°C for both heating and cooling seasons in the monitored buildings. Note that in rooms 10 to 19 the setpoint change requests might not be completely achieved during the cooling season as each of these rooms shares a VAV unit with three adjacent rooms (i.e., each VAV unit serves four rooms). Similarly, the thermostat setpoint increase requests during the cooling season or the setpoint

Table 3.1: List of rooms from which the data were collected.

	Rooms	Orientation	Shared/Private	Window	Profession
Building 1	1	South	Private	Fixed	Administrator
	2	East	Private	Operable	Professor
	3	West	Private	Operable	Professor
	4	East	Private	Operable	Administrator
	5	East	Private	Operable	Professor
	6	West	Private	Operable	Professor
	7	South	Private	Operable	Administrator
	8	South	Private	Operable	Administrator
	9	East	Private	Operable	Professor
Building 2	10	West	Private	Fixed	Professor
	11	West	Private	Fixed	Professor
	12	West	Private	Fixed	Professor
	13	West	Private	Fixed	Professor
	14	West	Private	Fixed	Professor
	15	West	Private	Fixed	Professor
	16	West	Private	Fixed	Professor
	17	West	Private	Fixed	Professor
	18	West	Private	Fixed	Professor
	19	West	Private	Fixed	Professor
	20	Southwest	Shared	Fixed	Grad Student

decrease requests during the heating season were often not achieved in the monitored offices, because the heating plant equipment was not functional during the cooling season and the cooling plant equipment was not operational during the heating season. This may have affected the frequency of thermostat interactions – however such coarse granularity in existing commercial buildings' terminal HVAC systems is not atypical. Regardless, the users' interactions with their thermostats provide invaluable insights with regard to their comfort.

In the dataset, the total occupied period was 11716 h. The breakdown of these occupied periods in different rooms is shown in Table 3.2. Sixteen of the rooms were used by academics (professors or graduate students) with flexible work hours. Consequently, the duration spent in individual rooms per unit time was low and diverse – by a factor of 18. In line with Karjalainen [30], occupants infrequently interacted with their thermostats. On average, one thermostat interaction tends to occur once every 35 h of occupancy. This value varied by a factor of 88 (between 7 and 648 h per thermostat override event). Over the monitored 11716 h occupied period, only 334 distinct thermostat override instances were recorded. At each unique thermostat override instance, the users pressed the thermostat keys at least twice (changing the temperature setpoint on average by 1.5°C). This is in line with Kempton et al. [31] that occupants overcompensate their thermostat control actions upon discomfort.

The breakdown of the distinct thermostat override instances in different rooms is also shown in Table 3.2. Of these thermostat override actions, two fifths were to increase the temperature setpoints and three fifths were to decrease the temperature setpoints. More importantly, about one third of the thermostat overrides were either to decrease the setpoint during the heating season or to increase the setpoint during the cooling season. This indicates that the conservatism in maintaining a minimum

Table 3.2: List of rooms from which the data were collected.

Rooms	Monitored occupied period (h)	Unique thermostat interactions	Average setpoint change ($^{\circ}\text{C}$)
1	420	57	1.0
2	384	3	0.7
3	457	5	0.6
4	523	37	0.8
5	388	1	0.6
6	73	2	0.6
7	582	12	1.1
8	422	28	0.8
9	455	45	0.6
10	129	7	0.9
11	285	2	2.0
12	1044	8	3.0
13	795	16	1.0
14	721	7	1.0
15	361	12	1.7
16	547	38	3.0
17	475	4	2.4
18	917	2	0.5
19	1296	2	1.0
20	1446	46	2.3
Total	11716	334	

of 22°C during the heating season and a maximum of 22°C during the cooling season not only wastes energy but also frustrates some occupants.

The occupants in the rooms with operable windows interacted with their thermostats about three times more frequently than the occupants in the rooms with fixed windows. However, a conclusive interpretation could not be made because of the small sample size. Figure 3.2 anecdotally presents a data record over a ten day period in February from room 4. The outdoor temperature during this period was less than -10°C . Despite the fact that this was the heating season, the occupant

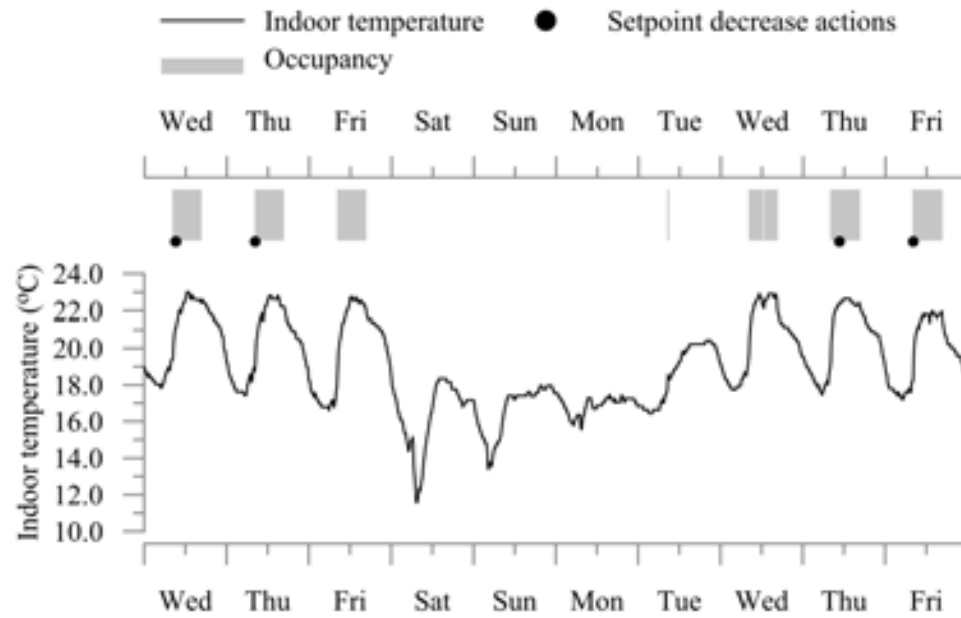


Figure 3.2: Thermostat use behaviour during the heating season over an illustrative ten day period in room 4.

consistently decreased the temperature setpoint as the indoor temperature increased over 22°C – except on the first Friday and the second Wednesday. Upon occupant's departure on the first Friday, the indoor temperature dropped as low as 12°C . Given that the central HVAC equipment were operational and the setback temperature was 18°C , this can be interpreted as the occupant – upon overheating on Friday – chose to open the window in lieu of decreasing the temperature setpoint and forgot the window open upon departure.

Figure 3.3 presents the distribution of the indoor temperatures just before the thermostat override instances. The results indicate that the indoor temperatures leading to a setpoint increase are different in heating and cooling seasons. The median of the thermostat overrides to increase the setpoint was at $\sim 20^{\circ}\text{C}$ in the heating season and $\sim 23^{\circ}\text{C}$ in the cooling season. On the other hand, the median of the thermostat overrides to decrease the setpoint was at $\sim 22^{\circ}\text{C}$ in the heating season




Figure 3.3: Distribution of the indoor temperatures just before the thermostat override instances.

and $\sim 25^{\circ}\text{C}$ in the cooling season. This indicates that the users' indoor temperature expectations change seasonally. This is mainly attributed to the changes in occupants' clothing insulation levels [17]. However, the seasonal changes in occupants' clothing assembly are likely to be affected by the default temperature setpoints imposed by the automation. For example, during the summer occupants may choose to wear heavier clothing than they would normally prefer knowing that it will be cold in their air-conditioned offices.

Parallel to these observations, Figure 3.4 presents the indoor and outdoor temperatures at the thermostat keypress instances. Results indicate that the mean of the indoor temperatures leading to a keypress action increases by about 1°C for each 10°C increase in the outdoor temperature. The temperature difference between setpoint increase and decrease instances was less than 1°C . The narrow band between the setpoint increase and decrease instances underlines the challenge in finding a single automation setpoint that would satisfy the comfort of the majority.

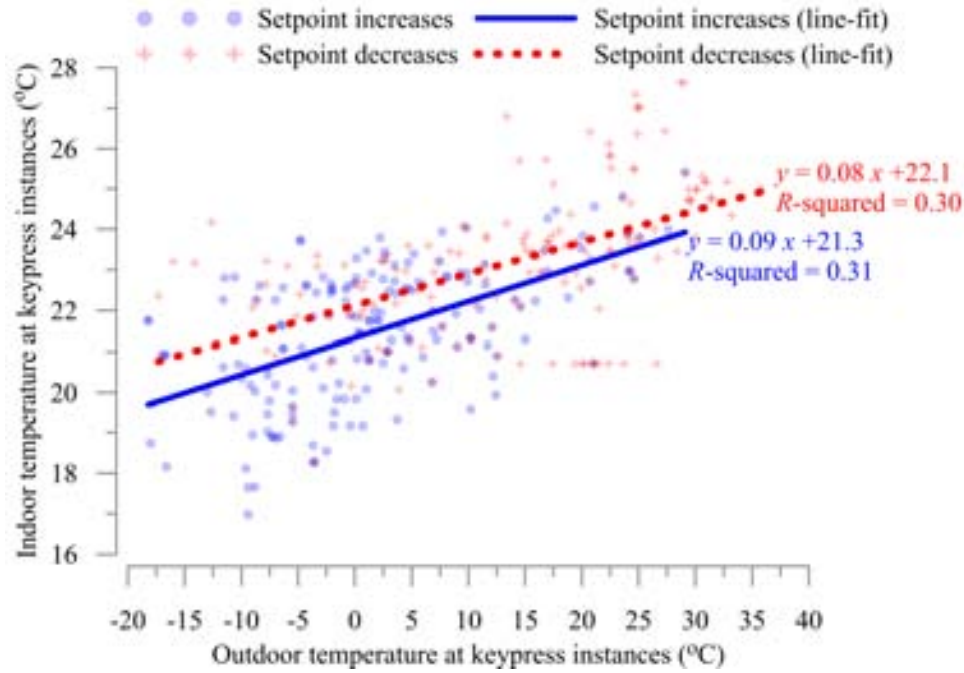


Figure 3.4: Indoor and outdoor temperatures at thermostat keypress instances.

Figure 3.5 present two discrete-time Markov models predicting the likelihood of observing a thermostat override action in the next one hour as a function of the indoor temperature. The areas of the bubble plots represent the occupied duration spent at a certain indoor temperature. The areas are relative to the scales shown for each plot. The probabilities were calculated by dividing the total number of thermostat override observations to the total number of occupied hours. The line-plots indicate the univariate logistic regression models fitted by employing the generalized linear modelling functionalities of an established statistical tool (R-statistics). While fitting the model, the regression method weighs the discrete likelihoods at different temperatures in accordance with the amount of observation used in building them. The results indicate that a setpoint increase action is about four times more likely when it is 20°C than 25°C; while a setpoint decrease action is about five times less likely when it is 20°C than 25°C. The frequency of thermostat interactions (both setpoint

increases and decreases) remained at a minimum of once every 40 h between 21 and 24°C for the studied occupant population. Based on the results of this data analysis, an occupant's frequency of thermostat interactions appeared to be a reflection of his/her temperature preferences; and it changes from one season to another.

In a similar fashion, Figure 3.6 presents two discrete-event Markov models predicting the likelihood of observing a thermostat override action at arrival (intermediate or first arrival of the day). The relatively small standard errors and p-values in Figures 3.5 and 3.6 can be interpreted that the occupants' overall setpoint increase/decrease behaviour can be approximated as both discrete-time and event univariate logistic regression models. The Akaike Information Criterion (AIC) values annotated on the figures indicate that the discrete-event Markov models – whereby the events were defined as the arrivals – appear to be more appropriate than the discrete-time Markov models (see Figure 3.5) in predicting the thermostat use actions. However, given that only 141 of the 334 thermostat use actions took place at the timestep following the arrival the discrete-time Markov models were employed in formulating the thermostat learning algorithm.

3.2.2 Thermostat learning algorithm

An approximate recursive method to develop logistic regression models inside local building controllers was formulated. This method alleviates the logistic regression model development process from storing an array of null and positive observations, and the need for iterative computations to converge to the parameter estimates – recall that the MLE for logistic regression does not have a closed-form solution.

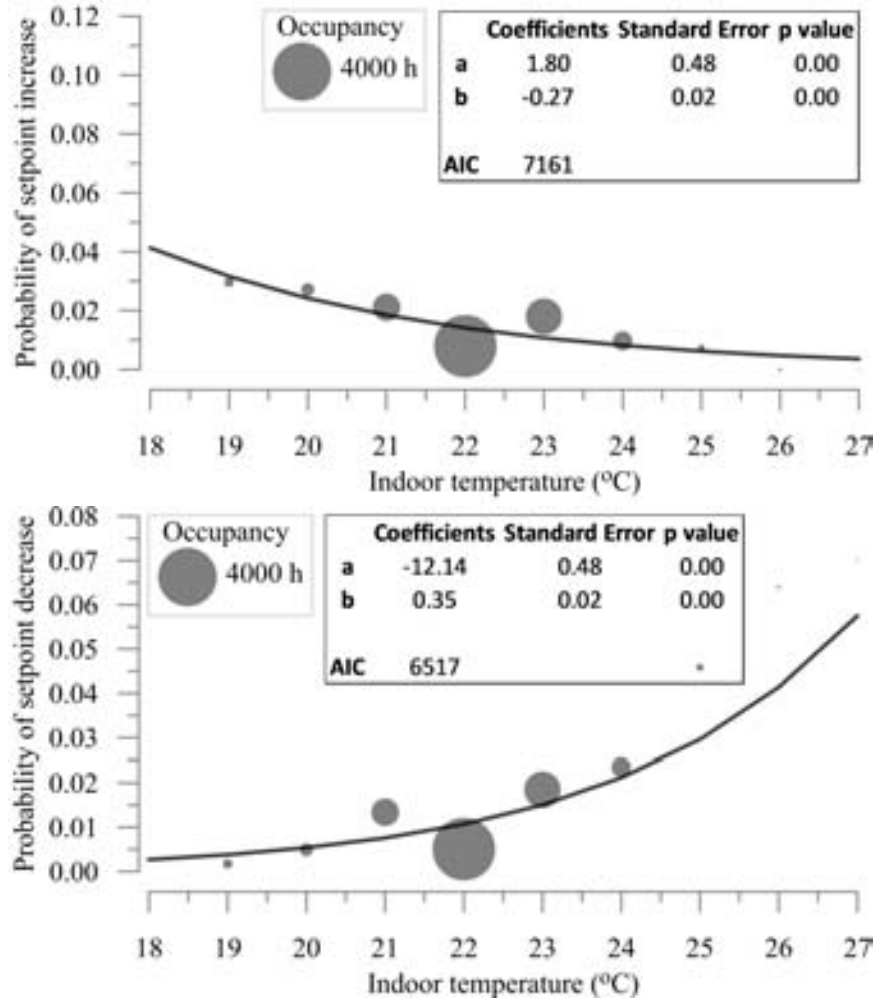


Figure 3.5: Probability of (a) increasing and (b) decreasing the temperature setpoint in the next one hour. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bT_{in})}}$. The properties of the regression parameters were annotated in the figure.

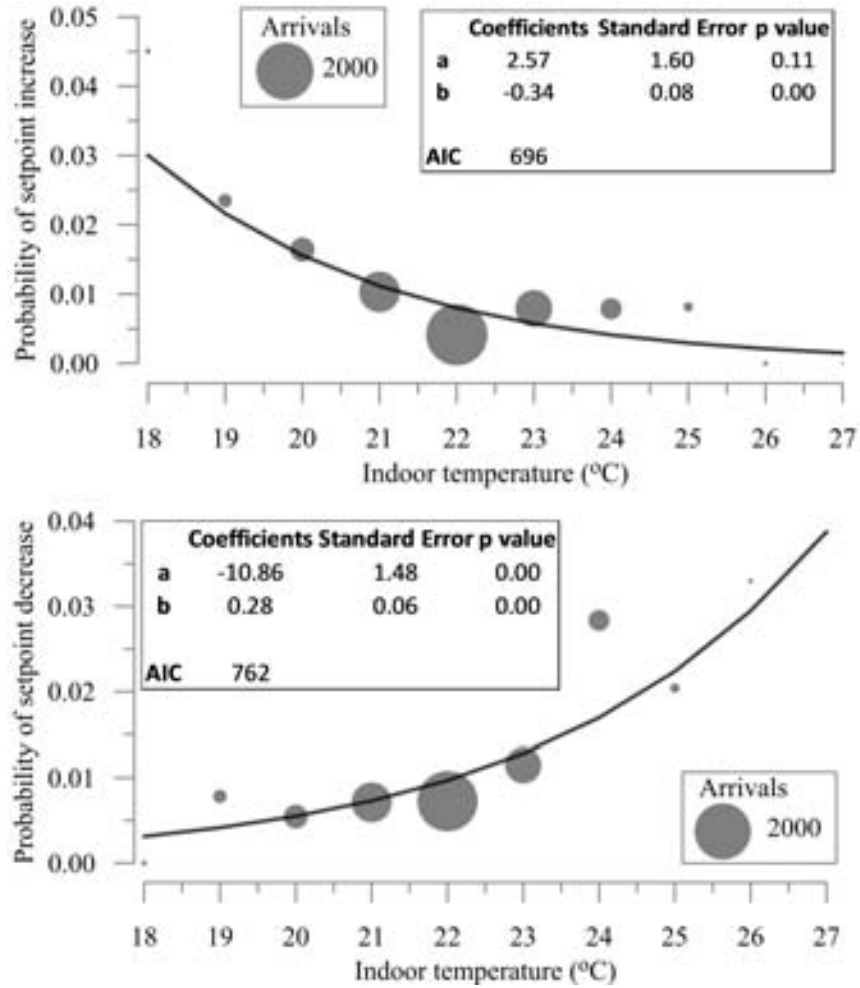


Figure 3.6: Probability of (a) increasing and (b) decreasing the temperature setpoint at arrival. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bT_{in})}}$. The properties of the regression parameters were annotated in the figure.




Figure 3.7: An illustration of the recursive parameter update process in a logistic regression model.

The algorithm initiates with a rough guess for the setpoint decrease and increase models. An example of this is illustrated in Figure 3.7 in solid black line. Recall that a univariate logistic regression model is in the following form: $p = \frac{1}{1+e^{-(a+bT_{in})}}$. A logical prior guess can be the parameter a is -6 and b is 0.25 for setpoint decrease, and the parameter a is 5 and b is -0.25 for setpoint increase.

When a setpoint increase or decrease event is registered, the model's prediction at that temperature is updated as illustrated in Figure 3.7. In addition, when an occupied hour spent without a setpoint increase or decrease (null observation), the model's prediction at that temperature is updated (see Figure 3.7). The updating power of a (null or keypress) observation depends on two factors: (1) the innovation – the difference between the model's prediction and the observation. For example, observing a setpoint increase when it is 15°C would not update the setpoint increase model as much as observing it at 20°C. (2) the number of (null or keypress) observations that the model has been built upon. As more observations are blended in the model, the updating power of a single observation will diminish.

Upon updating the model's prediction at the temperature the last (null or keypress) observation was acquired, the model parameters a and b need to be updated. The challenge in doing this is that there is an infinite number of a and b combinations that can satisfy the new model prediction at that temperature. A common method for this problem is to perturb the parameters towards the steepest gradient [131]. Conveniently, the derivative of a univariate logistic thermostat use model $p = \frac{1}{1+e^{-(a+bT_{in})}}$ with respect to its parameters a and b at an indoor temperature T_{in} equals: $p(T_{in})(1-p(T_{in}))$ and $p(T_{in})(1-p(T_{in}))T_{in}$, respectively. The change in the model's predictions at the temperature the last (null or keypress) observation acquired was reflected onto parameters a and b proportional to their respective derivatives. In machine learning literature, this simple procedure is known as the stochastic gradient descent for backpropagation in neural networks – where the learning rate is linked to the number of observations used in building the model [132].

Figure 3.8 illustrates the recursive evolution of a setpoint increase and decrease model predicting the likelihood of a thermostat override action over a year in room 16. The models rapidly change from their initial guesses and stabilize in time. On the other hand, Figure 3.9 presents the evolution of the parameter estimates of these models in time. Results indicate that parameter estimates a and b converge to relatively stable values in about a month – albeit rather slowly changing as more keypress and null observations are processed during the operation.




Figure 3.8: The recursive evolution of the setpoint decrease and increase models for room 16 in time over a year. The darkness of the line plots increases as time progresses and the thick red line plots represent the last state of the model.




Figure 3.9: The recursive evolution of the parameter estimates a and b for the models shown in Figure 3.8.

The recursively estimated models were compared with the measured discrete likelihood densities and the traditionally developed logistic regression models (using a batch of observations and the MLE method). Figure 3.10 illustrates a comparison between the models developed using these two methods and the measured discrete likelihood densities for room 16. The appropriateness of this recursive approximation method to develop logistic thermostat learning models was verified. The minor discrepancies between the batch and the recursively estimated models can be explained with the seasonal changes in the thermostat use behaviour. The recursive method weighs the keypress and the null observations with the expectation that these observations are homogeneously distributed in time. To tackle this issue, a constant drift coefficient can be introduced, providing a marginally larger weight to the latest observation than otherwise it would [133]. Alternatively, the parameter learning can be reinitiated at seasonal switchover to heating and to cooling.

The thermostat learning algorithm is intended to adapt the temperature setpoints during operating hours in the heating and the cooling seasons inside individual offices based on the recursively developed logistic regression models. The algorithm selects the setpoint during the heating season such that the likelihood of a setpoint increase event is 0.025 for the next one hour (once every 40 h of occupancy). Similarly, it selects the setpoint during the cooling season such that the likelihood of a setpoint decrease event is 0.025 for the next one hour (once every 40 h of occupancy). As the occupied time elapses during a heating season without a setpoint increase action, the setpoint satisfying this 2.5% chance criterion decreases slowly. Analogously, as the occupied time elapses during a cooling season in absence of setpoint decrease, the setpoint satisfying the 2.5% chance criterion increases slowly. The slow decrease in heating and the slow increase in cooling setpoints are interrupted by the thermostat

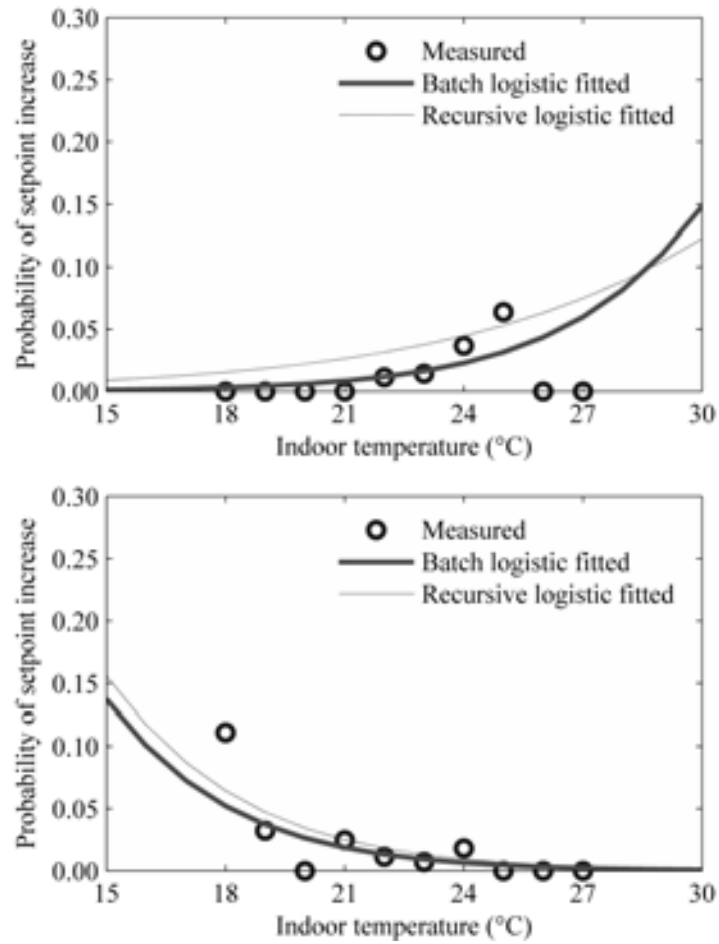


Figure 3.10: A comparison of the recursively estimated model with the measured discrete likelihood densities and the traditionally developed logistic regression models (using a batch of observations and the MLE method) for room 16.

keypress actions. Upon a setpoint increase action, the heating setpoint satisfying the 2.5% chance criterion increases; and after a setpoint decrease action, the cooling setpoint satisfying the 2.5% chance criterion decreases. When a setpoint change is registered, it is executed until the end of the operating hours that day. The new thermostat setpoint is applied the next day based on the logistic regression models – updated by the recursive thermostat learning algorithm. The thermostat learning algorithm was formulated as a standard BACnet program object – in scripting language GCL+ [93]. The code was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices. The details of the implementation results are presented in Chapter 5.

3.3 Development of the lighting and blinds learning algorithm

In this section, the light switch and internal blinds position data from ten West-facing private offices were analyzed with indoor illuminance, solar irradiance on the facade, and occupancy data. A recursive algorithm that learns from user interactions with light switches and BAS-integrated automated blinds (if available) to adapt the lighting and blinds automation setpoints was formulated as a standard BACnet program object – in the scripting language GCL+.

3.3.1 Analysis of the observational data

The light switch, internal blinds position, occupancy, indoor illuminance and solar irradiance data were gathered from ten West-facing private offices from an academic office building in Ottawa, Canada. They were used by ten full-time professors in the

Faculty of Engineering. The offices had identical geometry, construction, material, and occupancy characteristics. Each office had 4 m² West-facing double-glazed exterior windows with 70% visible transmittance. The window-to-wall ratio on the facade was 32% and the window-to-floor ratio was 24%. Each office was equipped with manually controlled top-down roller blinds. They had 14% visible transmittance. The reflectances (in the visible spectrum) of the floor, walls, and the ceiling were measured using a spectrophotometer as 0.2, 0.5, and 0.8, respectively.

The light switch, occupancy, and indoor illuminance data were extracted from the BAS archiver. The solar irradiance data on the facade were collected from a local weather station. The blinds positions were monitored through a time-lapse camera facing the facade of the building. Due to interruptions in the data acquisition, the length of individual data records were between five months to one year.

The indoor illuminance readings were taken by photodiode sensors (spectral response range 350 to 1100 nm and 60° field of view). The sensors were consistently mounted on the ceilings of the monitored offices 3 m away from the exterior window and above the workplane. They were able to measure illuminance between 0 and 1076 lux. In order to verify their precision, before the installation they were mounted closely on a single ceiling tile in the controls laboratory. The sensors were verified that they can consistently measure the indoor illuminance (± 15 lux). The indoor illuminance data were collected at 15 min intervals. Each room had a PIR motion sensor (5 m range and 100° horizontal and 80° vertical coverage). The PIR motion sensors' movement detections and the users' light switch actions were collected per event basis and stored in the same BAS archiver. The occupancy data records in each room were generated from the movement detections using the adaptive time delay algorithm (see Chapter 2).




Figure 3.11: The weekly lighting schedules in the ten monitored offices.

Each office was equipped with four recessed T8 fluorescent lighting with five state dimmable ballast. As the floor was covered by a low reflectance carpet (~ 0.2 based on the reflectometer measurements), the effect of electric lighting on the ceiling mounted photodiodes were minimal (less than 30 lux) in contrast to the operating range of the photodiodes (between 0 and 1076 lux). In the ten offices, occupants were exclusively responsible for turning on their lighting. The users were able to turn off or dim their lighting. Alternative to the manual switch off, upon vacancy (programmed as the absence of detectable movements for more than 30 min), the lights automatically turned off.

Figure 3.11 presents the lighting use schedules from the ten offices. Results indicate that lighting use in individual rooms vary substantially. Typically, during a day, a maximum of two or three of the ten offices used electric lighting simultaneously. This is a reflection of the diversity in individuals' lighting and blinds use behaviour, and their occupancy patterns. Figure 3.12 presents the lighting use in the ten monitored offices normalized against occupancy. Results indicate that the lighting use in individual rooms differed by a factor of 13 – from 10% of the occupied




Figure 3.12: Weekly lighting schedules in individual rooms plotted against weekly occupancy schedules.

time to 130% of the occupied time. In two of the rooms, the electric lighting use exceeded the total duration of occupancy. This was possible because during this monitoring period the lights were automated to turn off only 30 min after the last motion detection (i.e., lights remained on for 30 min after the rooms were vacated). Figure 3.13 groups the mean lighting use in the ten monitored offices normalized against occupancy when the peak daily solar altitude was higher than 40° (March to September) and lower than 40° (October to February). Results indicate that the occupants' lighting use change seasonally – with the changes in the solar geometry and daylight hours.




Figure 3.13: The mean weekly lighting use with respect to the mean weekly occupancy when the peak daily solar altitude was higher than 40° (March to September) and lower than 40° (October to February).




Figure 3.14: The ratio of lights on during occupied hours as a function of the solar irradiance on the facade in individual offices.

Figure 3.14 presents the ratio of lights on during occupied hours as a function of the solar irradiance on the facade. Results indicate that three of the ten occupants' lighting use behaviour exhibit no relationship with the daylight availability. Other occupants appeared more responsive to daylight availability, albeit at varying rates. Even when the solar irradiance on the facade reached to its maximum at about 800 W.m^{-2} , the lights were on about 40% of the occupied periods. This was about 70% when the solar irradiance was less than 25 W.m^{-2} . This can be interpreted as the occupants did not actively fine-tune their blinds to exploit daylight availability to reduce their electric lighting use. In line with this interpretation, the mean blind occlusion rate in the monitored offices does not change as a function of the solar irradiance (see Figure 3.15). The mean blind occlusion rate appears to be individual for each user – ranging from 0 to 90% in these identical office spaces.




Figure 3.15: Mean blind occlusion rate as a function of the solar irradiance on the facade.

Although the mean blind occlusion rate did not change as a function of the solar irradiance, the occupants' blinds closing actions were heavily influenced by the incident solar irradiance on the facade. Figure 3.16 presents the discrete likelihoods of observing a blinds closing action in the next 60 minutes – when the blinds were more than half open and the occupants were present. The area of the bubble plots indicate the duration of occupied periods observed at a particular solar irradiance level – when the blinds were more than half open. The blinds closing actions were events with a return period of 20 h occupancy when the blinds were more than half open and the solar irradiance was less than 50 W.m⁻². Above 800 W.m⁻², the return period becomes once every 3 h or less. Therefore, occupants actively respond to visual discomfort conditions due to glare.

The handicap in occupants' blinds use behaviour, which led to stagnant mean blind occlusions as shown in Figure 3.15, was that it sometimes took weeks before they reopen their blinds. This phenomena can be seen in the survival model shown




Figure 3.16: Probability of closing blinds in the next 60 min as a function of the incident solar irradiance on the facade. The univariate logistic regression model is in the following form: $p = \frac{1}{1+e^{-(a+bS_{\text{rad}})}}$. The properties of the regression parameters were annotated in the figure.




Figure 3.17: A survival model for the lifetime of a blinds position and a Weibull distribution fit for this model. The Weibull model is in the following form: $p = 1 - e^{-(t/A)^B}$. The properties of the regression parameters were annotated in the figure.

in Figure 3.17. After a blinds closing action, there was more than 50% chance that the reopening action would take place in more than three days. More than 30% of the reopening actions followed more than seven days after the closing actions. Consequently, the daylight potential to offset electric lighting during these periods were lost. As mentioned earlier, the occupants had characteristic blinds positions. They tend to partially close their blinds to avoid glare; and when they reopen them, they tend to return their blinds to their preferred position. This resulted in with mostly partial blinds opening and closing increments (either quarter or half opening/closing actions) shown in Figure 3.18. This is important because partly open blinds may allow daylight, avoid glare, and provide an opportunity to establish view and connection to outdoors. Note that the blinds positions were recorded in five states (open, quarter closed, half open, quarter open, closed).

The light switch on behaviour was modelled with respect to the indoor illuminance readings taken by the ceiling-mounted photodiode sensors. The results shown in




Figure 3.18: The distribution of the blinds closing and opening increments at four discrete levels.

Figure 3.19.a indicate that the frequency of the light switch on events depends heavily on the indoor daylight levels. When an occupant was present and the lights were off, a light switch on event was observed once every 50 h of occupancy at a ceiling illuminance of 200 lux. This became once every 30 min of occupancy when the ceiling illuminance was less than 50 lux. Therefore, the frequency of occupants' interactions with their light switches may be treated as a proxy to their preferred visual comfort conditions. Previous research [22] suggests that occupants are more likely to switch on their lighting at arrival than during intermediate occupancy. In line with this, in our dataset 90% of the light switch on events were observed right after the occupant's arrival (both the intermediate and the first arrivals of the day). Consequently, it was argued that modelling the light switch on behaviour as events linked to the occupants' arrival may be more appropriate than modelling it as processes linked to the time spent at a certain illuminance level [118]. However, a substantial difference in the model's quality in representing the light switch on behaviour as a discrete-event or a discrete-time random process was not observed. The Akaike Information Criterion (AIC) values annotated in Figure 3.19.a and b indicate that for the same dataset and

the number of degrees of freedom, the discrete-time model achieved slightly better representation of the observed behaviour.

When the light switch on models were developed for each room individually, it was realized that the light switch on behaviour was remarkably diverse (see Figure 3.20). It is worth noting that these models normalized the behaviour patterns with ceiling illuminance and occupancy. In addition, recall that these rooms had identical geometry, construction, orientation, and occupancy type. Arguably, these variations would only exacerbate the diversity amongst the occupants' light switch on behaviour. This underlines the differences in users' visual comfort preferences and the challenges in finding a fixed and prescribed daylight-integrated lighting automation setpoint for everyone – even when the photodiode sensors were precisely located at the same position. The regression parameters for the light switch on models from individual rooms are tabulated in Table 3.3.

More than 90% of the light switch off events were recorded at the departure instances. About 50% of these events were manual switch off actions. Other light switch off events were triggered upon the detection of vacancy. In other words, the occupants left their lights on at about 50% of the departure instances. More importantly, the occupants almost never turned off their lights during intermediate occupancy – regardless of the daylight availability. It is observed that the likelihood of a light switch off at departure was more likely when the occupants were leaving their offices at night or for the weekend (see Figure 3.21) than when they were leaving for an intermediate break. This behaviour has no logical connection to saving energy by consciously avoiding the lights to remain switched on overnight, during weekend or




Figure 3.19: Probability of switching on the lights (a) in the next 15 min (discrete-time) and (b) at arrival (discrete-event) as a function of the ceiling illuminance. The univariate logistic regression model is in the following form: $p = \frac{1}{1+e^{-(a+bE_{\text{lux}})}}$. The properties of the regression parameters were annotated in the figure.

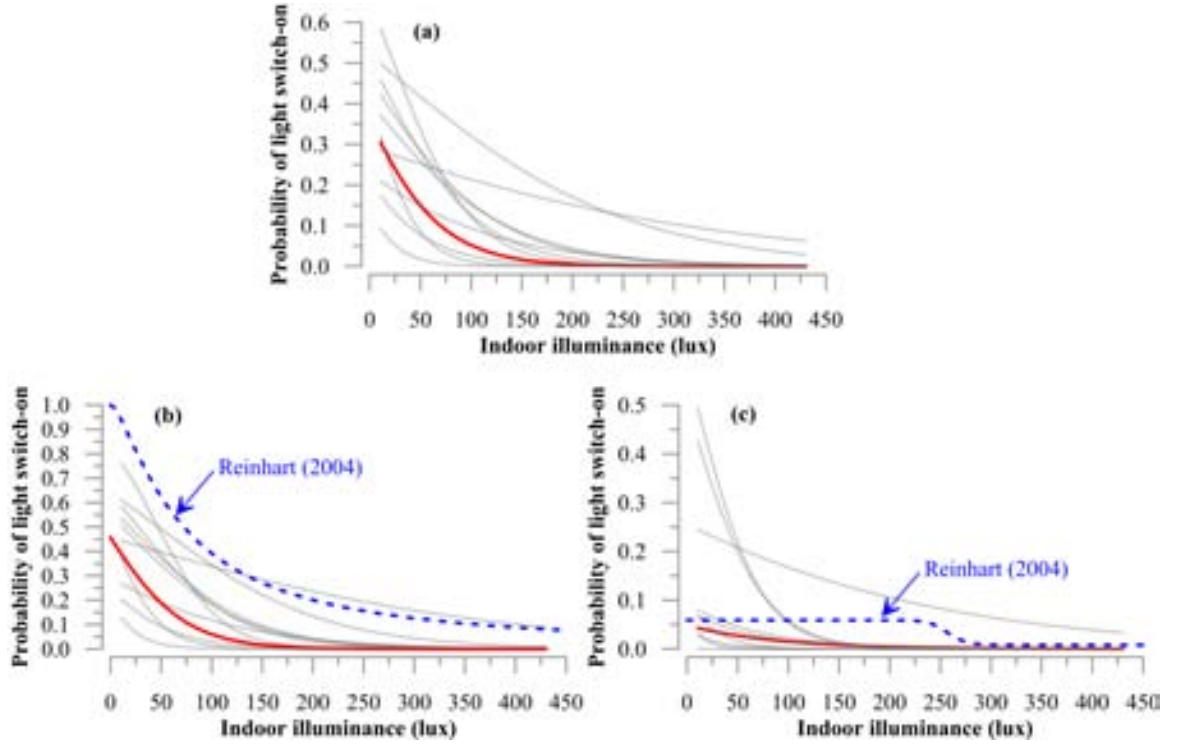


Figure 3.20: Probability of switching on the lights (a) in the next 15 min (discrete-time), (b) at arrival, (c) during intermediate occupancy in individual rooms. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bE_{\text{lux}})}}$. The thick red line represents the mean of the ten occupants. The regression parameters are tabulated in Table 3.3. For comparison purposes, Reinhart [37]'s lightswitch model for arrival and intermediate occupancy were included. Note that Reinhart [37]'s model was developed upon workplane illuminance measurements, whereas the models developed in this study were based on the indoor illuminance measurements taken from the ceiling.

Table 3.3: The logistic regression parameters for light switch behaviour in individual rooms as a function of the ceiling illuminance.

Rooms		Switch on probability at arrival			Switch on probability during intermediate occupancy in 15 min			Switch on probability in 15 min		
		Coefficients	Standard Error	p-value	Coefficients	Standard Error	p-value	Coefficients	Standard Error	p-value
R1	a	-1.094	0.652	0.094	-2.456	0.957	0.010	-1.306	0.636	0.040
	b	-0.024	0.012	0.045	-0.016	0.017	0.244	-0.023	0.012	0.045
R2	a	0.579	0.655	0.276	-14.648	0.352	0.100	0.085	0.607	0.289
	b	-0.012	0.010	0.199	0.000	0.937	0.100	-0.008	0.009	0.257
R3	a	-0.170	0.568	0.265	-1.072	0.635	0.092	-0.875	0.422	0.038
	b	-0.005	0.010	0.224	-0.005	0.012	0.352	-0.004	0.007	0.261
R4	a	0.327	0.918	0.222	0.317	1.552	0.338	-0.151	0.732	0.237
	b	-0.017	0.014	0.207	-0.032	0.024	0.193	-0.015	0.010	0.128
R5	a	0.208	0.933	0.223	0.006	1.235	0.296	-0.379	0.689	0.282
	b	-0.017	0.014	0.213	-0.028	0.021	0.173	-0.013	0.010	0.162
R6	a	1.522	0.141	0.000	-3.666	0.461	0.000	0.620	0.109	0.000
	b	-0.032	0.003	0.000	-0.005	0.005	0.232	-0.026	0.002	0.000
R7	a	-0.064	0.188	0.236	-3.253	0.694	0.000	-0.314	0.182	0.084
	b	-0.042	0.004	0.000	-0.034	0.015	0.022	-0.041	0.004	0.000
R8	a	-1.399	0.251	0.000	-2.511	0.574	0.000	-1.798	0.242	0.000
	b	-0.050	0.008	0.000	-0.080	0.029	0.006	-0.045	0.008	0.000
R9	a	0.559	0.209	0.007	-2.194	0.686	0.001	0.048	0.189	0.298
	b	-0.020	0.005	0.000	-0.039	0.023	0.092	-0.020	0.004	0.000
R10	a	-0.883	0.165	0.000	-2.291	0.349	0.000	-1.214	0.160	0.000
	b	-0.011	0.002	0.000	-0.015	0.004	0.000	-0.011	0.002	0.000




Figure 3.21: Probability of a light switch off at departure with respect to the duration of absence following the departure. The univariate logistic regression model are in the following form: $p = \frac{1}{1+e^{-(a+bD_{\text{abs}})}}$. The regression parameters are tabulated in the figure. For comparison, another light switch off model extracted from an independent study [45] is shown in the figure.

holidays, because the lights were automated to switch off in 30 min after the departure (time delay in the motion sensor). In other words, leaving without turning off their lights causes the same amount of electricity to be wasted for a lunch break or for a long vacation. Reinhart and Voss [45] analyzed the light switch off behaviour in office spaces and reached the same conclusions. As shown in Figure 3.21, their results were in line with our observations.

Insights gathered in this data analysis can be summarized as follows:

1. Occupants consciously and actively undertake blinds closing actions upon glare and light switch on actions at insufficient daylight. However, they often fail to turn

off their lighting or open their blinds once the conditions leading to visual discomfort disappears.

2. The indoor illuminance levels leading to visual discomfort vary substantially even in offices with identical geometry, orientation, and construction characteristics. It is therefore challenging to develop daylight-integrated blinds and lighting automation with fixed and prescribed illuminance setpoints.

3. The frequency at which the occupants close their blinds or turn on their lights can act as a proxy to their preferred visual comfort conditions. Discrete-time Markov logistic regression models that employ a ceiling-mounted photodiode sensor's readings to predict the likelihood of a light switch on or a blinds closing action during occupied hours appears to represent this phenomenon adequately.

3.3.2 Lighting and blinds learning algorithm

Analogous to the thermostat learning algorithm, a recursive algorithm was formulated to develop discrete-time Markov logistic regression models that predict the light switch on and the blinds closing behaviours inside a building controller. The models were intended to be used in selecting adaptive illuminance setpoints for daylight-integrated lighting and/or blinds automation (when available).

Recall that this recursive approximation was needed because today's building controllers do not have the analytical and computational capabilities for conducting a batch logistic regression. In addition, the datasets to train the batch logistic regression models become available in time only after the building becomes operational – with no guarantee that these datasets will remain ergodic.

The algorithm initiates with a rough guess for the light switch on and blinds closing models (when BAS-integrated automated blinds are available). Recall that a univariate logistic regression model is in the following form: $p = \frac{1}{1+e^{-(a+bE_{\text{lux}})}}$. A

logical prior guess can be where the parameter a is 1 and b is -0.001 for light switch on, and the parameter a is -2 and b is -0.004 for blinds closing.

Figure 3.22 illustrates the simple recursive learning process. As it was in the thermostat learning algorithm, there were two types of observations: (1) light switch on observations (or blinds closing actions for the blinds use model), (2) null observations in which the occupant spends a 15 min period without switching on the lights (or a 60 min period without closing the blinds). The updating power of an observation depends on two factors: (1) The difference between the model's prediction and the observation. For example, observing a light switch on when it is 20 lux would not update the light switch on model as much as observing it at 200 lux. This is reflected in the derivative of the quadratic error terms $\frac{de}{da}$ and $\frac{de}{db}$. (2) The number of observations that the model has been built upon. As more observations are blended in the model, the updating power of a single observation will diminish. This is reflected in the learning rate term η in the update process. With each new observation, the parameters a and b of the logistic regression model describing the light switch on or the blinds closing behaviour were perturbed. Figure 3.23.a illustrates the recursive evolution of a light switch on model predicting the likelihood of a light switch on action in one of the ten monitored offices. The model rapidly changes from the initial guesses and stabilizes in time. As illustrated in Figure 3.23.b, the recursively estimated models were verified against the measured discrete likelihood densities and the traditionally developed logistic regression models (using a batch of observations and the MLE method).

The lighting and blinds learning algorithm is intended to find out the brightest expected conditions that will result in a light switch on action ($E_{\text{light,sp}}$) and the darkest expected conditions that will result in a blinds closing action ($E_{\text{blinds,sp}}$) in




Figure 3.22: The recursive update procedure for the logistic regression models predicting the light switch on or the blinds closing behaviour.

individual offices. The blinds and lighting control algorithm (Chapter 6) employs these values in making automation decisions to switch off lighting and to open blinds during occupied hours. The algorithm defined the $E_{light,sp}$ as the illuminance level with 5% chance of a switch on probability in the next 15 min (an event with a return period of 5 h). The $E_{blinds,sp}$ were defined as the illuminance levels with 5% chance of a blinds closing action in the next 60 min (an event with a return period of 20 h).

As the occupied time elapses without a light switch on action (when the lights are off), the $(E_{light,sp})$ decreases slowly. As the occupied time elapses without a blinds closing action (when the blinds are open), the $E_{blinds,sp}$ increases slowly. The slow decrease in $E_{light,sp}$ and the slow increase in $E_{blinds,sp}$ are interrupted by the light switch or blinds closing actions. Upon a light switch action, the $E_{light,sp}$ increases; and after a blinds closing action, the $E_{blinds,sp}$ decreases. If the blinds are not integrated to the BAS and/or not automated, the learning process can be employed only for identifying $E_{light,sp}$.

The lighting and blinds learning algorithm was formulated as a standard BACnet program object – in scripting language GCL+ (see [93]). The code was implemented and tested inside the laboratory's controls network for the automation of both blinds and lighting. It was implemented inside three local controllers serving five private




Figure 3.23: (a) The recursive evolution of the light switch on models in one of the ten offices over a year. The darkness of the line plots increases as time progresses and the thick red line plots represent the final state of the model. (b) A comparison of the recursively estimated model with the measured discrete likelihood densities and the traditionally developed logistic regression models (using a batch of observations and the MLE method) for the same room.

offices only for lighting automation. The details of the implementation results are presented in Chapter 6.

3.4 Summary

This chapter presents the development of two recursive algorithms to learn from occupants' adaptive behaviours. Of them, one was intended for the thermostat use behaviour and the other one was intended for light switch and blinds use behaviour.

For the development of the thermostat-learning algorithm, the thermostat key-press data from 20 private offices were collected with concurrent occupancy and temperature data. The analyses of the data provided following insights with regard to the thermostat use behaviour of the occupants and their adaptive thermal comfort:

- The occupants infrequently interact with their thermostats (on average once every 35 h of occupancy).
- At each unique thermostat override instance, the users pressed the thermostat keys at least twice (changing the temperature setpoint on average by 1.5°C). This underlines that occupants likely overcompensate their thermostat control actions upon discomfort.
- About one third of the thermostat overrides were either to decrease the setpoint during the heating season or to increase the setpoint during the cooling season. This indicates that the conservatism in fixed and prescribed temperature setpoints not only wastes energy but also frustrates some occupants. However, future work should utilize larger datasets to better understand the setpoint decrease behaviour during the heating season and the setpoint increase behaviour during the cooling season.

- It is observed that the indoor temperatures leading to a setpoint increase or a decrease during the heating season are on average 3°C colder than they were during the cooling season.
- Discrete-time Markov logistic regression models for setpoint increase and decrease were developed independently for the heating and cooling seasons. Results indicate that the indoor temperatures substantially affect the frequency of the keypress actions – by a factor of 3 or more. Future research should investigate the link between the timing/frequency of thermostat keypress actions and the user comfort in depth.
- Upon this analysis, a recursive approximation of the maximum likelihood method was formulated to dynamically develop logistic regression models for the thermostat use behaviour inside local building controllers. This recursive thermostat behaviour learning algorithm was verified against the batch logistic regression models developed from the same dataset. The algorithm is employed by the temperature control algorithm (see Chapter 5) in selecting the operating setpoints. The code was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices. The details of the implementation results are presented in Chapter 5.

For the development of the lighting and blinds-learning algorithm, the light switch and blinds use behaviours in ten private offices were analyzed with concurrent solar irradiance, ceiling illuminance, and occupancy data. The analyses of the data provided following insights with regard to the light switch and blinds use behaviour of the occupants and their adaptive visual comfort:

- Occupants consciously and actively undertake blinds closing actions upon glare and light switch on actions at insufficient daylight. However, they often fail to

turn off their lighting or open their blinds once the conditions leading to visual discomfort disappears.

- Although the ten monitored offices were identical in geometry, orientation, and construction characteristics, the indoor conditions leading to visual discomfort were very different. Therefore, It is challenging to develop daylight-integrated blinds and lighting automation with fixed and prescribed illuminance setpoints.
- It was revealed that the frequency at which the occupants close their blinds or turn on their lights can act as a proxy to their preferred visual comfort conditions. A discrete-time Markov logistic regression model that employs a ceiling-mounted photosensor's readings to predict the likelihood of a light switch on or a blinds closing action during occupied hours was found adequate to represent this phenomenon.
- Upon this analysis, a recursive algorithm was formulated to develop discrete-time Markov logistic regression models that predict the light switch-on and the blinds closing behaviours inside a building controller. The models were intended to be used in selecting adaptive illuminance setpoints for daylight-integrated lighting and/or blinds automation (when available). The algorithms were implemented in the controls laboratory and in five private offices for further testing and analyses. The details of the implementation results are presented in Chapter 6.

Now that we have algorithms to learn from occupant's presence and comfort patterns, the next chapter will look at developing an inverse modelling algorithm for the zone level temperature response.

Chapter 4

Learning from the temperature response

This chapter presents the inverse modelling of the temperature response inside individual offices. The purpose of this exercise is to estimate the warming/cooling capacity of the terminal HVAC units subject to environment and occupant-driven loads. This information is needed for calculating the setback-to-setpoint transition time in each zone prior to the occupants' arrival. The analyses presented in this section were conducted upon the data gathered in three private offices in an academic office building.

The chapter covers the development of a recursive self-adaptive building learning algorithm. The laboratory and field implementation of the algorithm for verification and further testing are presented in Chapter 5.

4.1 Literature review

Data-driven models representing the heat and air-mass transfer processes in buildings have been increasingly used in controls and diagnoses of HVAC equipment [134–138]. Generally, a physical model of the thermal zone – driven by real-time sensory data – is employed to make short-term performance predictions [136]. These performance

predictions can be used to achieve optimal temperature setpoint scheduling (e.g., optimal start/stop algorithms) and/or to detect and isolate faults (e.g., faulty sensors, envelope degradation) from unexpected building and HVAC behaviour [136, 138, 139]. Despite the availability of established modelling methods (e.g., thermal resistance-capacitance networks), the data-driven model development process leaves control designers with immense decisions pertaining to the selection of the model inputs, model complexity, and parameter identification methodologies.

In these data-driven models, the model inputs are the data gathered from building sensors, meters, and the feedback from the building actuators. The data-driven models use them as a proxy to represent the impact of the major components of heat transfer and generation in office buildings. The major components of heat transfer and generation in offices are envelope transmission, air infiltration, solar gains, occupant-induced loads (e.g., internal gains, plug-in appliance loads, electric lighting), and HVAC loads [140]. In the reviewed greybox models, the influence of envelope and infiltration losses on the indoor temperature is accounted for by taking indoor and outdoor temperature measurements [139, 141–144]. The influence of the terminal HVAC systems on the indoor temperature has been inferred either through their state (e.g., on/off) or through the supply flow rate (e.g., air pressure drop in a VAV unit) and temperature (e.g., discharge air temperature in VAV unit or temperature drop across a hydronic heating panel) [138, 139, 142, 145]. The influence of solar gains on the indoor temperature is estimated by employing outdoor pyranometers measuring the global horizontal irradiance [144] or the irradiance on the facade [139, 142], or through photodiodes measuring indoor illuminance [145]. Occupants' influence on the indoor temperature – casual gains due to presence, electric lighting, plug-in appliance loads – has been rarely represented in these models as direct inputs. This can

be attributed to the fact that these occupant-interacting components are not always integrated into the BAS.

A greybox model simplifies the heat transfer and generation problem in buildings by dividing the building fabric into discrete elements (i.e., nodes). Because heat transfer is a continuous phenomenon, different portions of an office space attain different indoor temperatures. Particularly, in cases where the thermal mass is separated from the airspace by an insulating layer (e.g., thick carpet, raised access floor), the model inputs can influence the indoor air temperature and building fabric very differently [146]. This favours more detailed models with many nodes. On the other hand, the sparsity of indoor temperature sensors in an office space makes the development of detailed data-driven models challenging. Often, the zone temperature is controlled based on a single indoor temperature sensor. When a model with many nodes is employed without taking measurements for most of these nodes, the parameter estimation problem becomes highly dependent on the estimation of the unknown node temperatures. In other words, measurements from a small number of sensors cannot be used to improve our understanding of a detailed model with many parameters and unmeasured states (e.g., wall or slab temperatures). In an effort to mitigate this issue, several researchers chose to compute a portion of the parameter values from their physical descriptions (e.g., wall insulation value) and to keep these parameters constant during operation [143, 147, 148]. However, because a building's thermophysical and occupancy characteristics can change seasonally and during the life of a building (e.g., changes in furnishings and envelope degradation or retrofit), even parameters derived from detailed physical descriptions can become unrepresentative in time [149, 150]. Also, computation of the model parameters upon physical descriptions is labour-intensive and it can require the building's as-built geometry, thermophysical properties of its materials, air leakage and distribution characteristics, and internal

heat gains and operational schedules. Furthermore, when the parameters of a thermal zone model are derived from detailed physical descriptions, the model becomes non-transferable to other zones without manual fine-tuning. Therefore, control-oriented models should autonomously identify their parameters in recursion [151].

Beyond model complexity, the choice of recursive parameter estimation methodologies plays an important role over the robustness of the control-oriented models. As it was the case for model input and complexity choices, the reviewed literature is vastly fragmented among three broad recursive parameter estimation methodologies: the recursive least squares filters [152, 153], Bayesian filters such as the Extended and Unscented Kalman Filters [16, 141, 145, 148, 153], and Monte Carlo implementation of the Bayesian Filters such as the Ensemble Kalman Filters [154, 155]. The recursive least squares filter merely attempts to minimize the misfit between the measurements and model predictions by assuming that the model can perfectly capture the physics of the problem. Bayesian filters acknowledge the fact that a greybox model cannot make perfect predictions even when the measurements are perfect, however they assume that this modelling error is normally distributed. With the additional computational complexity, sequential Monte Carlo methods (e.g., Particle Filters) can accommodate non-Gaussian modelling errors. Note that batch parameter estimation methods are not within the scope of this study.

In brief, performance data in buildings emerge sequentially and are not often stored permanently – meaning that a large dataset to train controls-oriented models are usually unavailable. In addition, data in buildings are non-stationary. The computational power available to train controls-oriented models inside existing commercial building controllers in real-time is typically very limited [155]. Lastly, the building controls networks often provide only limited sensing and metering capabilities, and inputs on the existing controllers may not be readily available to upgrade

the sensing and metering. The data-driven controls-oriented models reviewed in the literature were not developed in recognition of all of these limitations.

4.2 Objectives and scope

The objective of this part of the thesis is to develop an algorithm to filter out the warming/cooling rate capacity of the terminal HVAC equipment subject to environment and occupant-driven loads from the temperature response during operation. This information is needed in calculating the near-optimal setback-to-setpoint transition time in each zone prior to the occupants' arrival. To this end, twelve different greybox models of varying degrees of complexity were developed and then trained with two different recursive parameter estimation approaches: the Extended Kalman Filter (EKF) and the Particle Filter (PF). The predictive performance and appropriateness of the models were assessed based on the data gathered from three private offices in an academic building. One of the models and the recursive parameter estimation approaches were selected. Subsequently, they were formulated as a standard BACnet program object in scripting language GCL+ [93]. The algorithm is intended to be executed recursively. It is called to update the model parameters in real-time. The code was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices. The details of the implementation results are presented in Chapter 5.



Figure 4.1: Photos of the monitored offices.

4.3 Methodology

4.3.1 Monitored offices and data

Unlike the parts pertaining with the occupant behaviour and occupancy, the diversity aspects were not as pronounced in the temperature response of these spaces. In an effort to present the analyses results concisely, the data from only three of the private perimeter offices (see Table 1.1) were employed in this part of the thesis. Figure 4.1 presents a photo for each of the offices. Recall that each office was used by one full-time faculty member. Each office had a floor area of 15 m². The concrete floor slab was covered with a carpet, and the ceiling slab is separated from the room air with a plenum space and ceiling tiles. The offices were neighboured by other conditioned spaces (adjacent offices and the corridor).

As part of the building's control network, each office was equipped with a number of sensors: a ceiling-mounted photodiode-based light intensity sensor (spectral response range 350 to 1100 nm and 60° field of view), a PIR motion sensor (5 m




Figure 4.2: Position of the sensors and actuators in the monitored offices.

range and 100° horizontal and 80° vertical coverage), two $10\text{ k}\Omega$ thermistor-based temperature sensors measuring the indoor and the VAV terminal unit's discharge air temperatures, a pressure sensor for the VAV terminal unit's airflow control, and a current sensor for plug-load monitoring. Because the electric lighting and the radiant panel heaters were integrated into the BAS, their states were available as variables in the building's control network. The outdoor air temperature was measured at the air inlet of the AHU; and it was also available as a variable in the building's control network. The solar irradiance on the facade was measured by a pyranometer sensor of an in-situ weather station. Figure 4.2 presents the sensor locations in the monitored offices.

As summarized in Table 4.1, the overlapping portion of the data from these sensors were collected between March 15 and June 26 in room 1 and between May 13 and June 26 in rooms 2 and 3. Because a control-oriented model needs to make predictions in presence of uncertainties typical of commercial BAS instrumentation, the instrumentation was not upgraded or calibrated specifically for this study – i.e., the instrumentation was assumed as-is. Simply put, if the errors due to the instrumentation and the models were excessive, the model would fail to represent the temperature

response in the prediction time horizon anyways. Regardless, the photodiode, plug load meter, and temperature sensors were verified against other sensors measuring the same variable – provided as supplemental files at [156].

Figure 4.3 presents the indoor air temperature and environment-driven model inputs. During the 103-day monitoring period, the ambient conditions varied substantially. The outdoor temperature varied between -15 and 40°C and diurnal peak solar irradiance on the facade orientation varied from 50 to 850 W.m^{-2} . This allowed us to test the control-oriented models' performance under substantially different environmental loading conditions. Note that the ceiling-mounted photodiodes could read light intensity between 0 and 1076 lux , and they reached 1076 lux on sunny afternoons – meaning that the true illuminance was 1076 lux or higher on these instances. The building was in the heating mode until May 6. During the seasonal switchover to cooling, when the cooling was needed, the AHU supplied outdoor air in economizer mode. The rooms did not have an operable window which resulted in the high indoor temperatures shown in Figure 4.3.

Figure 4.4 presents the occupant-driven model inputs: the occupancy state (present or absent), electric lighting state (on or off), and the plug-in appliance loads. The occupancy time-series data were created from the motion sensor data by using the adaptive time delay algorithm presented in Chapter 2. Figure 4.5 presents the measurements taken from the terminal HVAC systems: radiant panel heater state u_7 (0 to 100%), the VAV terminal unit's discharge air temperature and pressure. Rooms 1 and 2 share a VAV unit, and thus they have the same VAV unit discharge air temperature (u_8) and pressure (u_9) readings.

Table 4.1: Characteristics of the data.

Data type	Sensor type	Acquisition period	Sampling frequency
Illuminance on ceiling	Photodiode	March 15 to June 26 in room 1 May 13 to June 26 in rooms 2 and 3	Time-series data in 15 min timesteps
Occupancy (present or absent)	PIR	March 15 to June 26 in room 1 May 13 to June 26 in rooms 2 and 3	Event-based data
Indoor temperature	Thermistor	March 15 to June 26 in room 1 May 13 to June 26 in rooms 2 and 3	Time-series data in 15 min timesteps
Outdoor temperature sensor	Thermistor	March 15 to June 26	Time-series data in 15 min timesteps
Solar irradiance on the exterior window	Pyranometer	March 15 to June 26	Time-series data in 15 min timesteps
VAV discharge air temperature	Thermistor	March 15 to June 26 in room 1 May 13 to June 26 in rooms 2 and 3	Time-series data in 15 min timesteps
VAV discharge air pressure	Pressure sensor	March 15 to June 26 in room 1 May 13 to June 26 in rooms 2 and 3	Time-series data in 15 min timesteps
Plug load	Current sensor	March 15 to June 26 in room 1 May 13 to June 26 in rooms 2 and 3	Time-series data in 60 min timesteps




Figure 4.3: Measured indoor (y) and outdoor temperatures (u_1), indoor illuminance (u_2), and vertical solar irradiance on the facade (u_3).

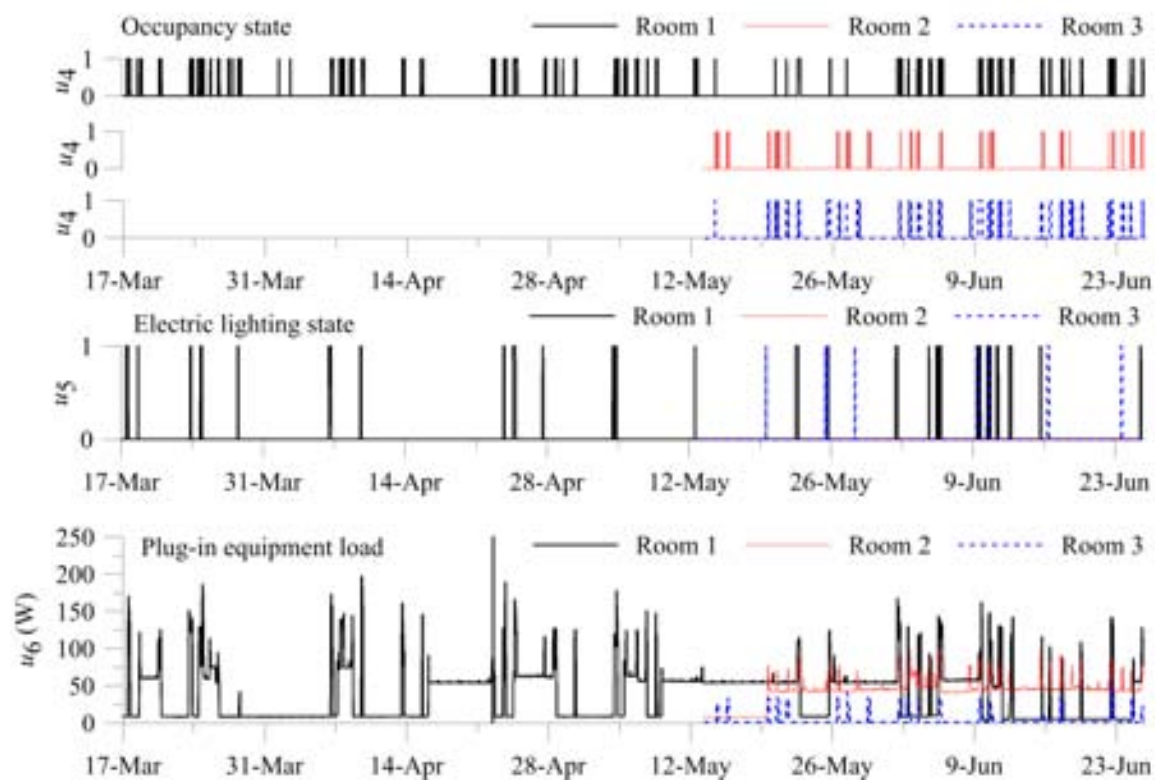


Figure 4.4: Monitored occupancy (u_4), electric lighting state (u_5), and plug-in equipment load (u_6).

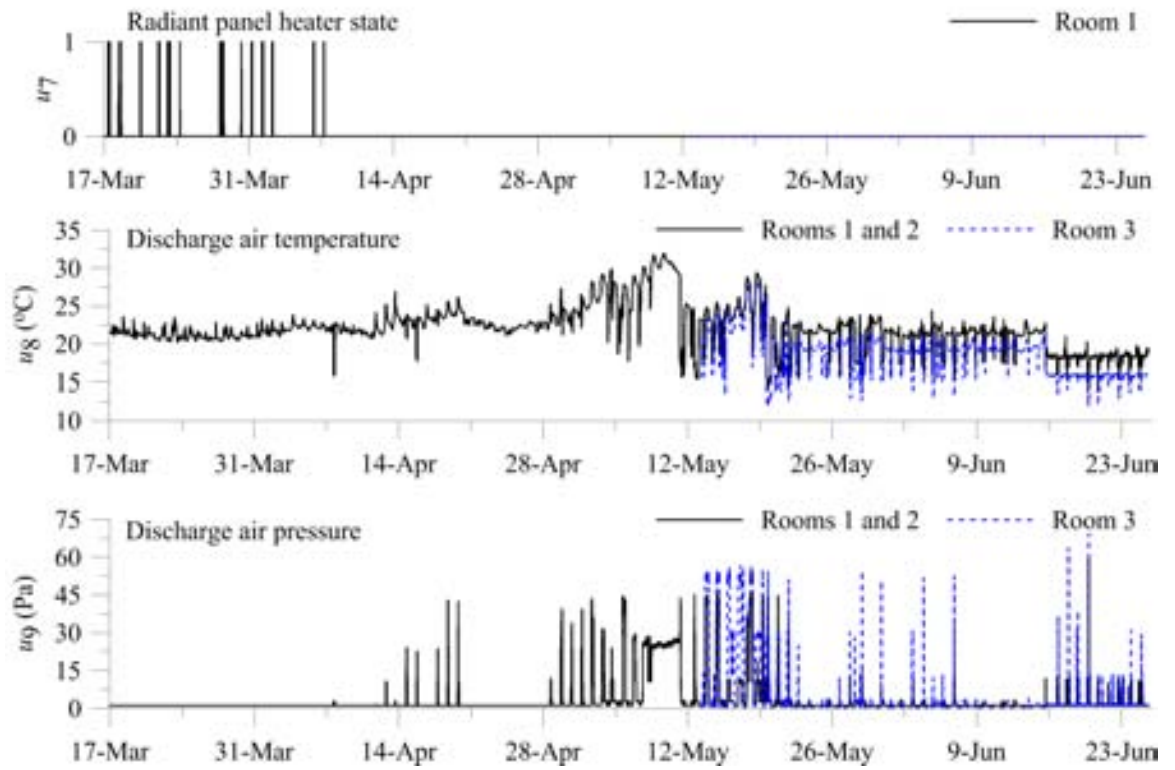


Figure 4.5: Measured radiant panel heater state (u_7), VAV terminal unit's discharge air temperature (u_8) and pressure (u_9).

Occupants were able to set their preferred indoor temperature setpoint through a wall mounted thermostat, and the rooms had different nighttime temperature setback schedules – as explained in Chapter 2. Furthermore, during the monitoring period, the mean blind occlusion rates (monitored via a timelapse camera) were 10, 35, and 5% in rooms 1, 2 and 3, respectively. These factors led to the variations in the transmitted solar radiation and daylight in different offices (see Figure 4.3). The variations in the interior design and furnishings between offices (see Figure 4.1) also played a nebulous role on the reflectance of the surfaces, and contributed to the diversity of the indoor daylight distribution. The mean plug-in equipment load during occupied periods was 95, 72 and 25 W in rooms 1, 2 and 3, respectively. During the unoccupied periods, it was 55, 43 and 14 W in rooms 1, 2 and 3, respectively. In each room, about 75% of the plug-in equipment electricity use took place during unoccupied periods because occupants tend to leave their plug-in equipment on as they leave their offices. Further information on the plug-in equipment load patterns in this building can be found elsewhere [157]. The occupants had diverse and flexible schedules. The mean occupied duration was 11, 8, and 13 h/week in rooms 1, 2 and 3, respectively. Lights were used only a fraction of the occupied duration (6 h/week in room 1 and 4 h/week in room 3). The electric lighting in room 2 was not used during the monitored time period. The diversity in occupants' behaviour and presence in the monitored offices allowed us to test the control-oriented models' performance under substantially different occupancy-driven loading conditions.

4.3.2 Control-oriented models

Figure 4.6 presents the first and second-order thermal network models of a perimeter office space. Clearly, many of the variables in these models are not directly observable through spatially limited building sensors (e.g., heat gains due to occupant-driven

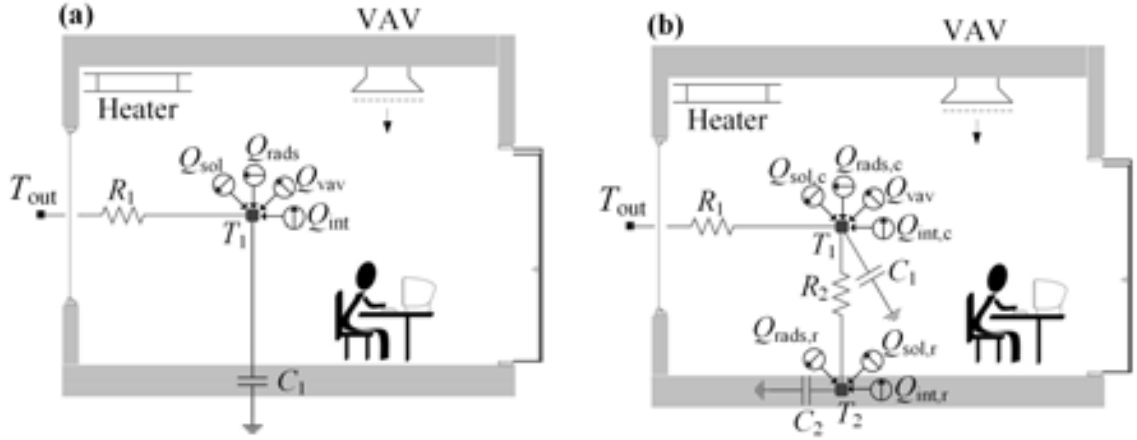


Figure 4.6: (a) First and (b) second-order lumped parameter thermal network model of a perimeter office space.

loads, air mass flow rate, solar gains). In order to use these models in controls applications, the parameters (R, C), states (T), and disturbances (Q) need to be mapped onto available sensory inputs shown in Figures 4.3 to 4.5.

In some cases, a building's control network may not provide any sensors to characterize a variable. For example, in absence of solar irradiation or illuminance data, solar gains cannot be characterized. On the other hand, in some cases, one may need to choose between multiple sensors (e.g., outdoor pyranometer or indoor photodiode) to characterize the same physical model input (e.g., solar gains). This results in different ways of developing controls-oriented models.

Based on the first and second-order thermal network models shown in Figure 4.6, twelve different controls-oriented greybox models at varying degrees of model complexities (number of inputs and parameters) were developed using the explicit finite difference formulation. In this modelling framework, the model inputs u and indoor temperature measurements y (see Figures 4.3 to 4.5), when multiplied by an array of parameters x , predict the rate of change in the indoor temperature \dot{T} .

Models 1 to 6 assume that the indoor temperature can be characterized as a single

temperature node – T_{in} . In particular, model 1 assumes that the indoor temperature can be characterized by the envelope losses and the terminal HVAC equipment states. It neglects the influence of the solar gains and the occupancy-driven heat gains/losses. The rate of change in temperature due to the envelope losses ($\frac{(T_{out}-T_1)}{R_1C_1}$ in Figure 4.6.a) was defined as $(u_1 - T_1)x_1$. The rate of change in temperature due to the radiant panel heaters ($\frac{Q_{rad}}{R_1C_1}$ in Figure 4.6.a) was defined as u_7x_7 . The rate of change in temperature due to the airflow from the VAV unit ($\frac{(T_{da}-T_1)\dot{m}C_{air}}{R_1C_1}$ in Figure 4.5.a) was defined as $(u_8 - T_1)\sqrt{u_9}x_8$. Note that $(u_8 - T_1)$ represents the difference between indoor and discharge air temperatures, and $\sqrt{u_9}$ represents a linear proportionality between the square-root of the pressure drop and the discharge air mass flow rate (i.e., flow rate is proportional to the square-root of pressure drop [8]). The model 1 in state-space form can be represented as follows:

$$\begin{aligned} \frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_7x_7 + (u_8 - T_1)\sqrt{u_9}x_8 + dw \\ y &= T_1 + v \end{aligned} \tag{4.1}$$

where the process noise w is a standard Wiener process and v is the measurement error. In a state-space representation, it is customary to choose a reasonable measurement error to avoid the process model overfit the observations – i.e., to slow down the parameter learning rate in recognition of the limited power of updating parameters with a single measurement. Thus, the measurements were assumed to be corrupted by an additive Gaussian white noise v with standard deviation 0.1°C . This 0.1 value was selected upon conducting a sensitivity analysis between 0 and 0.5°C .

Model 2 is identical to the model 1 with one exception that it attempts to characterize the influence of occupancy-driven heat gains on the rate of change in the indoor temperature. It assumes that the internal gains, plug-in appliance loads, and heat losses due to occupant's use of doors were subject to occupants' presence. Thus,

the rate of change in temperature due to occupancy-driven activities ($\frac{Q_{occ}}{R_1 C_1}$ in Figure 4.6.a) was defined as $u_4 x_4$ – i.e., occupancy u_4 multiplied by the nonphysical parameter (weight) x_4 . This implies that an occupant needs to be present in order to cause any disturbances in the temperature – neglecting the fact that occupant behaviours upon departure (e.g., leaving a desktop computer on) can influence the temperature response during the unoccupied periods. The functional form of the model 2 is as follows:

$$\begin{aligned} \frac{dT_1}{dt} &= (u_1 - T_1) x_1 + u_4 x_4 + u_7 x_7 + (u_8 - T_1) \sqrt{u_9} x_8 + dw \\ y &= T_1 + v \end{aligned} \quad (4.2)$$

Model 3 is identical to the model 2 except that it attempts to characterize the influence of solar gains on the rate of change in indoor temperature. It assumes that the temperature change rate due solar gains ($\frac{Q_{sol}}{R_1 C_1}$ in Figure 4.6.a) is linearly proportional to the photodiode sensor's readings u_2 . Thus, the rate of change in temperature due to the solar gains ($\frac{Q_{sol}}{R_1 C_1}$ in Figure 4.6.a) was defined as $u_2 x_2$ – i.e., photodiode sensor reading u_2 multiplied by the nonphysical parameter (weight) x_2 . The functional form of the model 3 is as follows:

$$\begin{aligned} \frac{dT_1}{dt} &= (u_1 - T_1) x_1 + u_2 x_2 + u_4 x_4 + u_7 x_7 + (u_8 - T_1) \sqrt{u_9} x_8 + dw \\ y &= T_1 + v \end{aligned} \quad (4.3)$$

Model 4 is identical to the model 3 with the exception that it attempts to characterize the influence of solar gains on the rate of change in indoor temperature ($\frac{Q_{sol}}{R_1 C_1}$ in Figure 4.6.a) by the pyranometer sensor's solar radiation readings on the facade u_3 – instead of the photodiode sensor's readings on the ceiling u_2 . The rate of change in the temperature due to the solar gains ($\frac{Q_{sol}}{R_1 C_1}$ in Figure 4.6.a) was defined as $u_3 x_3$ –

i.e., pyranometer sensor reading u_3 multiplied by the nonphysical parameter (weight) x_3 . The functional form of the model 4 is as follows:

$$\begin{aligned}\frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_3x_3 + u_4x_4 + u_7x_7 + (u_8 - x_1)\sqrt{u_9}x_8 + dw \\ y &= T_1 + v\end{aligned}\tag{4.4}$$

Model 5 is identical to model 4 except that it elaborates the role of occupancy-related activities over the rate of change in the indoor temperature ($\frac{Q_{occ}}{R_1C_1}$ in Figure 4.6.a) by introducing the electric lighting state u_5 as well as the occupancy u_4 . The rate of change in the temperature due to the lighting gains was modelled as u_5x_5 – i.e., electric lighting state u_5 multiplied by the nonphysical parameter (weight) x_5 . The functional form of the model 5 is as follows:

$$\begin{aligned}\frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_3x_3 + u_4x_4 + u_5x_5 + u_7x_7 + (u_8 - T_1)\sqrt{u_9}x_8 + dw \\ y &= T_1 + v\end{aligned}\tag{4.5}$$

Model 6 is identical to the model 5 except that it maps the occupancy-related activities onto the temperature response ($\frac{Q_{occ}}{R_1C_1}$ in Figure 4.6.a) in a more detailed way. The model 6 inputs the plug-in appliance loads u_6 as well as the states of occupancy u_4 and electric lighting u_5 to characterize the influence of occupancy-related activities over the rate of change in temperature. The rate of change in the temperature due to the plug-in appliance use was modelled as u_6x_6 – i.e., plug-in appliance load u_6 multiplied by the nonphysical parameter (weight) x_6 . The functional form of the model 6 is as follows:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_3x_3 + u_4x_4 + u_5x_5 + u_6x_6 + u_7x_7 + (u_8 - T_1)\sqrt{u_9}x_8 + dw \\
y &= T_1 + v
\end{aligned}
\tag{4.6}$$

Models 7 to 12 splits the problem domain into two temperature nodes: T_1 and T_2 . The node 1 (T_1) represents the portion of the fabric that can be described by the indoor air temperature measurements; and the node 2 (T_2) represents the portion of the fabric that does not follow the measured indoor temperature response. Note that measurements were only taken from node 1 (T_1).

The model inputs influence the two temperature nodes differently. This causes a temperature gradient between the two nodes – leading to a heat exchange between them. In the models 7 to 12, the rate of temperature change in node 1 due to this heat exchange ($(T_2 - T_1)/R_2C_1$ in Figure 4.6.b) was defined as $(T_2 - T_1)x_9$. The rate of temperature change in node 2 due to the heat exchange with node 1 ($(T_1 - T_2)/R_2C_2$ in Figure 4.6.b) was defined as $(T_1 - T_2)x_{15}$ – i.e., temperature difference between the two nodes ($T_2 - T_1$) multiplied by the nonphysical parameters x_9 and x_{15} . In recognition of the fact that the two temperature nodes can have different thermal capacitances, the parameters x_9 and x_{15} were defined as two different variables.

Because the solar gains, the heat emitted by an occupant, the heat gains due to the electric lighting and equipment, and the radiant panel heaters had a notable radiant component, they were assumed to influence the node 2 (T_2) as well as the first temperature node (T_1). The envelope and infiltration losses and the air-based cooling by the VAV unit were assumed to influence the temperature node 1 (T_1).

Model 7 has the same inputs as the model 1 with the exception of having two temperature nodes instead of one. The functional form of the model 7 is as follows:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_7x_7 + (u_8 - T_1)\sqrt{u_9}x_8 + (T_2 - T_1)x_9 + dw \\
\frac{dT_2}{dt} &= u_7x_{10} + (T_1 - T_2)x_{15} + dw \\
y &= T_1 + v
\end{aligned} \tag{4.7}$$

Model 8 has the same inputs as the model 2 except that it has two temperature nodes instead of one. The functional form of the model 8 is as follows:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_4x_4 + u_7x_7 + (u_8 - T_1)\sqrt{u_9}x_8 + (T_2 - T_1)x_9 + dw \\
\frac{dT_2}{dt} &= u_7x_{10} + u_4x_{11} + (T_1 - T_2)x_{15} + dw \\
y &= T_1 + v
\end{aligned} \tag{4.8}$$

Model 9 has the same inputs as the one-state model 3 with the following functional form:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1)x_1 + u_2x_2 + u_4x_4 + u_7x_7 + (u_8 - T_1)\sqrt{u_9}x_8 + (T_2 - T_1)x_9 + dw \\
\frac{dT_2}{dt} &= u_7x_{10} + u_4x_{11} + u_2x_{12} + (T_1 - T_2)x_{15} + dw \\
y &= T_1 + v
\end{aligned} \tag{4.9}$$

Model 10 has the same inputs as the model 4 except that it had two temperature nodes instead of one. It has the following functional form:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1) x_1 + u_3 x_3 + u_4 x_4 + u_7 x_7 + (u_8 - T_1) \sqrt{u_9} x_8 + (T_2 - T_1) x_9 + dw \\
\frac{dT_2}{dt} &= u_7 x_{10} + u_4 x_{11} + u_3 x_{12} + (T_1 - T_2) x_{15} + dw \\
y &= T_1 + v
\end{aligned} \tag{4.10}$$

Model 11 has identical inputs as the model 5 with the exception that it had two temperature nodes instead of one. It was defined as follows:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1) x_1 + u_3 x_3 + u_4 x_4 + u_5 x_5 + u_7 x_7 + (u_8 - T_1) \sqrt{u_9} x_8 + (T_2 - T_1) x_9 + dw \\
\frac{dT_2}{dt} &= u_7 x_{10} + u_4 x_{11} + u_3 x_{12} + u_5 x_{13} + (T_1 - T_2) x_{15} + dw \\
y &= T_1 + v
\end{aligned} \tag{4.11}$$

Model 12 has the same set of inputs as the model 6 except that it had two temperature nodes in lieu of one. It has the following functional form:

$$\begin{aligned}
\frac{dT_1}{dt} &= (u_1 - T_1) x_1 + u_3 x_3 + u_4 x_4 + u_5 x_5 + u_6 x_6 + u_7 x_7 + (u_8 - T_1) \sqrt{u_9} x_8 + (T_2 - T_1) x_9 + dw \\
\frac{dT_2}{dt} &= u_7 x_{10} + u_4 x_{11} + u_3 x_{12} + u_5 x_{13} + u_6 x_{14} + (T_1 - T_2) x_{15} + dw \\
y &= T_1 + v
\end{aligned} \tag{4.12}$$

Table 4.2 presents a summary of the parameters, their relationship with the models, and the effect of type of heat gains and losses on the temperature characterized with the parameters as a grid.

Recall that two simple thermal network models shown in Figure 4.6 represent the basis for the twelve models. The twelve models were developed by mapping

Table 4.2: Model parameters and inputs.

	Parameters															Inputs									
Models	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}	x_{12}	x_{13}	x_{14}	x_{15}	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	
First order models	1	✓					✓	✓									✓						✓	✓	✓
	2	✓		✓			✓	✓									✓			✓			✓	✓	✓
	3	✓		✓			✓	✓									✓			✓			✓	✓	✓
	4	✓	✓	✓			✓	✓									✓		✓	✓			✓	✓	✓
	5	✓	✓	✓	✓		✓	✓									✓		✓	✓	✓		✓	✓	✓
	6	✓		✓	✓	✓	✓	✓									✓		✓	✓	✓		✓	✓	✓
Second order models	7	✓					✓	✓	✓	✓					✓		✓						✓	✓	✓
	8	✓			✓		✓	✓	✓	✓	✓					✓	✓		✓				✓	✓	✓
	9	✓	✓		✓		✓	✓	✓	✓	✓	✓				✓	✓		✓				✓	✓	✓
	10	✓		✓			✓	✓	✓	✓	✓	✓				✓	✓		✓				✓	✓	✓
	11	✓		✓	✓		✓	✓	✓	✓	✓	✓	✓			✓	✓		✓		✓		✓	✓	✓
	12	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓		✓		✓		✓	✓	✓

Inputs:

- u_1 : outdoor temperature ($^{\circ}\text{C}$)
- u_2 : illuminance on the ceiling (lux)
- u_3 : vertical irradiance on the exterior window (W/m^2)
- u_4 : occupancy (present/absent)
- u_5 : electric lighting state (on/off)
- u_6 : plug load (W)
- u_7 : radiant panel heater state
- u_8 : VAV discharge air temperature ($^{\circ}\text{C}$)
- u_9 : VAV discharge air pressure (Pa)

different combinations of typical building sensors onto the simplified physics defined by the models shown in Figure 4.6. Although substantial simplifications were made during this procedure, the links to the physical problem were partly preserved. For example, the envelope and infiltration losses and the VAV unit's heat output were not applied onto the second temperature node T_2 – because of their negligible radiant components. Similarly, the non-linearities inherent in the relationship between the model inputs and the indoor temperature were also preserved. This procedure was sought in lieu of establishing mere input-output relationships between sensory inputs. Therefore, these models were categorized as greybox; yet it needs to be acknowledged that the models are much closer to being empirical than physical.

4.3.3 Recursive parameter estimation

At timestep $k-1$, the selected control-oriented model predicts the temperature \mathbf{T}_k at timestep k based on the parameters \mathbf{x}_{k-1} and inputs \mathbf{u}_{k-1} and \mathbf{T}_{k-1} . Note that model Eqns. 4.1 to 4.12 are stochastic, meaning that the model predictions for the temperature \mathbf{T}_k is a multidimensional likelihood distribution: $p(\mathbf{T}_k, \mathbf{x}_k | \mathbf{T}_{k-1}, \mathbf{u}_{k-1}, \mathbf{x}_{k-1})$. The letters $\mathbf{T}, \mathbf{x}, \mathbf{u}$ shown in bold font stand for the array of temperatures (T_1, T_2) , unknown parameters $(x_1 - x_{15})$, and model inputs $(u_1 - u_9)$ in Eqns. 4.1 to 4.12, respectively.

When the time elapses to k , an indoor temperature measurement y_k is taken. Almost certainly, the measurement will be different than the model prediction \mathbf{T}_k . The recursive parameter estimation problem attempts to update the parameters x by looking at this discrepancy between the model prediction \mathbf{T}_k and measurement y_k at each timestep.

The recursive Bayesian parameter estimation framework suggests that the updated likelihood distribution for the parameters \mathbf{x} after blending a single observation y_k

(i.e., posterior distribution) is proportional to the multiplication of the likelihood of observing y_k given the current set of parameters \mathbf{x} and the probability distribution of the parameters prior to observing y_k . Thus, the recursive Bayesian state and parameter update can be represented as follows:

$$p(\mathbf{T}_k, \mathbf{x}_k | y_{1:k}) \propto p(y_k | \mathbf{T}_{k-1}, \mathbf{x}_{k-1}, u_{k-1}) p(\mathbf{T}_{k-1}, \mathbf{x}_{k-1} | y_{1:k-1}) \quad (4.13)$$

If the distributions $p(\mathbf{T}_{k-1}, \mathbf{x}_{k-1} | y_{1:k-1})$ and $p(\mathbf{T}_k, \mathbf{x}_k | y_{1:k})$ can be approximated as conjugate distributions (in particular Gaussian), one can use the EKF to solve for the unknown parameters \mathbf{x}_k in recursion. However, if they are not conjugate distributions, one may need to use a computationally more expensive method: the sequential Monte Carlo (i.e., the PF). This may be particularly necessary in more complex models with many unknown parameters and states.

Recursive Parameter Estimation using the EKF:

The mean values for the temperature \mathbf{T}_k and the parameter array \mathbf{x}_k are predicted at time $k-1$ as follows:

$$\begin{bmatrix} \mathbf{T}_k \\ \mathbf{x}_k \end{bmatrix}^f = \begin{bmatrix} \mathbf{T}_{k-1} + f_i(\mathbf{T}_{k-1}, \mathbf{x}_{k-1}, u_{k-1}) \Delta t \\ \mathbf{x}_{k-1} \end{bmatrix}^u \quad (4.14)$$

where f_i represents the i^{th} model equation (see Eqns. 4.1 to 4.12). Note that the parameter array \mathbf{x}_k is augmented by the state array \mathbf{T}_k . Then, the covariance (\mathbf{P}_k) of the state and the parameters at time $k-1$ are forecasted as follows:

$$\begin{aligned}
\mathbf{P}_k^f = & \mathbf{J} \left(\begin{bmatrix} \mathbf{T}_{k-1} + f_i(\mathbf{T}_{k-1}, \mathbf{x}_{k-1}, \mathbf{u}_{k-1}) \Delta t \\ \mathbf{x}_{k-1} \end{bmatrix}^u \right) \mathbf{P}_{k-1}^u \\
& \cdot \mathbf{J} \left(\begin{bmatrix} \mathbf{T}_{k-1} + f_i(\mathbf{T}_{k-1}, \mathbf{x}_{k-1}, \mathbf{u}_{k-1}) \Delta t \\ \mathbf{x}_{k-1} \end{bmatrix}^u \right)^T + \mathbf{Q}
\end{aligned} \tag{4.15}$$

The states and the parameters forecasted using the Eqns. 4.14 and 4.15 are updated as follows:

$$\begin{aligned}
\begin{bmatrix} \mathbf{T}_k \\ \mathbf{x}_k \end{bmatrix}^u &= \begin{bmatrix} \mathbf{T}_k \\ \mathbf{x}_k \end{bmatrix}^f + \mathbf{P}_k \mathbf{H}^T (\mathbf{H} \mathbf{P}_k \mathbf{H}^T + \mathbf{R})^{-1} (\mathbf{y}_k - \mathbf{H} \mathbf{x}_k) \\
\mathbf{P}_k^u &= \left(\mathbf{I} - \mathbf{P}_k^f \mathbf{H}^T (\mathbf{H} \mathbf{P}_k^f \mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H} \right) \mathbf{P}_k^f
\end{aligned} \tag{4.16}$$

where \mathbf{R} is the covariance of measurement noise (square of v) in Eqns. 4.1 to 4.12 and \mathbf{H} equals $[1, 0, \dots, 0]_{1 \times \text{length}[\mathbf{T}_k, \mathbf{x}_k]}$. The superscripts f and u for the mean \mathbf{T}_k and \mathbf{x}_k , and the co-variance \mathbf{P}_k of the state-parameter array in Eqns. 4.14 to 4.16 stand for the forecast and the update stages of the filter, respectively. The states and the parameters updated in the previous timestep are used in forecasting the states and the parameters in the next timestep. As the Eqns. 4.14 to 4.16 are used with a greybox model defined in Eqns. 4.1 to 4.12 iteratively, the appropriate parameters \mathbf{x} will be learned in time. The mean of the initial parameters \mathbf{x} were set to physically reasonable values, and the initial covariance matrix \mathbf{P} was a diagonal matrix with entries reflecting our lack of confidence for the initial parameter estimates. The sensitivity of the diagonal covariance matrix assumption was studied and it was found that the influence of this initial assumption diminishes as more observations are processed in time. The initial diagonal entries were selected so that the distribution

of the initial parameter estimates cover a wide range of feasible values. With each measurement \mathbf{P} was updated and the influence of this initial guess diminished.

Recursive Parameter Estimation using the PF:

The PF recursive parameter estimation process is initialized by creating n number of parameter and state ensembles. These ensembles (i.e., particles) are randomly sampled from a reasonable flat prior $p(\mathbf{T}_{\mathbf{k}-1}, \mathbf{x}_{\mathbf{k}-1} | y_{1:k-1})$. Initially, each particle is assigned equal weights w_i where $\sum_{i=1}^n w_i = 1$. The subscript i stands for the index number of a particle.

The iterative parameter estimation process begins by employing one of the models shown in Eqns. 4.1 to 4.12. The temperature T_1 is predicted for each of the parameter and state ensembles. When the measurement y_k is taken at time k , the relative weight $p(y_k | \mathbf{T}_{\mathbf{k}-1}, \mathbf{x}_{\mathbf{k}-1}, \mathbf{u}_{\mathbf{k}-1})$ of each particle is updated as follows:

$$w_{i,k} \propto \exp \left(- \left((y_k - T_1^i) / v \right)^2 \right) w_{i,k-1} \quad (4.17)$$

The term v is the standard deviation of the assumed measurement error and it affects the learning rate. When it gets larger, the importance of the single measurement y_k becomes smaller and thus, the learning rate diminishes. Recall that v was taken as a Gaussian white noise with a standard deviation of 0.1°C . The particles that lead to temperature predictions closer to the measurement acquire larger likelihood weights than those yielding predictions that do not align with the measurement. Note that the relative particle weights form the kernel density; it is then normalized to satisfy $\sum_{i=1}^n w_i = 1$.

After a few iterations, all but one particle can have negligible weights. This phenomenon is also known as the particle degeneracy [158, 159] and it can be tackled by resampling. Note that the posterior distribution $p(\mathbf{T}_{\mathbf{k}}, \mathbf{x}_{\mathbf{k}} | y_{1:k})$ is not strictly a

generic unimodal distribution that one can directly sample from. Rather the inverse transform sampling over the cumulative distribution function of the particle weights were employed, when the number of effective particles (i.e., $\sum_{i=1}^n 1/(w_i)^2$) become less than one third of the total number of the particles [160, 161]. However, resampling introduces another challenge in the parameter estimation process. The diversity of the particles diminishes as the resampling favours particles with larger weights. This phenomenon is also known as the particle impoverishment [162]. Because a building's thermophysical and occupancy characteristics can change slowly in time, the parameter estimates acquired by the PF may become unrepresentative due to particle impoverishment. Here, the parameter ensembles were perturbed (i.e., regularized) at the resampling stage in order to ensure that parameter estimates remain time variant [163]. Furthermore, the number of particles plays an important role over the performance of the PF. As more particles are used, the posterior distribution can be better represented albeit at an increased computational cost. In this study, the PF was trained five times with each model by using 50, 100, 500, 1000, 5000, 10000, 20000 number of particles. 500 particles were found adequate to achieve repeatable model predictions for the single state models (models 1 to 6) and 1000 particles were found adequate for the two state models (models 7 to 12). The computational demand of using 500 to 1000 particles clearly restrains the PF's ability to be embedded inside building controllers – while it may be more logical to use it inside a central microcomputer serving multiple thermal zones. For more detailed information on the EKF and the PF, the readers may refer to Arulampalam et al. [163].

4.4 Predictive accuracy of the models

Each of the unknown parameters of the twelve models (see Eqns. 4.1 to 4.12) were estimated by using the EKF and PF methods. Note that the parameter estimation procedure was conducted recursively; meaning that at each timestep only one new set of sensory measurements (e.g., indoor and outdoor temperature) were made available for the EKF and the PF algorithms. The timestep size Δt was 15 min.

In order to evaluate the predictive accuracy of the models, after two weeks the recursive parameter learning process was interrupted once a day at a random time. The parameter estimates at that time-instance were used to make an offline prediction of the temperature response over a two day time-horizon. For example, parameters learned until the 20th day were used to predict the temperature response between the 20th and 22nd day. The predictive accuracies of the models were assessed based upon their ability to make these offline predictions. If a model underfits or overfits the data, its predictions will not be able to explain the characteristics of the data record – e.g., mismatch in the timing and the magnitude of the predicted and observed local extrema. Because we were interested in assessing the accuracy of the models, the model inputs during the prediction time horizon were assumed known. This isolated the modelling related errors from the errors emerging from model input forecasting process (e.g., the error in weather forecasts over the next two days). In fact, developing accurate forecasts to estimate occupancy and weather-driven model inputs over the prediction time horizon is another research challenge [164].

As illustrated in Figure 4.7, the prediction error is the difference between the observed and the predicted indoor temperature at each forecasting timestep over the two day offline prediction horizons. Figure 4.8 presents the distribution of the models' prediction errors in room 1. Over 90% of the predictions made by models 3 to 6 and




Figure 4.7: An illustration of the prediction error (by one of the models) at two random two day prediction time horizons in room 1.

9 to 12 with the EKF resulted in errors less than 1°C in 48 h. On the contrary, when the same models were trained with the PF about 30% of the prediction errors exceeded 1°C . The models 1, 2, 7, and 8 failed to provide accurate indoor temperature predictions. Note that these models do not use sensory inputs that characterize the solar gains (u_2, u_3 in Figure 4.3) and the occupancy-driven loads (u_5, u_6 in Figure 4.4).

With both EKF and PF the two-state models performed worse or marginally better than their single-state counterparts with identical inputs. Furthermore, the computations with the PF took 3 to 10 times longer than they were with the EKF. In recognition of the computational burden of conducting recursive parameter estimation with the PF method, it is hard to justify its usage in small thermal zone models. It




Figure 4.8: Histogram plots illustrating the distribution of the offline prediction errors over the two day time horizons (bin size 1°C) in room 1.

may find application in models representing large thermally-massive thermal zones – where the use of complex models may be necessary. Also, the PF method can become an alternative in cases where the initial parameter values can be guessed confidently – due to the availability of prior information regarding the geometry and material properties of the modelled zone. The EKF is an efficient recursive parameter estimation method and it appears to be adequate in zone level building controls applications.

Table 4.3 presents the MAE in predicting the indoor temperature over the two day time horizons with respect to the number of inputs and parameters in all three rooms. Results underline the fact that increasing the number of parameters without increasing the number of sensory inputs available for the model can increase – instead of decrease – the MAE. In other words, models with a large number of parameters and a small number of sensory inputs tend to overfit the observational data – which undermines the predictive accuracy. This is evident in comparisons between models 2 and 8 or models 3 and 9.

Based on the average of the MAE in the three rooms, model 10 appears to outperform all other models – including two models (models 11 and 12) with a larger number of sensory inputs and parameters. Model 6 – a model with 7 parameters and 8 sensory inputs – provided the second best predictive accuracy.

The simplest feasible models that can capture the timing and magnitude of the local extrema were models 3 and 4 trained with the EKF. They both had six inputs and five parameters. They could robustly predict the indoor temperature at less than $\pm 0.6^\circ\text{C}$ MAE over the two day prediction horizons in all three offices.

A comparison of the MAEs between models 4 and 5 or models 10 and 11 indicates that the inclusion of the electric lighting state as a model input marginally improved

Table 4.3: The MAE of the model predictions.

Models	Number of inputs	Number of parameters	MAE (°C)					
			Room 1		Room 2		Room 3	
			EKF	PF	EKF	PF	EKF	PF
Model 1	4	3	1.05	1.09	1.19	1.25	1.27	1.31
Model 2	5	4	1.06	1.11	1.02	1.07	1.05	1.10
Model 3	6	5	0.55	0.64	0.59	0.43	0.53	0.65
Model 4	6	5	0.57	0.65	0.56	0.46	0.51	0.60
Model 5	7	6	0.53	0.70	0.57	0.55	0.52	0.59
Model 6	8	7	0.45	0.65	0.54	0.68	0.53	0.53
Model 7	4	6	1.64	1.16	1.17	1.29	1.10	1.04
Model 8	5	8	1.63	1.18	1.10	1.08	1.07	1.08
Model 9	6	10	0.57	0.60	0.58	0.36	0.54	0.54
Model 10	6	10	0.40	0.61	0.65	0.53	0.35	0.56
Model 11	7	12	0.48	0.67	0.65	0.39	0.43	0.57
Model 12	8	14	0.42	0.68	0.51	0.57	0.58	0.68

or worsened the predictive accuracy of the models. In contrast, the inclusion of the plug-in appliance loads noticeably improved predictive accuracy particularly in rooms 1 and 2 (e.g., models 5 and 6). This can be attributed to the fact that a significant portion of the plug-in equipment loads can extend beyond the occupied periods [165]. In this study, about 75% of the plug-in equipment loads were measured during the unoccupied time periods, whereas occupancy sensor-based automated electric lighting could remain on only during the occupied periods. Consequently, occupancy (u_4) alone can be an adequate proxy for the time varying bias in the temperature response due to electric lighting use. However, due to the plug-in appliance loads during the unoccupied hours, it can be beneficial to incorporate the model input u_6 – if high resolution plug-load meters (i.e., current sensors) are part of the controls network. In other words, the model inputs need to be orthogonal; while lighting and occupancy violate this because occupancy-based automated-lighting is strictly dependent on the occupants' presence. However, it is worth noting that when the plug-in equipment

loads are negligible (e.g., laptop users), the addition of plug-in equipment loads as a sensory model input may detrimentally influence the predictive accuracy. This became evident when the data gathered in room 3 were employed in models 6 and 12 (see Table 4.3). Recall that the peak plug-in equipment load in room 3 was about 40 W (see Figure 4.4).

A comparison between models 3 and 4 and models 9 and 10 indicates that the solar irradiance on the facade – despite neglecting the effect of the day-to-day blind position variations – is as effective as the ceiling-mounted photodiode sensor's light intensity readings to act as a proxy to mimic the influence of solar gains over the rate of change in the indoor temperature. This can be partly explained with the fact that occupants rarely changed their blind positions in the monitored offices – i.e., on average two adjustments per month. This is in line with the earlier research suggesting that occupants infrequently interact with their blinds [39]. Because the blind positions tend not to change during the prediction time horizon (typically less than 48 h [138]), model predictions made upon solar radiation incident on the facade were often not affected by the change in the blind positions. Also, the internal roller blinds are usually better at rejecting daylight than the solar radiation [166]. In other words, when the roller blinds are closed, they can reject up to 95% of the daylight – while still permitting between 20 to 40% of the solar irradiance [166]. Furthermore, the solar heat gains – not just transmitted solar irradiation – are typically higher than these values with the blinds closed. In buildings with external shading systems – where the shading systems are more effective in controlling solar gains – the ceiling-mounted photodiode sensors may be a better proxy to characterize the influence of solar gains over the rate of change in the indoor temperature than the solar irradiance on the facade. A practical advantage of using the solar irradiance on the facade is that the data from a single pyranometer sensor are adequate for all the thermal zones in

that orientation. The pyranometer sensor can be located in a local rooftop weather station or the data can be extracted from a weather forecasting service (e.g., The North American Mesoscale Forecast System). In contrast, because a ceiling-mounted photodiode sensor's readings depend on numerous indoor peculiarities (e.g., furniture position and internal blind position), the correlation between the rate of change in the indoor temperature due to solar gains and photodiode sensor's readings becomes unique in each thermal zone. Consequently, the sensors need to be installed in each thermal zone. The advantage of the indoor photodiode sensors in contrast to the ambient pyranometer sensors is that they can also be used in daylight-integrated lighting and blinds automation.

4.5 Residual analysis and parameter estimates

Residuals are the differences between the timestep ahead model predictions and the measured temperature. Recall that the timestep size was 15 min in this study. A residual analysis provides valuable insight to determine a model's appropriateness to a dataset. An appropriate model's residuals should roughly be a zero mean and constant variance Gaussian white-noise under all loading conditions (e.g., when it is cold or hot outside and when it is sunny or cloudy). If a model underfits or overfits the observational data, it is likely to see that the residuals differ significantly from one portion of the solution space to another. To this end, the scatter plots shown in Figure 4.9 were built. They illustrate the models' prediction residuals in the three rooms at different loading conditions – when they were trained with the EKF. The solid lines in the figure represent the least-squares line fit through these scatter plots. Results indicate that the residuals of models 1, 2, 7, and 8 change substantially with respect to the solar irradiation intensity on the facade and with respect to the outdoor air

temperature. This variation is significantly more pronounced when the models were trained with the PF instead of the EKF (see Figure 4.10). For all other models, the characteristics of the residuals remained relatively unaffected with the change in the solar irradiation intensity on the facade and with the change in the outdoor air temperature.

The models were designed to predict the temperature in the next timestep by using the sensory model inputs in the current timestep. If this assumption holds, the model prediction residuals at each timestep should be uncorrelated to the earlier residuals (i.e., residual independence). For example, if a positive prediction residual is likely to be followed by another positive prediction residual or a negative residual tends to be followed by another negative residual, the residuals tend to accumulate and the prediction error propagates over the prediction time horizon. The correlation of a model's prediction residuals with itself at different points in time can be best seen by an autocorrelation plot. The faster the autocorrelation converges to zero, the better the model will satisfy the residual independence assumption. If the error follows diurnal loading cycles (e.g., solar radiation, outdoor air temperature), it will exhibit periodicity and deviate to non-zero values. This type of behaviour is evident in the residuals of the models 1 and 2 trained with the EKF and PF (see Figure 4.11.a and c). Similarly, the two-state counterparts of the models 1 and 2 (i.e., models 7 and 8) failed to fulfill the residual independence criteria (see Figure 4.11.b and d). The residuals of the other models' predictions (models 3 to 6 and 9 to 12) trained with the EKF did not exhibit such periodicity. However, the residuals of the models' predictions tend to remain correlated much longer, when the models were trained by the PF rather than the EKF method.

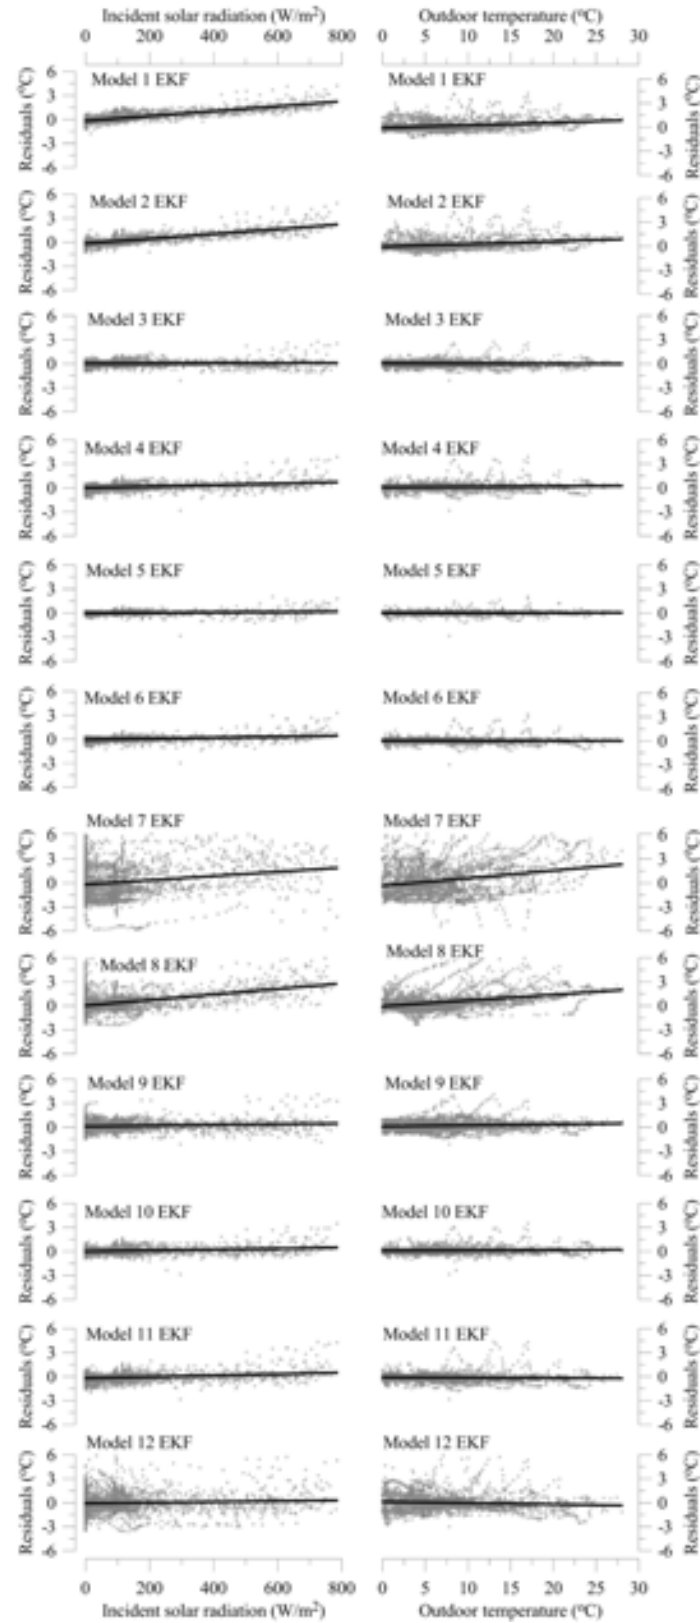


Figure 4.9: Scatter plots illustrating the models' prediction residuals as a function of the incident solar radiation and the outdoor temperature, when they were trained with the EKF. The solid lines represent the least squares line fit through these scatter plots.

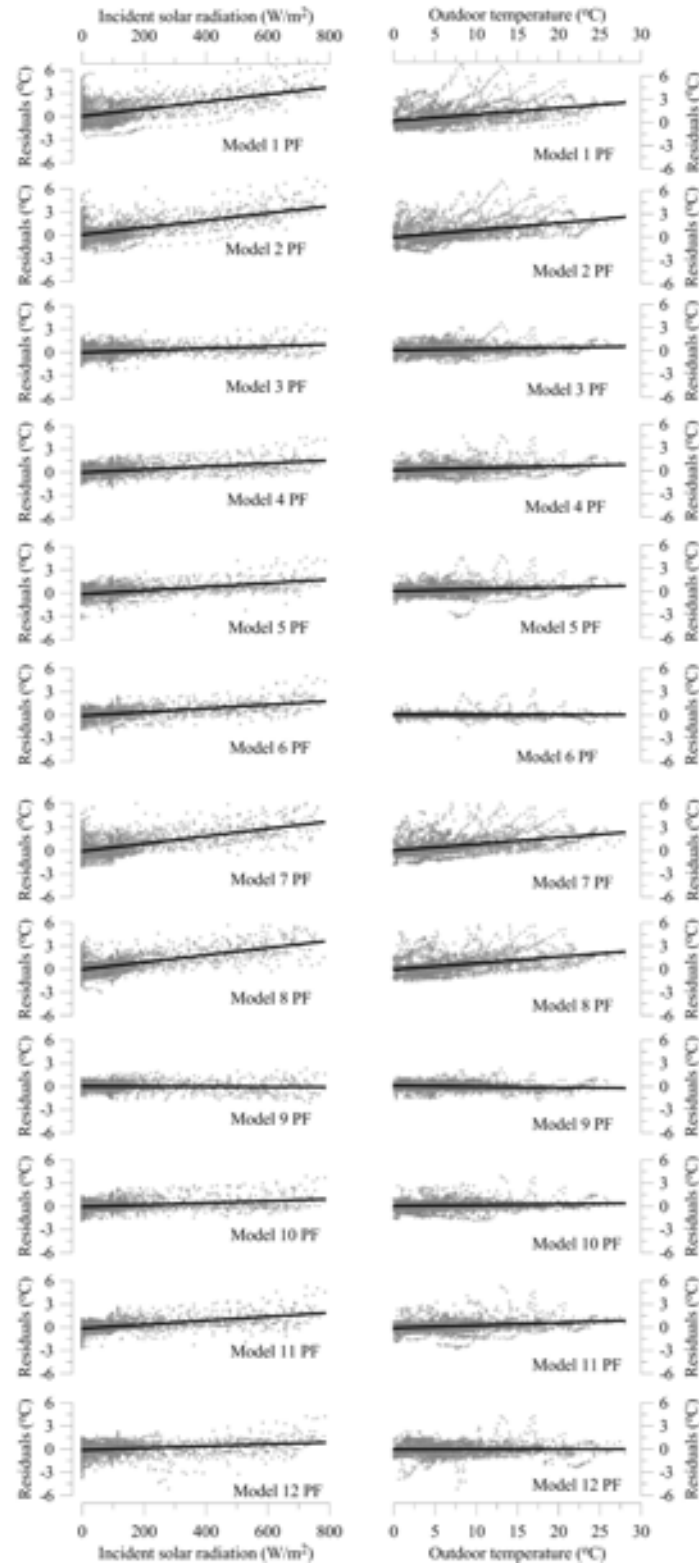


Figure 4.10: Scatter plots illustrating the models' prediction residuals as a function of the incident solar radiation and the outdoor temperature, when they were trained with the PF. The solid lines represent the least squares line fit through these scatter plots.

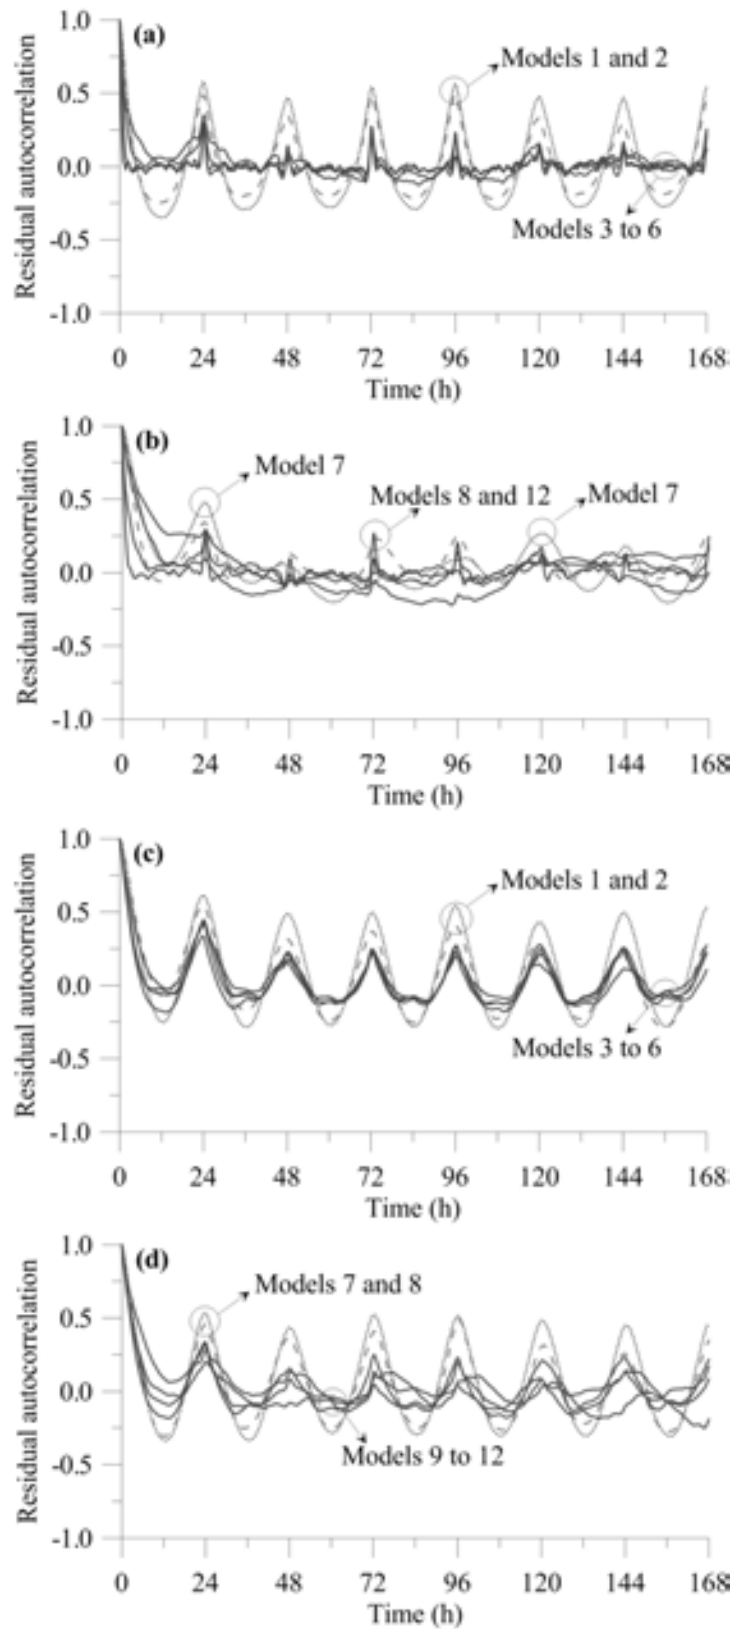


Figure 4.11: Autocorrelation plots for the models' prediction residuals. (a) The single-state models trained by the EKF, (b) the two-state models trained by the EKF, (c) the single-state models trained by the PF, and (d) the two-state models trained by the PF.

Recall that the models trained with the PF performed worse than the models trained with the EKF with respect to a number of performance metrics (e.g., offline prediction accuracy, model whiteness and independence). This can be explained with the particle impoverishment problem associated with frequent resampling in sequential combined state and parameter estimation problem [167]. As it is expected from a self-adaptive control-oriented model, the prior knowledge of the parameters before the model training was very limited. Herein, crude – yet identical – initial values for the EKF and PF were provided. For example, the rate of temperature change due to a constant 1000 W solar irradiation on the facade was initially defined as uniformly distributed between 0 and 4°C/h (i.e., initial value of the parameter x_7). With the importance resampling, the particles (i.e., parameter ensembles) with higher likelihood weights were replaced the others, and most of the particles sampled from the uninformative prior became unrepresentative during the model training. Due to the time-varying nature of the parameters, this has led to suboptimal parameter estimates. Other forms of the PF such as the fixed-lag PF could be pursued to employ the sequential parameter estimation [167], however it was found that the EKF could reach accurate parameter estimates robustly and efficiently for control-oriented models representing heat and air-mass transfer processes in perimeter office spaces.

Table 4.4 tabulates the R_1C_1 values in Figure 4.6 (reciprocal of the model parameter x_1) estimated by the models on June 15. The R_1C_1 values estimated for models 1, 2, 7 and 8 were unrealistically large. This is because they fail to characterize the daily temperature swings in absence of the inputs characterizing the solar gains. In lieu of predicting daily indoor temperature response, the models rather attempted to characterize the mean daily temperature trends to minimize their misfit to the data

Table 4.4: Reciprocals of the model parameters x_1 .

Models	Room 1 R_1C_1 (h)		Room 2 R_1C_1 (h)		Room 3 R_1C_1 (h)	
	EKF	PF	EKF	PF	EKF	PF
1	374	597	358	538	335	537
2	228	172	204	202	189	200
3	96	101	90	114	88	131
4	92	111	92	93	95	97
5	106	109	102	109	98	121
6	81	85	78	75	73	77
7	163	121	171	182	159	163
8	323	164	286	242	298	224
9	106	94	112	106	94	100
10	102	76	95	79	85	83
11	77	75	68	46	45	68
12	96	85	89	79	83	71

– which resulted in these unrealistically large R_1C_1 values. In addition, R_1C_1 estimates for a given model type were in agreement with PF and EKF methods. R_1C_1 values estimated using the data gathered from different rooms with identical geometry and construction characteristics were also similar. Results shown in Figure 4.12 indicate that when the EKF was employed for model 3 for each office independently, the parameter estimates pertaining to the envelope characteristics x_1 , solar gains x_2 , radiant panel heaters x_7 , and the VAV units x_8 converged to physically meaningful, stable, and consistent values. However, the parameter that estimates the influence of the occupancy-driven activities over the indoor temperature x_3 converged to different values in different offices. This can be attributed to the fact that the heat gains due to occupant-driven activities and the tendency to keep office door open during the occupied periods vary significantly among different occupants.

Based on the analyses of the predictive accuracy and the model appropriateness, model 3 with the EKF method was selected as the simplest feasible model achieving




Figure 4.12: The evolution of the model 4's parameters in time when it is trained by the EKF.

an acceptable level of accuracy for executing short-term control decisions in zone level HVAC equipment. Subsequently, the model's algorithm was formulated as a standard BACnet program object – in scripting language GCL+ [93]. The algorithm is intended to identify the unknown parameters of model 3 and update them in real-time. The algorithm was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices. The details of the implementation results are presented in Chapter 5.

4.6 Summary

Twelve different greybox models that predict the temperature response of a perimeter office space were developed at varying numbers of model inputs and parameters. Two different recursive parameter estimation methodologies were employed for each of the twelve models: the Extended Kalman Filter and the Particle Filter. A systematic way of selecting one of these models and recursive parameter estimation methodologies were introduced by using the sensory data collected from three perimeter office spaces.

It was found that increasing the number of parameters without increasing the number of sensory inputs available for a model can decrease – instead of increase – the models' predictive accuracy. However, increasing the number of parameters and the sensory model inputs characterizing different modes of heat transfer in tandem may improve the predictive accuracy. For example, the model with 10 parameters and 6 inputs, when trained with the EKF, outperformed all other models. Also, the model with 7 parameters and 8 inputs, when trained with the EKF, made more accurate predictions than more complex models with 12 and 14 parameters.

The inclusion of the electric lighting state as a model input only marginally influenced the predictive accuracy. In contrast, the inclusion of the plug-in appliance loads

improved predictive accuracy noticeably in two of the rooms. This was explained with the fact that a significant portion of the plug-in equipment load occurs during the unoccupied hours, whereas the occupancy-based automated electric lighting was only used during the occupied hours. As a result, the model input occupancy state alone can be an adequate proxy for the time varying bias in the temperature response due to the electric lighting use. However, due to the plug-in equipment loads during the unoccupied hours, it can be beneficial to incorporate the model input plug loads – when plug-load meters are present.

The solar irradiance on the facade – despite neglecting the effect of the day-to-day blind position variations – was found as effective as the ceiling-mounted photodiode sensor's light intensity readings to act as a proxy to capture the influence of solar gains over the rate of change in the indoor temperature.

The models trained with the PF performed worse than the models trained with the EKF. This was explained with the particle impoverishment problem associated with frequent resampling in sequential combined state and parameter estimation problem. The EKF was found to be an efficient recursive parameter estimation method adequate in zone level building controls applications. However, in zones with large exposed thermal mass (unlike the offices studied in this thesis) where more complex models are needed, the PF may still become more appropriate.

The simplest feasible models that can capture the timing and magnitude of the local extrema were the models with 5 parameters and 6 sensory inputs (i.e., indoor and outdoor temperature sensors, an indoor light intensity sensor or a pyranometer, a motion sensor, a radiant panel heater's state, a VAV unit's discharge temperature and pressure) – i.e., models 3 and 4. They could robustly predict the indoor temperature at less than 0.6°C MAE over the two day prediction horizons in all three offices.

Model 3 was selected and the model's algorithm was formulated as a standard

BACnet program object – in scripting language GCL+ [93]. The algorithm was intended to identify the unknown parameters of model 3 and update them in real-time. The algorithm was implemented and tested inside the laboratory's controls network and inside four local controllers serving eight private offices.

Now that the recursive algorithms to learn from occupants' adaptive behaviours and presence, and an inverse model to recursively learn from the temperature response have been developed, the methods are applied to the offices to improve operations in Chapters 5 and 6.

Chapter 5

Adaptive control of the indoor temperature

This chapter presents an adaptive indoor temperature control algorithm. It employs the recursively derived information from the learning algorithms presented in Chapters 2 to 4. Laboratory and field implementation results and comfort implications were presented. Energy-savings potential was investigated through a simulation-based study.

5.1 Literature review

A building's temperature control algorithm needs to leverage three types of information: (a) building-related constraints (HVAC equipment capacity, thermal capacitance, and envelope's thermal resistance), (b) occupant-related constraints (occupants' schedule and temperature preferences), and (c) disturbances (casual and solar gains, outdoor temperature). Traditionally, they are heuristically considered by the controls technicians while programming the temperature setpoints and schedules

inside a reactive controller (on-off, proportional, proportional-integral, proportional-integral-derivative controllers [168]). On the other hand, a model-based predictive controller (MPC) renders the potential to blend this information systematically [134].

An MPC for temperature control employs a model of the thermal zone to predict the temperature response over a time horizon subject to occupant and environment-driven loads (i.e., disturbances) and HVAC loads [169, 170]. An optimization algorithm is trained to identify the optimal control sequence minimizing a particular cost function. The cost function can be formulated to minimize the total heating/cooling loads or the utility costs over the time horizon (i.e., natural gas or electricity) [138]. The constraints of the optimization problem are typically the setpoint schedules and the HVAC equipment's heating/cooling capacity.

Researchers have extensively studied and demonstrated the viability of different optimization algorithms: (1) linear programming [169], (2) dynamic programming [171], (3) quadratic programming [172], (4) particle swarm optimization [173], (5) evolutionary algorithm [174], (6) firefly algorithm [175], and (7) Newton's method of optimization [152, 154]. The researchers have employed these methods inside modern computers with established statistical tools [176]. However, their implementation inside commercial building controllers – with limited computational and analytical capabilities – has remained an open research challenge.

An important variable in MPC design is the length of the prediction time horizon. Regardless of the model accuracy, the disturbances – i.e., inputs of the model – over the prediction time horizon need to be predicted. In the reviewed literature, the prediction horizons for temperature control range from 1 to 48 hours [138, 152]. Depending on the length of the prediction time horizon, different approaches of disturbance forecasting can be employed. For short prediction horizons (e.g., less than 3 hours), persistence-based forecasts can be appropriate [152]. In a persistence-based

forecast, the optimization algorithm plans for the next control decision by assuming that the disturbances will remain identical over the prediction horizon [177]. A common alternative to this is to infuse external forecasts inside the MPC [134, 178].

The first limitation of this approach is that a reliable source of external forecasts for environmental disturbances may not be available. For example, Environment Canada [179] forecasts cloud cover ratio, but not the solar irradiance. Instead of importing environmental forecasts (e.g., from Environment Canada API), the disturbance predictions can be created internally using established statistical modelling methods (e.g., moving average models, neural networks) [178]. For example, if a controls-oriented model employed by a zone temperature MPC inputs a ceiling-mounted photodiode sensor's readings as a proxy to characterize the solar gains such as Gunay et al. [16], the next a few hours' readings can be estimated by looking at the historical trends of this sensor.

The second limitation of infusing external forecasts is the lack of external forecasts for occupant-driven disturbances. Even if the external forecasts are perfect, an MPC for zone level indoor temperature control will make suboptimal decisions when an occupant changes the thermostat setpoints, window blind or operable window position, lighting or plug-in equipment state over the prediction time horizon [152]. In such cases, it is necessary to limit the prediction time horizon length to a time period whereby the occupant-driven disturbances are unlikely to change [152]. However, the reviewed literature almost consistently overlooked the influence of occupant-driven disturbances over the performance of MPC for indoor temperature control – with two exceptions [180, 181].

5.2 Objectives and scope

The objective of this chapter is to develop an adaptive temperature control algorithm. It is intended to use the information recursively derived by the learning algorithms presented in Chapters 2 to 4 to make control decisions for zone level HVAC equipment that improve the energy efficiency and maintain user comfort. Similar to the other algorithms presented in this thesis, it is intended to be embedded inside a typical commercial building controller. This chapter presents the structure of this temperature control algorithm and its relationship with the learning algorithms (see Chapter 2 to 4). The laboratory and field implementation of the control algorithm in tandem with the learning algorithms are presented. Energy and comfort related implications of the control algorithms were analyzed through the implementation results. Alternative control scenarios (and their energy and comfort-related implications) were studied through EnergyPlus simulations.

5.3 Adaptive temperature control algorithm

The temperature control algorithm inputs the outputs of the occupancy detection and learning, thermostat learning, and building learning algorithms. Recall that the occupancy detection algorithm (Chapter 2) outputs whether or not the zone is occupied. For a given zone, the occupancy-learning algorithm (Chapter 2) outputs four parameters: (1) the earliest expected arrival time, (2) the latest expected arrival time, (3) the latest expected departure time, and (4) the longest expected duration of intermediate absence. The thermostat learning algorithm (Chapter 3) outputs two parameters: (1) the lowest expected temperature to decrease the setpoint and (2) the highest expected temperature to increase the setpoint. The building learning algorithm (Chapter 4) outputs a data-driven model with five parameters that can

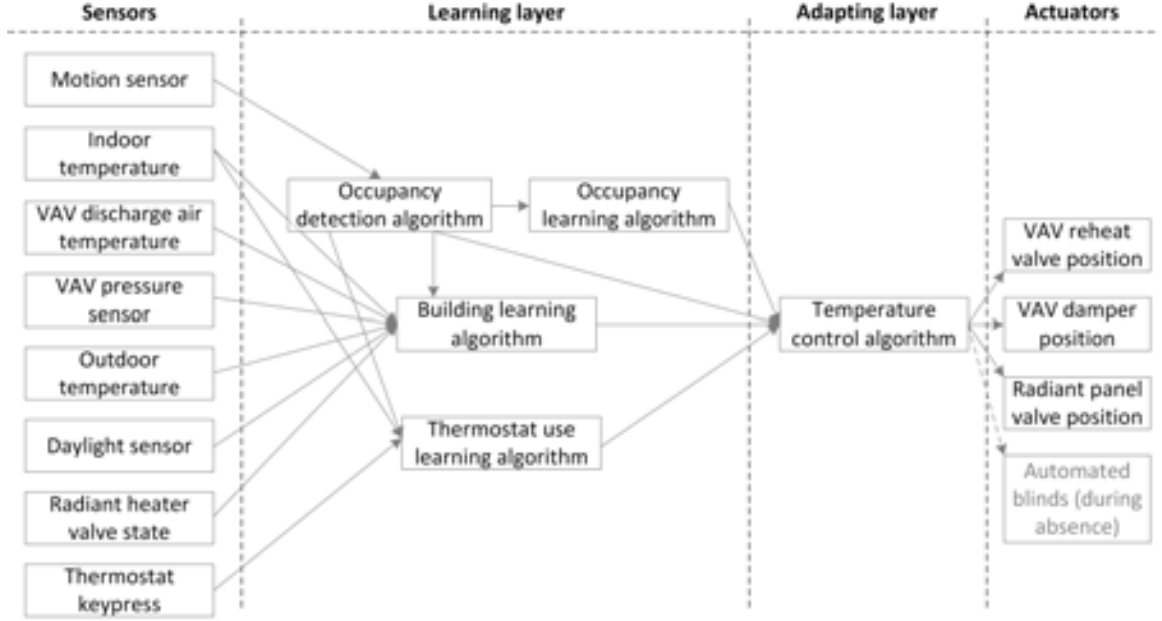


Figure 5.1: The temperature control algorithm and its relationship with the learning algorithms.

predict the temperature response of the zone. Figure 5.1 illustrates the relationship of the temperature control algorithm with these learning algorithms.

The temperature control algorithm employs a discrete optimization algorithm to minimize following quadratic cost function:

$$\text{minimize } \mathbf{J} = (\mathbf{T}_{\text{sp}} - \mathbf{T})^T (\mathbf{T}_{\text{sp}} - \mathbf{T}) + \Delta \mathbf{U}^T \mathbf{R} \Delta \mathbf{U} \quad (5.1)$$

where \mathbf{T}_{sp} ($^{\circ}\text{C}$) is the array of temperature setpoint information and \mathbf{T} ($^{\circ}\text{C}$) is the temperature predicted by the data-driven model presented in Chapter 4 over the prediction time horizon. $\Delta \mathbf{U}$ ($^{\circ}\text{C}$) is an array of the temperature change due to the heating/cooling loads. \mathbf{R} is a diagonal matrix carrying penalizing factors for different time-instances within a prediction time horizon. For the case that all diagonals in \mathbf{R} equal zero, this can be interpreted as the situation where we would not want to pay

any attention to how large the ΔU will be. For simplicity, \mathbf{R} is taken as an identity matrix in this study. The optimization problem is constrained by the capacity of the terminal heating and cooling systems to change the indoor temperature.

In the building where the algorithms were implemented, the central heating plant was operational between October 5 and May 6 (heating season). The cooling season was between May 6 and October 5. During the heating season, ΔU is constrained as follows:

$$\text{subject to } \Delta U \leq (x_7 + \sqrt{p_{da,sp}} (T_{da,sp} - T) x_8) \Delta t \quad (5.2)$$

where $p_{da,sp}$ (Pa) is the differential pressure setpoint of the VAV unit when the damper is at open and $T_{da,sp}$ ($^{\circ}\text{C}$) is the discharge air temperature setpoint of the VAV unit when the reheat coil valve is open. These variables are needed for the control of the VAV units' damper and the supply air temperature, and thus they are readily available in the VAV controller database. They were 30°C and 70 Pa in the building where the algorithms were implemented. Recall that x_7 and x_8 are the data-driven model's parameters (Chapter 4). The rate of temperature change in time, when the radiant panel is turned on, is x_7 . The parameter x_8 accounts for the effect of the VAV unit on the indoor temperature – when multiplied by the supply air temperature and the square-root of the supply air pressure. During the cooling season, ΔU is constrained as follows:

$$\text{subject to } \Delta U \leq (\sqrt{p_{da,sp}} (T_{da,sp} - T) x_8) \Delta t \quad (5.3)$$

In the cooling season, the $p_{da,sp}$ (Pa) (when the damper is at its open position) was still 70 Pa, but the $T_{da,sp}$ ($^{\circ}\text{C}$) was 13°C . Note that the humidity control is absent in the methodologies developed in this thesis. This was because the humidity control logic in a typical AHU VAV system is to maintain the return air relative humidity

between 30 and 60% at the air-handling unit – not at the VAV level. The zone level equipment is often not programmed to control humidity and temperature concurrently – partly because many of the existing buildings do not have humidity sensors in individual thermal zones. This may become problematic if a portion of the building disproportionately generates too much humidity than the rest. As an extension of the current work, future research should look into modifying the algorithms for the control of humidity and temperature in tandem.

The temperature in the prediction time horizon (\mathbf{T}) can be represented as a function of the current indoor temperature (T_k) as follows:

$$\begin{aligned}
 T_{k+1} &= \underbrace{(1 - x_1 \Delta t)}_A T_k + \underbrace{(x_1 u_1^k + x_2 u_2^k + x_4 u_4^k)}_{B_k} \Delta t + \Delta U_k \\
 T_{k+2} &= \underbrace{(1 - x_1 \Delta t)}_A (AT_k + B_k + \Delta U_k) + \underbrace{(x_1 u_1^{k+1} + x_2 u_2^{k+1} + x_4 u_4^{k+1})}_{B_{k+1}} + \Delta U_{k+1} \\
 T_{k+3} &= A(A(AT_k + AB_k + \Delta U_k) + B_{k+1} + \Delta U_{k+1}) + B_{k+2} + \Delta U_{k+2} \\
 &\vdots \\
 T_{k+n} &= A^n T_k + (A^{n-1} B_k + A^{n-2} B_{k+1} + \dots + A^0 B_{k+n-1}) + (A^{n-1} U_k + A^{n-1} U_{k-1} + \dots + A^0 U_{k+n-1})
 \end{aligned} \tag{5.4}$$

where u and x are the inputs and the parameters of the data-driven model developed in Chapter 4. Recall that u_1 ($^\circ\text{C}$) is the outdoor air temperature, u_2 (lux) is the ceiling light intensity, and u_4 is the binary occupancy indicator (present or absent). Note that the cost function J is differentiable with respect to the control output array ΔU and yields the following optimal control output array:

$$\Delta U = \underbrace{(\varphi^T \varphi + R)}_H^{-1} \varphi^T (\mathbf{T}_{\text{sp}} - \mathbf{T}_f) \tag{5.5}$$

where \mathbf{T}_f ($^\circ\text{C}$) is the free running temperature in the prediction time horizon subject to environment and occupant-driven loads and φ is the compact form of the following

matrix:

$$\varphi_{n,n} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ (1 - x_1 \Delta t) & 1 & 0 & 0 & 0 & 0 \\ (1 - x_1 \Delta t)^2 & (1 - x_1 \Delta t) & 1 & 0 & 0 & 0 \\ (1 - x_1 \Delta t)^3 & (1 - x_1 \Delta t)^2 & (1 - x_1 \Delta t)^1 & 1 & 0 & 0 \\ \vdots & & & & 1 & 0 \\ (1 - x_1 \Delta t)^{n-1} & (1 - x_1 \Delta t)^{n-2} & \dots & & & 1 \end{bmatrix} \quad (5.6)$$

If any of the ΔU values in control input array ΔU exceeds the constraints related with the capacity of the HVAC equipment, the portion of the control input ΔU that cannot be met by the equipment at that timestep is redistributed to the earlier timesteps in the prediction horizon. This condition typically happens during the setback-to-setpoint transition periods – where preheating or precooling can be necessary. A more detailed description of this procedure can be found in Wang [182]. The ΔU values were employed inside a P control loop with a proportional band of 0.5°C – i.e., if ΔU equals 0.4, the radiant panel and reheat coil valves and VAV damper are 80% open; they are 100% open at ΔU equals ΔU . Lastly, if there are BAS-integrated motorized blinds and the occupants are not present, the blinds are set open when there is daylight and ΔU is positive – or set closed when ΔU is negative.

Two important considerations for the controller are the timestep size (Δt) used in discretizing the prediction time horizon and the length of the prediction time horizon. Should the timesteps get smaller or the prediction time horizon gets longer, the computational burden of undertaking these matrix operations will exacerbate. Furthermore, as shown in Gunay et al. [152], regardless of the model accuracy, the




Figure 5.2: Illustration of the receding time horizons employed in the temperature control algorithm.

stochastic nature of an occupant's behaviours and presence makes it hard to predict the temperature response beyond just a few hours. On the other hand, if the prediction time horizon is too short, preheating or precooling decisions can be made too late and the temperature setpoints may not be met before occupants begin to arrive. In the control of the zone level HVAC equipment, the nighttime setback to daytime setpoint transitions require the longest planning periods [16]. As reported by Candanedo et al. [139], unless the equipment is undersized, setback-to-setpoint transition periods typically range from a few minutes to a couple hours. In line with this, a prediction time horizon of three hours was chosen in this study. The timestep size on this prediction horizon was selected 30 min. Note that the timestep size used by the controller algorithm does not equal to the control time horizon. The control time horizon was five minutes; meaning that the controller plans for the entire three hours, yet it executes the control signal for only five minutes. The five minute control horizon was the default control time horizon for the terminal HVAC equipment in the building where the algorithms were implemented. Figure 5.2 illustrates the prediction (with its timesteps) and the control horizons on a receding time window.

Note the term $(\phi^T R \phi)^{-1}$ in Eqn. 5.5 requires the inversion of a six by six matrix.

Although carrying out the six by six matrix multiplications through nested for loop blocks are possible inside GCL+ programming environment, ordinary zone controllers fail to provide the analytical and computational means to undertake this matrix operation. In the optimization literature, the term $(\phi^T R \phi)^{-1}$ is called the Hessian matrix. Note that it is a function of a single variable x_1 (reciprocal of the time-constant). The Hessian matrix was computed offline for a range of time-constant values (i.e., 50 h for low-mass and 400 h for high-mass thermal zones) [183] – while keeping the prediction time horizon 3 h and timestep size 0.5 h constant. It was noticed that the variation in the entries of the Hessian matrix as a function of the time-constant was negligible ($\sim 1\%$). This underlines that the temperature response for any thermal zone over a 3 h window subject to constant loading conditions is nearly linear. The potential for better control sequences is in characterizing the setpoint values over the prediction horizon \mathbf{T}_{sp} , the free running response of the thermal zone \mathbf{T}_f subject to environmental loads and occupant-driven activities, and the terminal equipment capacity for warming and cooling rate.

The setpoint array \mathbf{T}_{sp} over the prediction time horizon is determined dynamically based on the outputs of the occupancy and thermostat learning algorithms. Weekday temperature schedule can follow three scenarios: (1) temperature setback is applied on occupied weekdays before the earliest expected arrival time $t_{arr,early}$ and after the latest expected departure time $t_{dpt,late}$, (2) temperature setback is applied, if current time exceeds the latest expected arrival time $t_{arr,late}$ yet an arrival has not observed, (3) temperature setback is applied on occupied weekdays, if an occupant takes an intermediate break longer than the longest expected intermediate break D_{abs} . Figure 5.3 illustrates these three scenarios. In the building where the algorithms were implemented, the default setback temperature setpoints were preserved – i.e., 18°C during heating season and 27°C during the cooling season. The viability of more aggressive

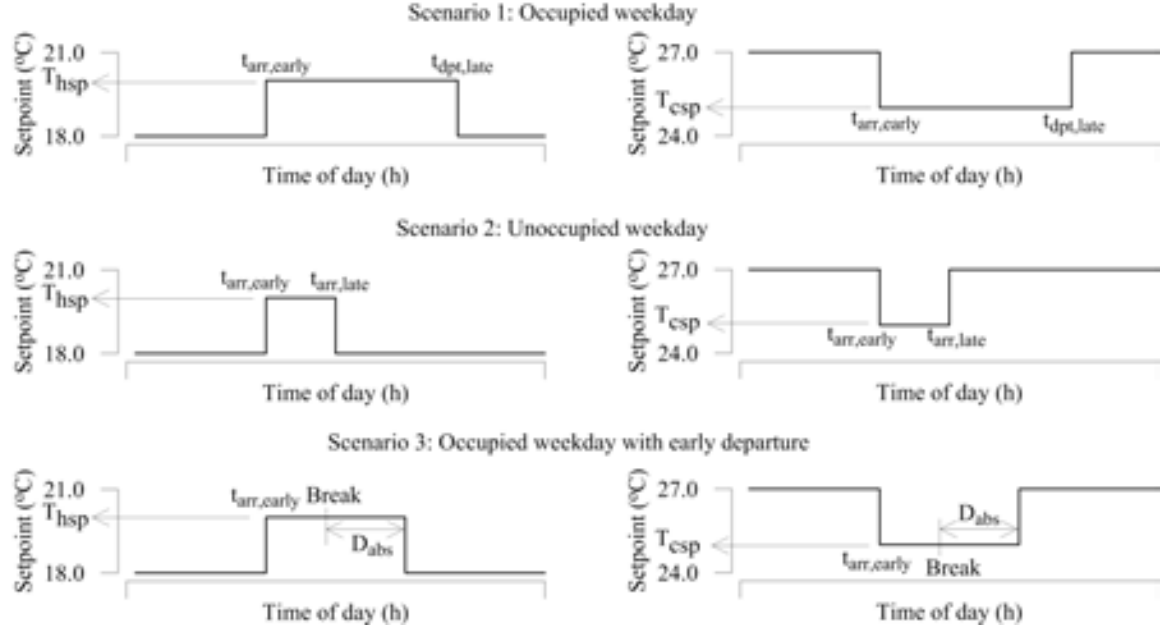


Figure 5.3: The three scenarios the weekday temperature schedule can follow in the temperature control algorithm.

setback temperatures was investigated as well, albeit only in simulation. The temperature setpoints during operation are the outputs of the thermostat learning algorithm T_{hsp} (in the heating season) and T_{csp} (in the cooling season). The temperature setpoints were also set back during weekends. In the unlikely event of an occupant presence detection during any of the temperature setback periods (e.g., weekends), the temperature setback period is terminated for as long as the office remains occupied. Upon an occupant override on the temperature setpoints, the temperature automation is temporarily halted for four hours. During these four hour override periods, the temperature setpoint is maintained as it is specified by the occupant.

Recall that \mathbf{T}_f is the free running temperature in the prediction time horizon subject to environment and occupant-driven loads. This implies that we need to predict the environment and occupant-driven model inputs over the prediction time horizon.

The influence of the occupant-driven loads on the indoor temperature over the prediction horizon is estimated using the presence/absence predictions of the occupancy-learning algorithm (Chapter 2) and the parameter x_4 of the building-learning algorithm (Chapter 4). The environment-driven model inputs are the daylight intensity on the ceiling and the outdoor temperature. If the prediction time horizon (3 h) is short enough that these sensory model inputs tend not to change substantially in that period, one can use the persistence method. The persistence method assumes that the outdoor temperature and the daylight sensor readings remain identical for the next three hours. Figure 5.4 presents the distribution of the mean absolute error in making 3 h ahead predictions by employing the persistence method. It is built by using a year's worth of data gathered from a ceiling-mounted photodiode sensor in one of the monitored offices (west-facing) and an outdoor temperature sensor. When it is assumed that current photodiode sensor readings remain the same over a 3 h horizon, this results in errors larger than 200 lux in more than 10% of the prediction instances. Note that this is about 20% of the range of the indoor daylight sensors used in this study. Based on the range of the parameter estimate x_2 of the data-driven model selected in Chapter 4, a 200 lux error would drift the predictions by 0.15 - 0.20°C/h in these offices. On the other hand, the outdoor temperature predictions made by the persistence method resulted in errors less than 5°C in 90% of the prediction instances. The influence of this error on the model's predictive accuracy is $\sim 0.05^\circ\text{C/h}$ – based on model parameter x_1 .

Florita and Henze [178] reviewed different short-term forecasting methods in different climatic conditions, and recommended the prior moving average method. The method relies on building a characteristic profile of the forecasted variable in recursion. The trend in this characteristic profile in tandem with the current sensory




Figure 5.4: Error in predicting the outdoor temperature and the indoor illuminance three hours ahead by using the persistence and the prior moving average methods.

reading is employed to make a prediction over the horizon. For example, if the daylight sensor reading has been increasing by about 40% from 10 am to 1 pm for the past couple weeks, and the sensor reading currently equals 100 lux at 10 am, the method suggests that the sensor reading will become 140 lux at 1 pm. Figure 5.5 presents the characteristic profiles for the daylight and the outdoor temperature on different months. Note that the characteristic profiles can be generated inside the controller by updating the running mean of the hourly daylight intensity and the outdoor temperature in recursion. When they were employed in the prior moving average method, the accuracy of the 3 h ahead predictions improved substantially for the daylight sensor readings and they remained unaffected for the outdoor temperature. Consequently, in the temperature control algorithm, the outdoor temperature sensor readings over the prediction horizon are computed using the persistence method, whereas the daylight sensor readings are computed using the prior moving average method.

5.4 Implementation results and comfort implications

The adaptive temperature control algorithm – like other algorithms developed in this thesis – is programmed as a standard BACnet [168] program in the scripting language GCL+ – provided as supplemental files [93]. It is implemented in a shared-office space (the controls laboratory presented in Chapter 1) and in the eight private offices (the field implementation presented in Chapter 2).

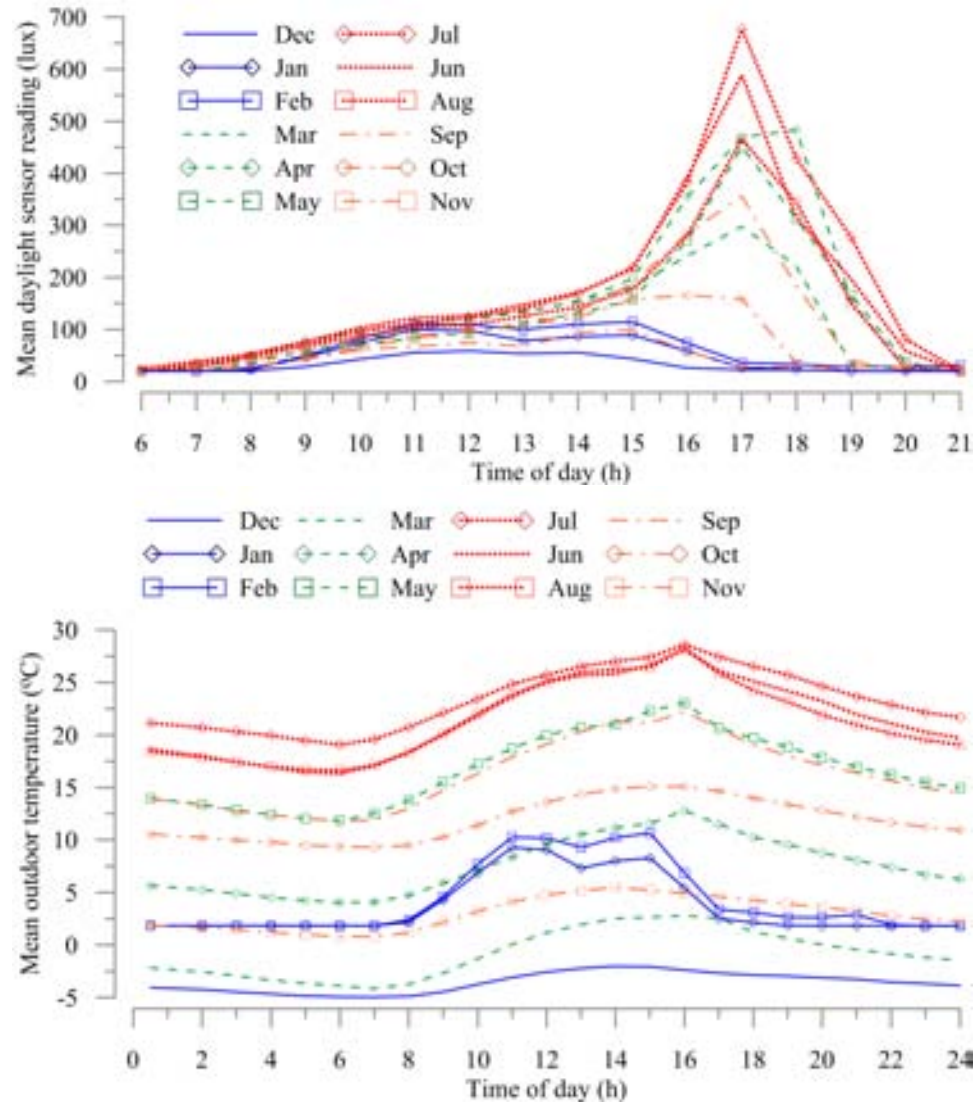


Figure 5.5: Illustrative characteristic profiles for (a) daylight and (b) outdoor temperature sensor readings for different months of the year.

5.4.1 Laboratory implementation

The physical properties of the construction of controls laboratory were determined from the as-built architectural drawings and detailed surveys conducted in an effort to develop a building information model of the building [184]. The steel-framed exterior walls have $3.8 \text{ m}^2\text{C}/\text{W}$ insulation between them. The aluminum-framed windows have a U-value of $3.1 \text{ W}/\text{m}^2\text{C}$, a SHGC of 0.4, and a visible transmittance of 0.6. The window-to-wall ratio on the Southwest-facing facade is 0.5. The cooling and ventilation are provided by a terminal VAV unit (without a reheat coil). The heating is supplied by the ceiling-mounted hydronic perimeter heaters. The temperature control algorithm modulates the VAV units' damper through a proportional control loop, the perimeter heaters' valve, and adjusts the automated-blinds' position during unoccupied periods. Recall that the laboratory has a standalone controls network – which permitted us to conduct rigorous verification tests without interfering with the rest of the building's network.

Note that GCL+ is a programming environment with limited analytical functionalities and no formal debugging environment. As a first step, the scripts of the learning algorithms presented in Chapters 2 to 4 were verified. To this end, the algorithms were also developed inside a more robust programming environment Matlab. The sensory data gathered inside the network were downloaded and employed offline within the Matlab version of the algorithms [119, 185, 186]. The parameters learned by the algorithms inside the controller were verified against those estimated offline.

Figure 5.6 presents the occupancy-learning algorithms' parameter estimates extracted from the laboratory's controls network database on February 2016 – after one year of operation. The results indicate that the earliest expected time of arrival

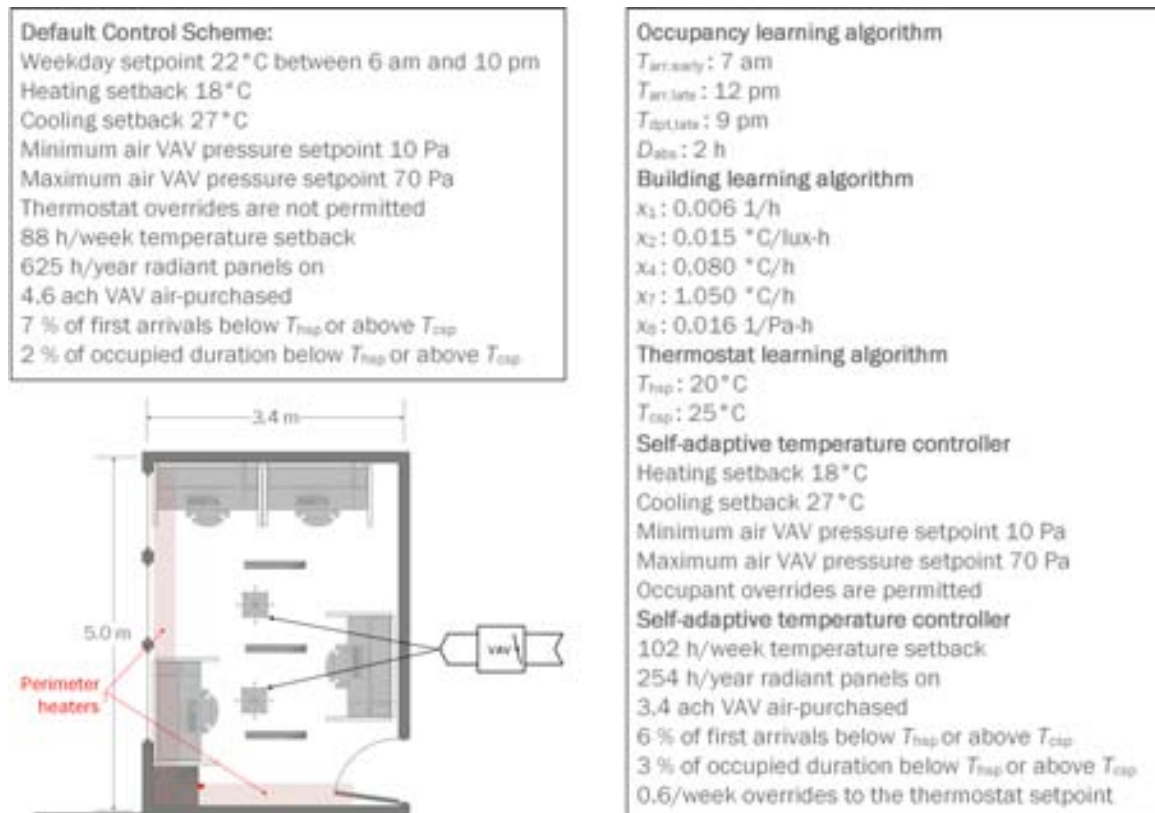


Figure 5.6: An infographic summarizing the implementation results in the controls laboratory.




Figure 5.7: Mean weekly occupancy schedule in the controls laboratory.

$t_{arr,early}$ in this office was 7 am, and the latest expected arrival time was 12 pm. The results also indicate that the latest expected departure time was 9 pm.

Subsequently, the raw motion sensor data were extracted from the laboratory's network data archiver and generated Figures 5.7 and 5.8 using the Matlab function library [119, 186]. This shared-office space was used on average for 43 h/week. The first arrivals and the last departures were unimodal – meaning that late arrivals were uncommon (see Figure 5.8). The duration of intermediate absence periods were almost always less than 2 h (see Figure 5.9). These results were in line the recursively estimated parameter values by algorithms. Similarly, Figure 5.6 shows the building learning algorithms' parameters in February 2016 extracted from the controller database. The parameter estimates were also verified against a batch data analysis conducted using the Matlab function library [185].

The heating and cooling setpoints (T_{hsp} and T_{csp}), despite actively changing based on occupants' interactions, were estimated as ~ 20 and 25°C by the thermostat learning algorithm (Chapter 3). In order to verify this, the timestamps of the thermostat




Figure 5.8: Distribution of the first arrivals and the last departures in the controls laboratory.




Figure 5.9: Distribution of the duration of the intermediate vacancy periods in the controls laboratory.

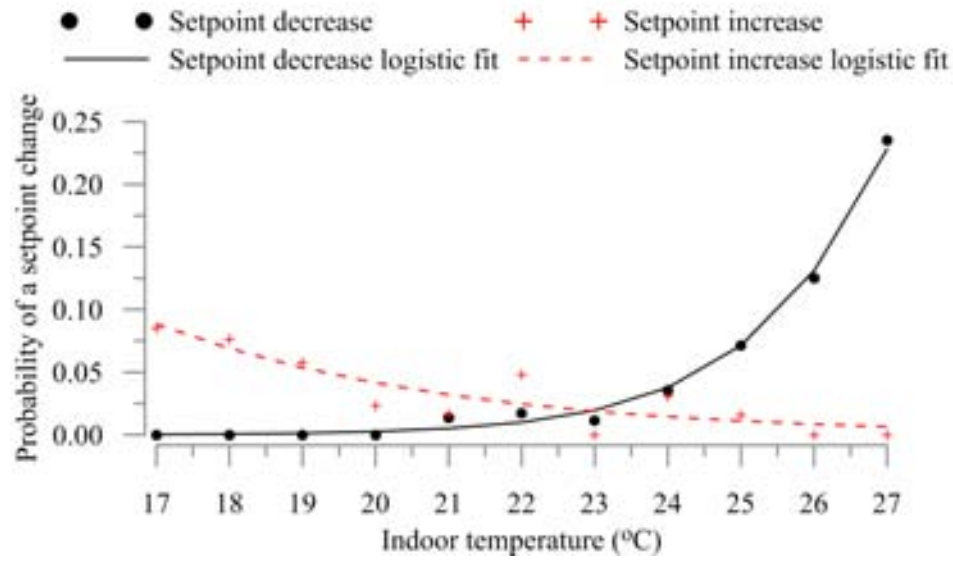


Figure 5.10: Discrete-time Markov thermostat use model predicting the likelihood of a thermostat override in the next 30 min for the controls laboratory.

keypress actions, the occupancy, and the concurrent indoor temperature measurements were analyzed upon a batch data processing using the Matlab function library [119,186]. Figure 5.10 presents the likelihood of observing a thermostat override action in the next 30 mins as a function of the indoor temperature. Results indicate that occupants in this shared-office space respond more actively against warmer temperatures than colder temperature. This can be explained with the fact that increasing clothing insulation is a viable option in the heating season; while opportunities to decrease clothing insulation level can be limited particularly in shared-offices. Figure 5.11 presents the distribution of the indoor temperatures that lead to the thermostat override actions, and Figure 5.12 presents the number of thermostat override actions in 100 occupied hours between different temperature ranges. The analyses confirmed that the likelihood of observing a thermostat override action above 20°C and below 25°C was low.




Figure 5.11: Distribution of the indoor temperature at the keypress instances in the controls laboratory.




Figure 5.12: Number of thermostat override actions in 100 h of occupancy between different temperature ranges in the controls laboratory.

In the controls laboratory, the algorithms continuously learned from the occupants and the building as described in Chapters 2 to 4. However, two different temperature control algorithms were employed at alternating weeks: (1) the adaptive learning control method developed in this thesis and (2) the existing temperature control method in the building where the laboratory is located. The laboratory experiments were conducted over a year – while each week one of the two temperature control algorithms followed the other. In the end of the year, it was confirmed that the distributions of the outdoor temperatures and the solar irradiance during the operation of the two temperature control algorithms were nearly identical. Both temperature control algorithms had an 18°C setback temperature in the heating season and a 27°C setback temperature in the cooling season. In both algorithms, the airflow was controlled by the pressure sensor: (a) 10 Pa for the damper in closed position (~ 40 L/s for the VAV unit with 150 mm diameter [8]) and (b) 70 Pa for the damper in open position (~ 110 L/s for the VAV unit with 150 mm diameter [8]). The default control algorithm's temperature setpoint was 22°C between 6 am and 10 pm on weekdays – it employs the setback temperature setpoints outside this period.

The default temperature control algorithm employs the temperature setback for 88 h/week. By applying the characteristics of the occupants and the building's temperature response, the adaptive learning control algorithm extended the temperature setback periods by 14 h – to 102 h. In addition, with information gathered by the thermostat learning algorithm, the heating setpoint could be lowered by 2°C (from 22 to $\sim 20^\circ\text{C}$) and the cooling setpoint could be increased by 3°C (from 22 to $\sim 25^\circ\text{C}$). As a result, the supply air use by the zone for cooling and ventilation reduced by 25% (from an average of 4.6 to 3.4 air changes per hour). Also, the duration the radiant panel heaters remained on decreased by 60% (from 625 to 254 h/year).

Figure 5.13 compares the distribution of the indoor temperature measurements

taken at the first arrival of the day. At more than 93% of the first arrivals with both the default and the adaptive learning temperature control algorithms, the temperature was between 20 and 25°C. In a few cases, the temperatures at the first arrivals in the heating season were less than 20°C with the default control algorithm. This occurred on cold days (less than -10°C) with an early arrival (earlier than 8 am). Figure 5.13 also compares the distribution of the indoor temperatures during occupancy in the heating and cooling seasons. With both the default and the adaptive learning temperature controllers, only less than 3% of the occupied durations were less than 20°C or more than 25°C. The basic principle of the occupant-learning controls is that the occupants need to willingly interact with the adaptive building systems – so that we can learn from the conditions that they dislike. However, if they interact with their thermostats too frequently, this indicates that they are chronically annoyed by the automation decisions. In this shared-office space (i.e., the controls laboratory), the occupants overrode the temperature setpoints at a frequency of 0.6 times per week. Arguably, this can be considered as a reasonable frequency.

5.4.2 Field implementation

The adaptive temperature control algorithm together with the learning algorithms were implemented inside four commercial controllers serving eight private office spaces in March 2015. These offices were denoted as R1 to R8 in Figure 5.14. Seven of the offices were used by full-time professors in the Faculty of Engineering and Design. One of them was empty (R6). Each office has a 4 m² West-facing exterior window (U-value of 3.1 W/m²°C, SHGC 0.4) and a 10 m² South-facing exterior wall (RSI 3.8 m²°C/W insulation). Recall that the eight offices and the laboratory are located in the same building. Each office has an individually controlled ceiling-mounted radiant

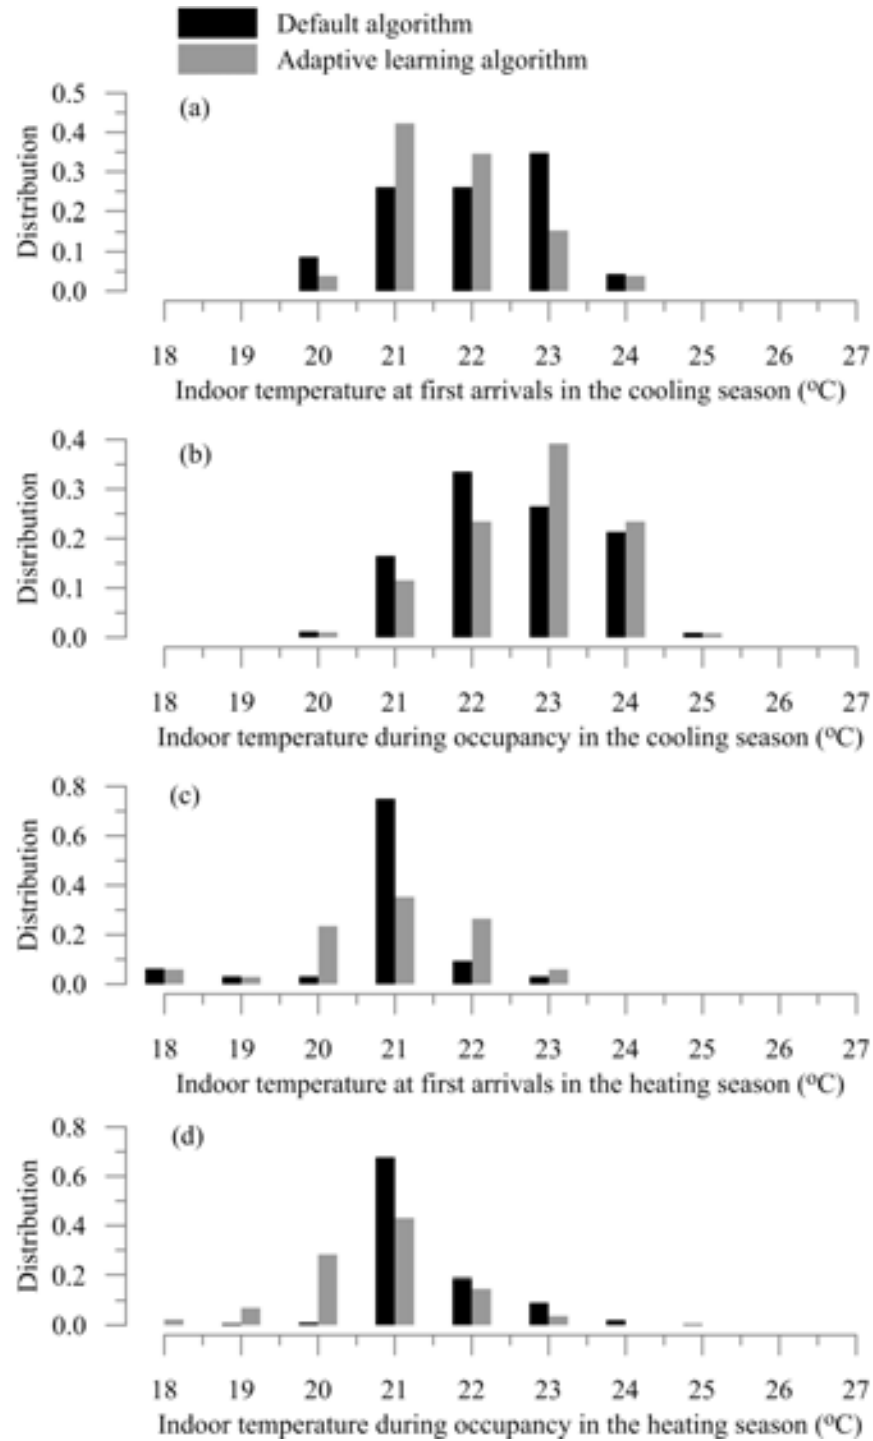


Figure 5.13: Comparison of the indoor temperatures with the default and the adaptive learning temperature control algorithms

panel heater. However, there are only two VAV units serving them. The offices R1 to R4 and R5 to R8 share two VAV units. Each VAV unit has an inlet diameter of 250 mm – modulated the airflow between 200 L/s (10 Pa pressure setpoint) and 700 L/s (70 Pa pressure setpoint) [8]. The VAV units have reheat coils – available during the heating season (October 5 and May 6). When the reheat coil is turned on, it is controlled to maintain a supply air temperature of 30°C setpoint. The temperature control algorithm chooses the states of the terminal HVAC equipment (i.e., the radiant perimeter heaters' valves, the VAV unit reheat coils' valves, and the VAV units' dampers) upon the information recursively derived by the learning algorithms (Chapters 2 to 4).

Figure 5.14 provides a summary of the parameters independently estimated by the occupancy, thermostat, and building learning algorithms in the eight offices. The occupancy and the thermostat learning algorithms' parameters for the unoccupied office (R6) remain as the initial values. Given the learned probability of absence in this office is one, both the earliest and the latest expected first arrival time are 6 am. It remained under temperature setback conditions, unless an unlikely presence event was detected. For the seven occupied offices, the earliest expected first arrival times $t_{arr,early}$ were estimated between 7 and 8 am. The latest expected first arrival times $t_{arr,late}$ and the latest expected last departure times $t_{dpt,late}$ were diverse – $t_{arr,late}$ distributed from 10 am to 2 pm and $t_{dpt,late}$ were between 6 pm and 10 pm. The longest expected intermediate break durations D_{abs} were estimated between 2 and 4 h in different rooms.

The building learning algorithm for each room was trained independently in the zone controllers. Given that these rooms have identical geometry and construction; it

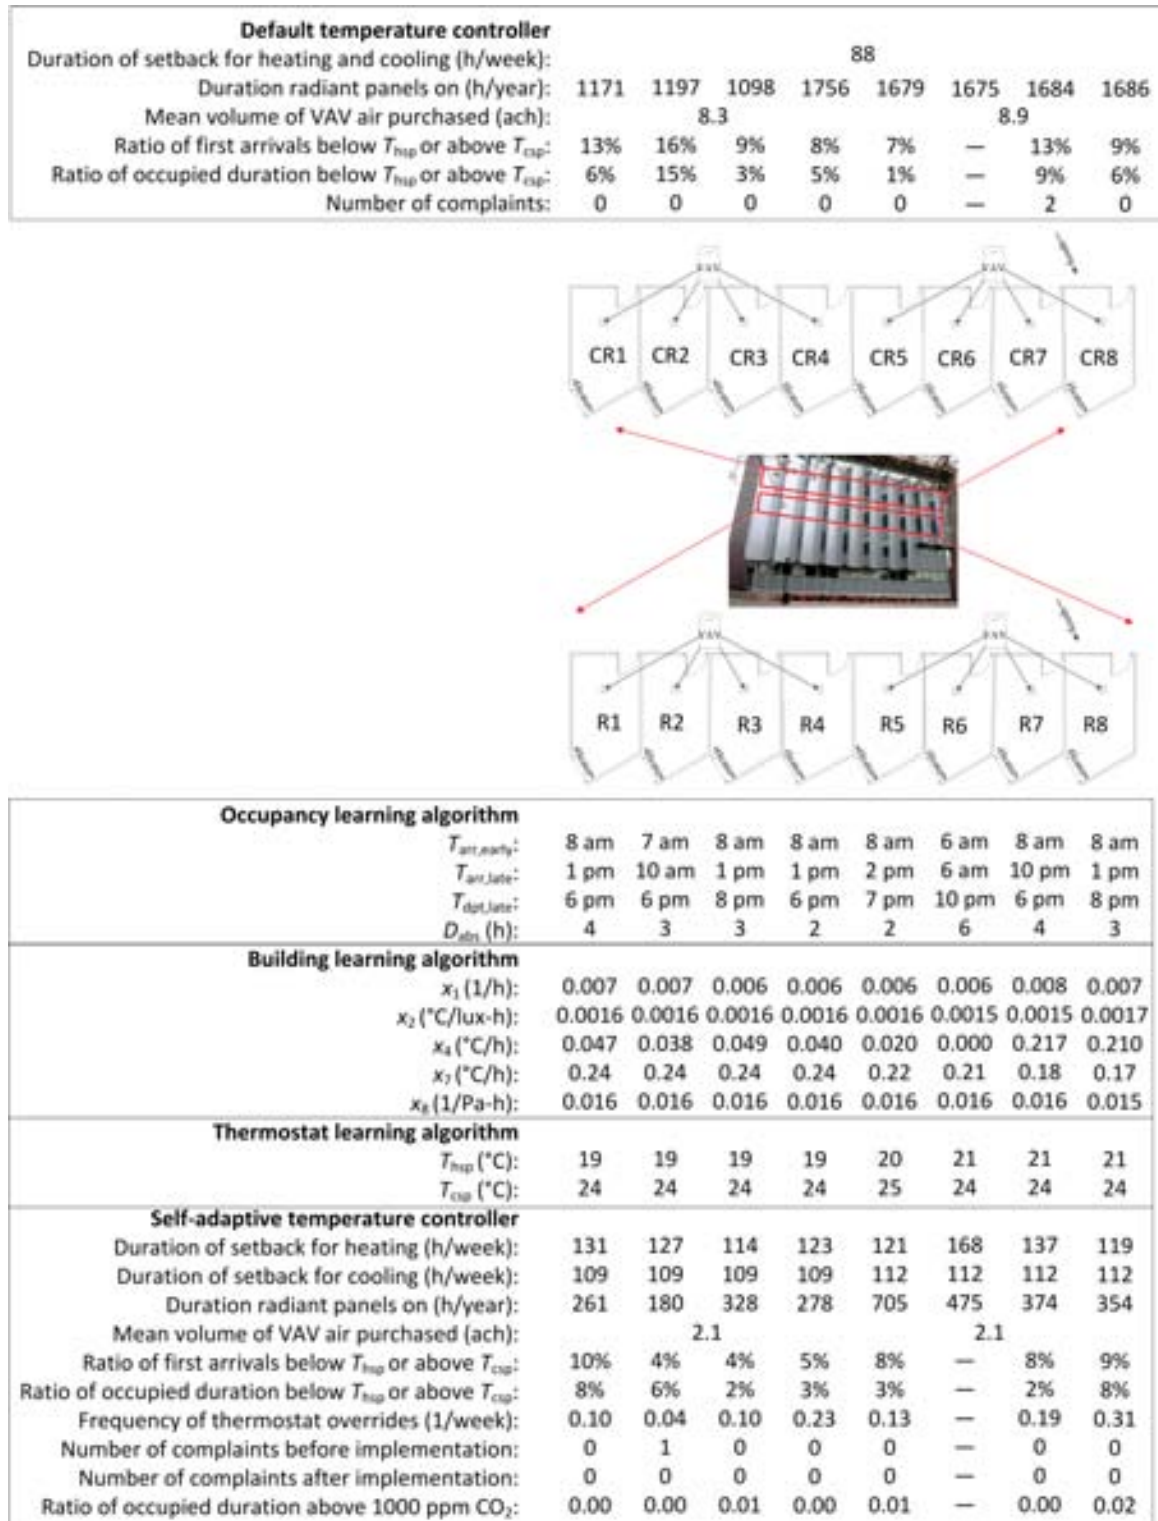


Figure 5.14: An infographic summarizing the field implementation results. The order of the rooms R1 to R8 and CR1 to CR8 were shuffled to preserve the anonymity of the occupants.

is expected that the model parameters characterizing the influence of the environment-driven loads on the temperature response (i.e., envelope losses x_1 and solar gains x_2) would converge to near-identical values in these offices. Similarly, given that these rooms have identical terminal HVAC equipment, it is expected that the model parameters characterizing the influence of the HVAC loads on the temperature response (i.e., radiant panels x_7 and VAV unit x_8) would converge to similar values in these offices. In line with these expectations, the parameter estimates (x_1, x_2, x_7, x_8) in individual rooms attained consistent values (see Figure 5.14). On the other hand, the parameter that characterizes the influence of the occupant-driven loads on the temperature response (i.e., x_4) varied by a factor of ten in the occupied seven offices. Beyond the difference in plug-in equipment choices, the usage patterns can help us explain this variation. The building learning algorithm predicates that occupants need to be present in order to cause any disturbances in the indoor temperature. This may be a reasonable assumption where the occupant and/or a software management software turn off the idling office equipment. However, if the occupants keep their computers on habitually, the model parameter x_4 may become ineffective. For example, both R5 and R8 had similar plug-in equipment load intensities when the occupants were present (~ 80 W). However, the occupant of R8 has a greater tendency to turn off his computer upon departure (the mean plug load during absence periods ~ 30 W) than the occupant of R5 (the mean plug load during absence periods ~ 60 W). Consequently, the parameter x_4 in R5 was considerably smaller than it was in R8 (see Figure 5.14).

The thermostat learning algorithm was employed to estimate the setpoints for the heating and cooling seasons (T_{hsp} and T_{csp}) in each room (see Figure 5.14). Upon the implementation of the adaptive temperature control algorithm, the occupants were




Figure 5.15: An illustrative example of the setpoint evolution process inside the controller database serving R5.

informed that they can change the heating and cooling setpoints. Figure 5.15 presents the dynamic heating setpoint evolution process in a representative office (R5). As intended the occupants infrequently interacted with their thermostats providing the opportunity for the algorithm to finetune the setpoints. Upon an override, temperature setpoint changes abruptly as shown in Figure 5.15. As the occupant spends more time in the office without interacting with the thermostat, the algorithm keeps processing this information to slowly relax the setpoints (i.e., decrease heating or increase cooling setpoint). Recall that there were two reasons for the slow setpoint relaxation process: (1) Occupants when they override, they set the temperature to a value that they think will be comfortable. However, they tend to overcompensate when increasing or decreasing the setpoints – perhaps hoping to accelerate the heating/cooling process according to two thermostat use surveys [31, 187]. (2) Occupants' thermal comfort preferences change in time – primarily because their clothing habits change seasonally [17, 20]. This iterative learning process continued and the setpoint information is converged independently for the heating and the cooling seasons. In the seven offices, this resulted in a 2-3°C reduction in the setpoints during the heating season and a 2-3°C increase in the setpoints during the cooling season with respect

to the default 22°C setpoint in both seasons. In an effort to minimize the thermal interactions between the test offices and the rest of the building, the temperature in the adjacent corridor was controlled to minimize the temperature gradient between the test offices. Also, it was noted that the professors in these private offices tend to keep their doors closed both during absence and presence.

A storey above these eight offices, there were eight other offices with identical geometry, orientation, construction, HVAC equipment and occupancy characteristics. The average weekly occupied duration in CR1 to CR8 was 11 h/week (ranging from 6 to 21 h/week in individual offices), whereas this was 10 h/week in R1 to R8 (ranging from 2 to 17 h/week in individual offices). We will refer to them as CR1 to CR8 (Figure 5.14). Like the rooms R1 to R8, seven of them were used by full time engineering professors and one of them was empty. They also had two VAVs with reheat option (250 mm inlet diameter) and eight ceiling-mounted radiant panel heaters. Both R1 to R8 and CR1 to CR8 were served by the same system level HVAC equipment (same AHU, chiller and boiler plants). Their main difference from the rooms R1 to R8 was that they were controlled by the default temperature control algorithm. The default temperature control algorithm was programmed by one of the largest and most established building controls companies in Eastern Canada – with over 500 employees. Through their continuous commissioning contract with the facilities management and planning, the same company verified the effectiveness of the default temperature control algorithm. Therefore, it is considered that their workmanship can represent a benchmark in quantifying the comfort and energy implications of the adaptive temperature control algorithm implemented in offices R1 to R8.

Figure 5.14 provides a summary of the results one year after the implementation of this adaptive temperature setback algorithm. Recall that the eight offices could be controlled as eight thermal zones for heating (eight radiant panel heaters) and

two zones for cooling (two VAV units). Consequently, the setback schedules could be tailored for each room for heating. However, because each VAV unit serves four offices, the uniqueness of the learned occupancy profiles reduces and the time period during which the temperature setback could be employed in the cooling season becomes shorter than it was for heating. This is because the temperature setback for cooling should be ended for the room with the earliest arriving habits and for the room with the latest departing habits. Regardless of this, the mean temperature setback period increased substantially for both heating and cooling. In contrast to the default control algorithm, the mean temperature setback period increased from 88 to 130 h/week for heating (50%) and from 88 to 110 h for cooling (24%).

The extension in the temperature setback periods can be mainly attributed to the learning of the occupancy and the warming and cooling characteristics of the individual spaces subject to terminal HVAC equipment (Chapters 2 and 4). In addition, the individually selected and learned thermostat setpoints lowered the heating setpoint by $\sim 2^{\circ}\text{C}$ and increased the cooling setpoints by $\sim 2^{\circ}\text{C}$ from the default 22°C setpoint – for both the heating and the cooling. As a result the radiant panel heaters in CR1 to CR8 on average were operational for ~ 1400 h, whereas the radiant panel heaters were on for ~ 400 h in the offices where the adaptive learning algorithms were implemented (R1 to R8). This large reduction in the heating load can be partly attributed to the thermal autonomy of these spaces. Between December 25 and 31 in 2014, the radiant heaters' and the VAV units' heating coils were overridden to turn off by taking advantage of the absence period during the holiday season. As shown in Figure 5.16, it was noted that the indoor temperature with the solar heat gains and the envelope losses balances about 14 to 16°C – even when the outdoor temperature was about -10°C . Consequently, extending temperature setback periods in these intermittently used spaces resulted in these substantial reductions in the space heating




Figure 5.16: Illustration of the free-running temperature response in R8 prior to implementation.

loads. Also, it is worth noting that the default 22°C setpoint both during the heating and cooling season can be considered an extremely conservative and wasteful control decision. Based on ASHRAE Standard 55 [20] – without the thermostat-learning algorithm – the operating setpoints could be selected 21°C in the heating season and 24°C in the cooling season while maintaining the RH between 30 and 60%. However, as noted in Chapter 1, operators and controls technicians are often motivated to find a single temperature setpoint value that minimizes the complaints for the whole year (see Figure 1.3). Despite uncommon in the modelling practice, the year-round 22°C setpoint is not uncommon in other commercial buildings [66].

Similarly, the volume of chilled air use by individual thermal zones between June and August in 2015 was $\sim 75\%$ lower in R1 to R8 than CR1 to CR8. Note that despite the efforts to maintain the corridor at a similar temperature, it is likely that the tested offices R1 to R8 exchanged considerable amounts of heat from the upper and lower storeys. Therefore, the reported load reduction values can be distorted due

to the non-ideal conditions in these field studies. The results should be interpreted cautiously. Regardless, the load reduction potential appeared to be significant – even when the potential end-use savings associated with the distribution of the hot and chilled water and the ventilation air were considered. A secondary benefit of reducing the operating hours of the terminal equipment can be extending the lifetime of the equipment. Anecdotally the ceiling-mounted radiant panel heater valve in this five years old building ruptured and damaged three of faculty offices (rooms CR5 to CR8) with the hot water-glycol solution. Future research should investigate the role of operation scheme on the risks and benefits associated with the ceiling-mounted hydronic heating (e.g., saving floor space; potential hazard against safety, loss of productivity and equipment, comfort issues, and radiant asymmetry).

Another benefit of fine-tuning the setpoints and schedules to user preferences and habits is the diversification of the demand profiles. With the traditional control method, the temperature setback begins and ends at the same time in all spaces with identical setpoints. This tends to create abrupt changes in the space heating and cooling load demands – detrimentally affecting the peak demand utility costs. Figure 5.17 presents the total number of active radiant panel heaters in rooms CR1 to CR8 (rooms controlled by the default algorithm) and in rooms R1 to R8 (rooms controlled by the adaptive algorithm) over a cold week (between -10 and -30°C). These rooms – despite being used by occupants with identical professions, having identical geometry, HVAC and envelope characteristics – resulted in a very diverse set of heating demand profiles (see Figure 5.17.a). On the other hand, the heating demand profiles were nearly identical in rooms CR1 to CR8 – where all rooms were controlled with the identical setpoint schedule. Because there were only two VAV units serving these eight rooms, the diversity discussion could not be extended for the cooling season operation. Future work should involve larger scale implementation




Figure 5.17: Diversification of the radiant panel heater profiles with the adaptive control algorithm.

campaigns to better understand the role of adaptive controls on the load diversity (and peak load reduction and shifting).

Table 5.1 presents the mean and standard deviation of the measured indoor temperatures in rooms R1 to R8 and in rooms CR1 to CR8 during different periods. The rooms R1 to R8 were on average 0.9°C colder than the CR1 to CR8 during the occupied periods in the heating season. During the unoccupied periods, this difference was 1.4°C . In the cooling season, the temperatures were on average 1.5°C warmer in R1 to R8 than CR1 to CR8 both during occupied and unoccupied periods. When contrasted with their learned indoor temperature preferences (T_{hsp} and T_{csp}), less than 10% of the first arrivals took place outside the temperature range between T_{hsp} and T_{csp} ; and less than 8% of the occupied periods were outside the learned temperature

Table 5.1: Mean and standard deviation of the indoor temperature in CR1 to CR8 and R1 to R8.

Indoor temperature (°C)	Default control algorithm		Adaptive control algorithm	
	Mean	Standard Deviation	Mean	Standard Deviation
Heating season during occupancy	21.5	0.9	20.6	1.0
Heating season during absence	21.3	1.4	19.9	1.3
Cooling season during occupancy	21.3	1.7	22.8	1.6
Cooling season during absence	21.5	2.1	23.0	2.1

ranges for these offices. In rooms CR1 to CR8, at $\sim 10\%$ of the first arrivals were either below 19°C or above 25°C ; and $\sim 8\%$ of the occupied duration was outside this range. Note that these 19 and 25°C values were conservatively selected as the lowest T_{hsp} and the highest T_{csp} in R1 to R8. Considering that the setpoint value for the default temperature controller was 22°C for both heating and cooling seasons, the adaptive controller was at least as successful as the default controller in maintaining the indoor temperature within the designated (learned comfortable) temperature range during occupied hours.

Figure 5.18 contrasts the indoor temperature distributions during occupied periods in R1 to R8 and in CR1 to CR8. The empty offices in both groups (R6 and CR6) were not included in this analysis. The top and bottom of the whiskers were set to enclose 90% of the occupied periods. Results indicate that the lowest expected indoor temperatures during the heating season in R1 to R8 were between 19 and 20°C . On the other hand, these were between 20 and 21°C in CR1 to CR8 (see Figure 5.18.a). At 90% of the first arrivals during the heating season, the indoor temperatures were 19°C or larger (except R8) in rooms R1 to R8. These were 20°C or larger (except CR2) in rooms R1 to R8 (see 5.18.c). During the cooling season, 90% of the indoor

temperatures during occupied periods were recorded between 20 and 26°C in R1 to R8. This range was between 18 and 24°C in CR1 to CR8.

The unsolicited feedback of the occupants regarding their thermal comfort was sought through the analysis of the institutional historical complaint logs (from the facilities management and planning) and the frequency of thermostat interactions. A survey to solicit information regarding their thermal comfort was refrained for three reasons: (1) the sample size in the implementation was too small to retrieve reliable information through a transverse survey (7 in control group and 7 in test group), (2) the lack of interest/commitment to fill out longitudinal surveys (e.g., filling out weekly surveys), (3) the phenomena known as the Hawthorne effect [113] that the subjects can change their attitudes, habits, and behaviours when they know they are being monitored and they have direct access to someone (the researchers) who can change the mode of operation.

In CR1 to CR8, there were only two temperature-related complaints – by the same occupant CR7 reported being too cold. In R1 to R8, there was a single temperature-related complaint prior to the implementation of the algorithm – the occupant in R2 reported being too cold. After the implementation of the adaptive temperature control algorithm, there were not any temperature-related complaints recorded in R1 to R8. This can be interpreted as the change from the default temperature control algorithm to the adaptive temperature control algorithm did not cause any extreme conditions leading to chronic complaints.

The working principle of the thermostat learning algorithm was that occupants are seldom intentionally exposed to mild discomfort conditions; such that they infrequently interact with their thermostats as we learn from their comfort preferences and adapt the automation setpoints accordingly (e.g., Figure 5.15). Occupants in




Figure 5.18: Indoor temperatures during intermediate occupancy in (a) the heating season and (b) the cooling season. Indoor temperatures at first arrivals in (c) the heating season and (d) the cooling season.

R1 to R8 could interact with their thermostats to change their setpoints temporarily. If they were to override the temperature settings frequently, this would imply that they were not satisfied with the temperatures selected by the algorithm – would have meant that the temperature control algorithm caused annoyance and failed to learn their preferences. As the design intend of the algorithms, the occupants in R1 to R8 interacted with thermostats very infrequently – one in every three weeks to one in every twenty-five weeks. This can interpreted as the occupants were overall satisfied with the temperatures learned by the thermostat learning algorithm. The literature on occupant use of thermostats is vastly limited to survey-based studies [30, 188]. Future research should investigate the link between the timing/frequency of thermostat keypress actions and the user comfort in depth.

The control algorithm – despite not changing the minimum and maximum limits of the VAV units' airflow rate – could affect the IAQ detrimentally as it reduced the amount of ventilation air use by each zone. To this end, the distribution of the measured CO₂ concentration was analyzed during occupied hours in the offices where the control algorithms were implemented (R1 to R8). Results shown in Figure 5.19 indicates that the CO₂ concentration in these rooms were almost always less than 900 ppm. The IAQ standard ASHRAE [77] recommends CO₂ concentration levels lower than 1000-1100 ppm for office spaces (i.e., less than 700 ppm above the local outdoor CO₂ concentration). Simply put, the ventilation rates were adequate. It is also worth noting that a VAV AHU system recirculates 50 to 70% of the return air. Consequently, the CO₂ concentration of the unoccupied spaces such as R6 is affected by the occupancy in the building. Recent research [189] identified that human decision-making progress may further improve at CO₂ concentrations below the ASHRAE standard limit [77] of 1000 ppm. However, the dominant form of HVAC configurations in North American commercial buildings (i.e., VAV AHU) limits our




Figure 5.19: CO₂ concentration during the occupied hours in R1 to R8 – except the unoccupied office R6.

ability to optimally control for zone level CO₂ and temperature concurrently. Future research should investigate the viability of other HVAC configurations such as the dedicated outdoor air systems (DOAS) in conjunction with the VAVs to control for IAQ and thermal comfort in tandem.

5.5 Estimating energy-savings potential

The implementation results provided promising insights on the performance of the adaptive temperature control algorithm. However, the measured performance in CR1 to CR2 was limited to a single building control company's programming assumptions, and the performance of R1 to R8 was limited to a number assumptions made in developing the learning algorithms. In an effort to analyze a range of potential technician decisions in the controls programming for the default and the adaptive learning temperature control algorithms, a simulation-based investigation was conducted by employing a physical model of the eight offices. The model was developed by using the BPS tool EnergyPlus.

An overview of the model parameters are shown in Figure 5.21. The energy simulations were conducted using the standard weather data CWC [69] for Ottawa.

The plug-in equipment loads and electric lighting use schedules used in this BPS model were built upon a year's worth of current sensor and light switch data gathered from each office. Each office was modelled with 140 W lighting power use for 4 recessed T8 florescent light bulbs. A year's worth of occupancy data from each office was imported as an occupancy schedule in 15 min timesteps. The averages of the weekly occupancy, plug-in equipment load, and lighting in the eight offices are shown in Figure 5.22. The blind occlusion rate in the offices were monitored for eight months by a time-lapse camera. It was noted that once the blinds were positioned to avoid glare, they were not moved for weeks (an average of 1.7 blinds use actions in a month). As a result, the blind occlusion rate did not change as a function of solar irradiance. Thus, the mean blind occlusion rate in each office was set as a constant schedule, as shown in Table 5.2. Because EnergyPlus cannot actuate blinds with partial openings – it can actuate them as open or closed only – each window was divided into vertically stacked four pieces such that the blinds could be actuated as open, one fourth closed, half closed, three fourths closed, and fully closed. The wall insulation (RSI), SHGC and U-factor (USI) values used in the model shown in Figure 5.21 were taken from the BIM database of the building. The solar transmittance value for the blinds were taken as 21% – tested and reported by the manufacturer in compliance with ASHRAE Standard 74 [190]. The infiltration rate was set as 0.2 ach [191]. The sensitivity of the EnergyPlus predictions to this assumption was tested by selecting a run case with the default algorithm and a run case with the adaptive algorithm. The infiltration rate parameter was changed from 0.1 to 0.5 ach at 0.1 ach increments for both of these run cases. Results shown in Figure 5.20 indicate that even a 0.1 ach change in the infiltration rate assumption accounts for great uncertainty in the heating and cooling load intensity predictions of the EnergyPlus model. However, the percentages of the load reductions with the adaptive temperature control algorithm

remained relatively unaffected with the changes in the infiltration rate. The reader should interpret the simulation results accordingly. Similar to the actual building, four adjacent offices were modelled to share a VAV unit with reheat option. The VAV was modelled to have a supply air temperature setpoint of 13°C in the cooling season and 30°C in the heating season. The maximum airflow rate for the VAV was set to 700 L/s. During the operating hours, the VAV unit was modelled to provide outdoor air at a minimum of 40 L/s. When outdoor air temperature was advantageous for sensible cooling, an economizer was set to increase the outdoor airflow rate. When the exhausted air is advantageous for sensible heating, a heat recovery unit of 70% efficiency was modelled to temper the cold outdoor air. The heating was set to be available during the heating season from October 5 to May 6 and the cooling was modelled to be available during the remainder of the year. Each office was modelled to have a 500 W radiant panel heater. The 500 W per perimeter heating capacity was calculated by using the hot water mass flow rate and the hot water temperature drop across the inlet and the outlet of the radiant panel heaters under steady-state conditions. Evidently, the HVAC modelling assumptions do not represent the system efficiency of the ceiling-mounted radiant panel heater system served by a central boiler plant, but note that the BPS model was intended to compute the relative changes in the heating and cooling loads – not the HVAC equipment energy use. In order to avoid unrealistic temperature gradients in the air-node, it is assumed that the thermal mass of the furniture that joins the thermal response of the airspace was 10 times the thermal mass of the room air [192]. The simulation time-step size was 15 minutes. The temperature control algorithms (both the adaptive and the default) were implemented inside the EMS application of EnergyPlus.




Figure 5.20: Sensitivity of the EnergyPlus results to the infiltration rate assumption.




Figure 5.21: An overview of the EnergyPlus model.

Table 5.2: The mean blind occlusion rates in R1 to R8 – rounded to the nearest multiple of 0.25.

RM 1	RM 2	RM 3	RM 4	RM 5	RM 6	RM 7	RM 8
0.25	0.25	1.00	0.00	0.25	0.00	0.50	0.00




Figure 5.22: Average weekly occupancy, plug-in equipment load, and lighting schedules.

For the default temperature control algorithm, the influence of six decisions made by the controls technicians were investigated: (1) Start time of the operation (i.e., end of the nighttime setback), (2) Stop time of the operation (i.e., start of the nighttime setback), (3) Temperature setpoint during the heating season, (4) Temperature setpoint during the cooling season, (5) Temperature setpoint during the setback periods in the heating season, (6) Temperature setpoint during the setback periods in the cooling season. A wide range of typical technician decisions for each of these parameters were randomly generated. The simulated technicians were permitted to program the start time of the operation anytime between 1 am and 10 am, and they were modelled to stop the operation anytime between 4 pm and 12 am. They were able to choose temperature setpoints during the heating season amongst three discrete temperatures: 20, 21, or 22°C. They were able to choose temperature setpoints during the cooling season amongst four discrete temperatures: 22, 23, 24 or 25°C. They were modelled to select the temperature setpoint during the setback periods in the heating season amongst nine discrete values: 15, 16, 17, 18, 19, 20, 21, 22°C. The temperature setpoint during the setback periods in the cooling season were selected amongst

nine discrete values: 22, 23, 24, 25, 26, 27, 28, 29, 30°C. Note none of these values can be considered unusual or atypical – they simply reflect different perspectives in a controls technician's programming style. A Matlab script was written to randomly generate and to write the six parameters representing the technician decisions inside the EnergyPlus model. Thus, the EnergyPlus simulations were repeated 600 times with 600 different technician decisions for setpoints and schedules. Conservatively, it is assumed that all controls technicians would program the weekends as setback periods – which was not the case for the monitored offices in the data analyses portion of this thesis (Chapters 2 to 4).

In addition, the impact of two assumptions made in developing the adaptive temperature control algorithm were investigated. In the implementation, the default 3°C setback was preserved – e.g., reducing the setpoint from 21 to 18°C in the heating season and increasing the setpoint from 24 to 27°C in the cooling season. In this simulation-based investigation, the sensitivity of the algorithm's performance to this assumption was examined at nine different discrete setback values: 1, 2, 3, 4, 5, 6, 8, 9°C. In the implementation, the second assumption pertained to the definition of an expected occupancy event. Rather arbitrarily, an expected occupancy event was defined when its likelihood was equal or more than 10% (z-score 1.3) – e.g., latest expected departure time was estimated as follows: $t_{dpt,late} = \mu_{dpt} + 1.3\sigma_{dpt}$. Through simulation, the sensitivity of this assumption was studied by selecting amongst six discrete z-score values: 0.5, 1.0, 1.5, 2.0, 2.5, 3.0. The Matlab script was employed to write each of the 54 combinations of the two parameters inside the EnergyPlus model and repeat the simulations.

The simulation results were evaluated based on three metrics: (a) the annual heating load intensity, (b) the annual cooling load intensity, and (c) the discomfort

which was defined as the ratio of the occupied duration the learned temperature setpoints shown in Figure 5.14 not met (i.e., below T_{hsp} or above T_{csp}).

The scatterplot matrices shown in Figure 5.23 present the correlation between individual operational parameters and the cooling load/discomfort with the default temperature control algorithm. Each scatter point represents one of the 600 run cases. The results indicate that the start and stop times do not affect the cooling loads substantially. This can be attributed to the fact that the cooling loads in Ottawa are concentrated during the day – not during arrival and departure periods. For the same reason, increasing the temperature setpoint during the setback periods in the cooling season beyond 27°C resulted in only marginal reductions in the cooling load. For the modelled offices, the most influential parameter affecting the cooling loads appears to be the setpoint choices during the operating hours. The simulation results indicate that a 3°C change in the cooling setpoint choices during the operating hours accounts for a 20% ($\sim 9 \text{ kWh.m}^{-2}$) variation in the cooling loads. On the other hand, a 9 h increase in the start time extended the occupied duration during which the setpoints were not met to an average of 17% from an average 10%. Shifting the stop time from 10 pm to 4 pm increased this discomfort metric by only 1%. This can be attributed to the fact that occupied durations in these offices mostly end between 4 and 5 pm. Note that in any of the 600 run cases, the discomfort metric did not fall below 5% of the occupied time. This can be partly attributed to the rather late seasonal switchover to cooling (i.e., May 6) and the weekend occupancy.

In a similar fashion, the scatterplot matrices shown in Figure 5.24 present the correlation between the parameters and the heating load/discomfort. Unlike the cooling load, the start time appears to account for a substantial variation in the heating load intensity. This can be attributed to the fact that heating loads in these




Figure 5.23: Correlation between the setpoint/scheduling decisions and the cooling load/discomfort in the default temperature control algorithm. The histogram plots present the distribution of the cooling load and discomfort metric results with individual run cases.

offices were concentrated in the morning hours – when the casual and solar gains were typically negligible. For the same reason, changes in stop time did not affect the heating load as much as the start time. Shifting the start time by an hour earlier seems to reduce the heating load intensity by 1 kWh.m^{-2} . A 2°C reduction in the setpoints during the heating season from 22 to 20°C accounted for an 15 kWh.m^{-2} variation in the heating loads. Overall, simple setpoint and scheduling decisions made with little – if any – guidance by the controls technicians accounted for great uncertainty in a building's energy and comfort performance – $\sim 40 \text{ kWh.m}^{-2}$ in heating loads, $\sim 15 \text{ kWh.m}^{-2}$ in cooling load, and potentially causing chronic discomfort conditions up to 35% of the occupied duration.

Figure 5.25.a provides a visualization of the Pearson correlation coefficients between the parameters representing the controls technician decisions and the cooling load/discomfort. Within the studied range, it appears that the temperature setpoint choices during the setback periods in the cooling season do not affect the studied user comfort metric; however it inversely correlates with the cooling load intensity.




Figure 5.24: Correlation between the setpoint/scheduling decisions and the heating load/discomfort in the default temperature control algorithm. The histogram plots present the distribution of the heating load and discomfort metric results with individual run cases.

Therefore, a reasonable energy saving strategy could be to increase cooling setback temperatures. However, changing any of the four parameters affecting the heating load (i.e., heating setpoint, heating setback temperature, start and stop times) to save energy adversely affects the proportion of the unmet setpoints (see Figure 5.25.b).

One could argue that a controls technician can choose the optimal operational parameters by looking at these 600 run cases, and therefore there is no need for the adaptive controls that choose these parameters autonomously by learning from its occupancy and temperature response. However, the controls technicians almost never have access to the right data analytics tools pertaining to the occupancy and occupant comfort preferences. In addition, at the implementation stage of the controls, the performance data will not be available to them. Even when the right data analytics tools and the performance data are available, programming each zone to a different set of setpoints and schedules requires diligence. For example, despite the availability of a competent energy manager and routine commissioning practices, the university campus in which these analyses were conducted had a single operating schedule and

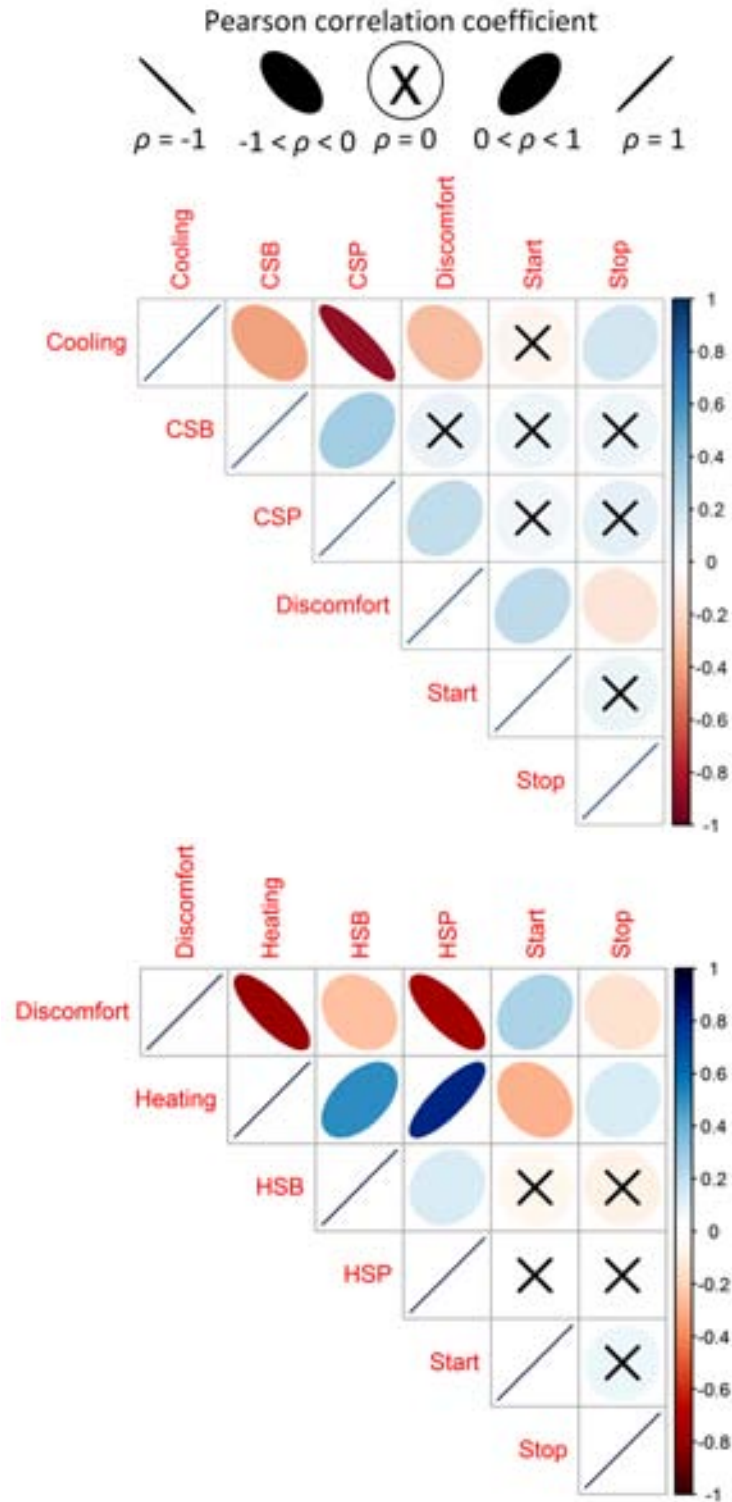


Figure 5.25: Pearson correlation coefficients between the setpoint/scheduling decisions and (a) the cooling/discomfort and (b) heating/discomfort. The axes captions were abbreviated as follows: heating setpoint (HSP), cooling setpoint (CSP), heating setback (HSB), cooling setback (HSB).

a single temperature setpoint for all academic buildings. This is because the facilities management and planning departments' focus in commissioning is often on identifying malfunctioning equipment and addressing user complaints – not on optimizing the setpoints and schedules.

The scatterplots matrices shown in Figure 5.26 illustrates the correlation between the parameters affecting the performance of the adaptive temperature control algorithm and the performance of the algorithm in pairs. Recall that these parameters were the difference between the setpoint and the setback temperatures, and the confidence in terms of the z-score in the expected occupancy events (e.g., earliest expected arrival time). Changing the difference between the setpoint and the setback temperatures from 1 to 10°C appears to account for a 5 kWh.m⁻² variation. The discomfort predictions become more sensitive to the setback magnitude as the z-score values were less than one. The two studied parameters overall accounted for a variation less than 2 kWh.m⁻² (~5%) in cooling load intensity and less than 10 kWh.m⁻² (~15%) in heating load intensity.

Figure 5.27 contrasts the heating and cooling load intensities of the traditional and the adaptive temperature control algorithms with each run case employing parameters covering their feasible operational range. Results indicate that among the run cases those with the ratio of the occupied duration with setpoints unmet of 10% predicted, the adaptive temperature control algorithm had an average of 17% lower heating load and 8% lower cooling load than the default temperature control algorithm. A challenge in the implementation of the traditional temperature control algorithm is that a technician would not know which set of setpoints and schedules would be appropriate for each zone – so that it would cause less than a permissible level of discomfort. As it was in the case for the building in which this study was conducted, this can result




Figure 5.26: The correlation between the parameters affecting the performance of the adaptive temperature control algorithm and the performance of the algorithm.




Figure 5.27: Comparison of the simulation run cases with the default and the adaptive temperature control algorithms.

in energy intensive decisions such as maintaining a setpoint of 22°C in both heating and cooling seasons and applying this setpoint on all seven days of the week between 6 am and 10 pm. As shown in Figure 5.27, unlike the adaptive temperature control algorithm, the default temperature control algorithm run case results spread over a wide range of heating/cooling load and discomfort values. Therefore, the adaptive temperature control algorithm not only tends to reduce the expected heating and cooling loads but also it reduces the risk of poor operational decisions leading to discomfort conditions and/or extravagance in heating and cooling end uses.

5.6 Summary

This chapter introduces an adaptive temperature control algorithm. The algorithm employs the learning algorithms presented in Chapters 2 to 4. The algorithm is intended to choose the state of the terminal HVAC equipment such as VAV units' damper position, reheat coil or radiant heater valve positions.

This adaptive algorithm was implemented in a controls laboratory – a shared office with a standalone controls network. The laboratory experiments were conducted such that the control algorithm was employed only on the alternating weeks. By applying the recursively learned characteristics of the occupancy and the laboratory's temperature response, the adaptive learning control algorithm extended the temperature setback periods by 14%. In addition, with information gathered by the thermostat learning algorithm, the heating setpoint could be lowered by 2°C (from 22 to $\sim 20^\circ\text{C}$) and the cooling setpoint could be increased by 3°C (from 22 to $\sim 25^\circ\text{C}$). As a result, the ventilation air use by the zone for cooling and ventilation reduced by 25% (from an average of 4.6 to 3.4 air changes per hour). Also, the duration the radiant panel heaters remained on decreased by 60% (from 625 to 254 h/year).

In March 2015, the adaptive temperature control algorithm together with the learning algorithms were implemented inside four commercial controllers serving eight private office spaces. A storey above these eight offices, there were eight other offices with identical geometry, orientation, construction, HVAC equipment and occupancy characteristics. These offices were used as a benchmark to assess the performance of the algorithms. By employing the recursively learned occupancy and temperature response information in individual offices, the setback periods were increased by 50% for the heating and 24% for the cooling season. In addition, the individually selected and learned thermostat setpoints lowered the heating setpoint by $\sim 2^\circ\text{C}$ and

increased the cooling setpoints by $\sim 2^{\circ}\text{C}$ from the default 22°C setpoint – for both the heating and the cooling. As a result, the average duration the radiant panels were operational was 400 h in the offices controlled by the adaptive temperature control algorithm; in contrast they were operational for 1400 h in the offices controlled by default temperature control algorithm. Similarly, the volume of chilled air used by the offices between June and August in 2015 was $\sim 75\%$ lower in the offices controlled by the adaptive algorithm than the offices controlled by the default algorithm. In an effort to gather unbiased information about the user comfort, in lieu of conducting surveys, the institutional historical complaint logs, the frequency of the thermostat keypress actions, the indoor air temperature and the CO_2 concentration during the occupied hours were analyzed. However, no evidence was found that the user comfort was adversely affected in the offices controlled by the adaptive temperature control algorithm. Note that occupants were uninformed about the study.

In an effort to analyze the influence of a range of potential technician decisions in the controls programming for the default and the adaptive learning temperature control algorithms, a simulation-based investigation was conducted by employing a physical model of the eight offices. It is shown that simple operational decisions (e.g., setpoint/setback temperature choices and schedules) that are made with little guidance account for great uncertainty in heating/cooling load intensities and user comfort. It was found that the adaptive temperature control algorithm not only tends to reduce the expected heating and cooling loads but also it reduces the risk of poor operational decisions leading to discomfort conditions and/or extravagance in heating and cooling end uses.

The energy savings reported in this study were estimated in office spaces whereby occupants had flexible work hours. Unlike these faculty office spaces, many occupants cannot change the length of workdays. More commonly, office occupants can

shift their arrival and/or departure times or adjust their intermediate break durations. Thus, it is likely that the autonomy of having personal work schedules and the potential for reducing space conditioning loads by increasing the duration of the temperature setback periods were exaggerated in this study. In addition, the unideal testing conditions due to the heat exchange between the test offices and the rest of the building were likely exacerbated the energy-savings potential. The results should be interpreted cautiously.

A secondary benefit of the adaptive temperature control algorithm is the diversification of the load profiles, as it employs different learned setpoints and schedules in individual zones. This is demonstrated anecdotally in the eight test offices. Future work should investigate the demand diversification impact of the control algorithm through large scale implementations. Future work should also investigate the effects of the adaptive temperature control algorithm over the plant level HVAC equipment (e.g., boilers and chillers).

Now that we have finished presenting the adaptive temperature control algorithm, the next chapter will be covering the adaptive lighting and blinds control using the occupant-learning algorithms developed in this thesis.

Chapter 6

Adaptive control of the indoor illuminance

This chapter presents an adaptive indoor illuminance control algorithm. It employs the recursively derived information from the learning algorithms presented in Chapters 2 and 3. The laboratory and field implementation results and comfort implications were presented. Energy-savings potential was investigated through a simulation-based study.

6.1 Literature review

In the reviewed literature two common configurations of daylight-integrated blinds and lighting systems were found: closed-loop [193–196] and open-loop [45, 46, 197, 198] control systems. Figure 6.1 provides an illustration of these two configurations. In closed-loop systems, blinds and lighting automation decisions are based on illuminance measurements taken inside the control zone. Although the occupant comfort research favours illuminance measurements taken at the workplane [13, 40], taking measurements at the workplane can be susceptible to occupants' actions (e.g., covering the sensor surface by papers, shading effect of shelves, monitors). Consequently, the closed-loop systems commonly employ ceiling-mounted illuminance sensors (e.g., photodiodes). When the lighting and blinds can be controlled continuously, the

closed-loop systems can maintain the sensor readings at the intended setpoint limits [196]. However, in practice the shade motors are programmed to stop at a discrete number of positions, and often the lights are non-dimmable or dimmable at a few discrete states. As a result, the control decisions may require the knowledge of the effects of changing the blinds and lighting states on the sensor readings. Otherwise, frequent and conflicting control decisions can be executed. As reported in Lee and Selkowitz [199], an important design consideration in blinds and lighting automation should be to avoid too frequent control actions. The industry practice is to change the blinds and lighting states during the commissioning of the systems, and observe the effects of these changes and incorporate them in the controls programming. For example, if switching on the lights increases the sensor reading by ~ 20 lux and the design intent is to maintain a minimum of ~ 200 lux on the ceiling, the lights are switched off when the sensor reading exceeds ~ 220 lux. However, post-occupancy commissioning was often found necessary to account for the changes in the indoor environment [62, 200]. Because floors in offices are designed to have low reflectances [201], ceiling-mounted sensors typically attain near-open-loop characteristics for lighting controls. Meaning that the influence of electric lighting on a ceiling-mounted illuminance sensor is much smaller than the influence of it on an illuminance sensor placed on the workplane.

On the other hand, the open-loop systems employ illuminance measurements on the window surface through sensors mounted on the window pane between the internal blinds and the window or exterior sensors measuring the incident solar irradiance on the facade [45, 46, 197, 198]. As a single sensor is adequate to automate all the rooms in the same orientation of a building, this is a cost effective way of automating window blinds and lighting. The weakness of this approach is that it ignores the variations in the indoor environment and the blind position. Differences in occupants' position

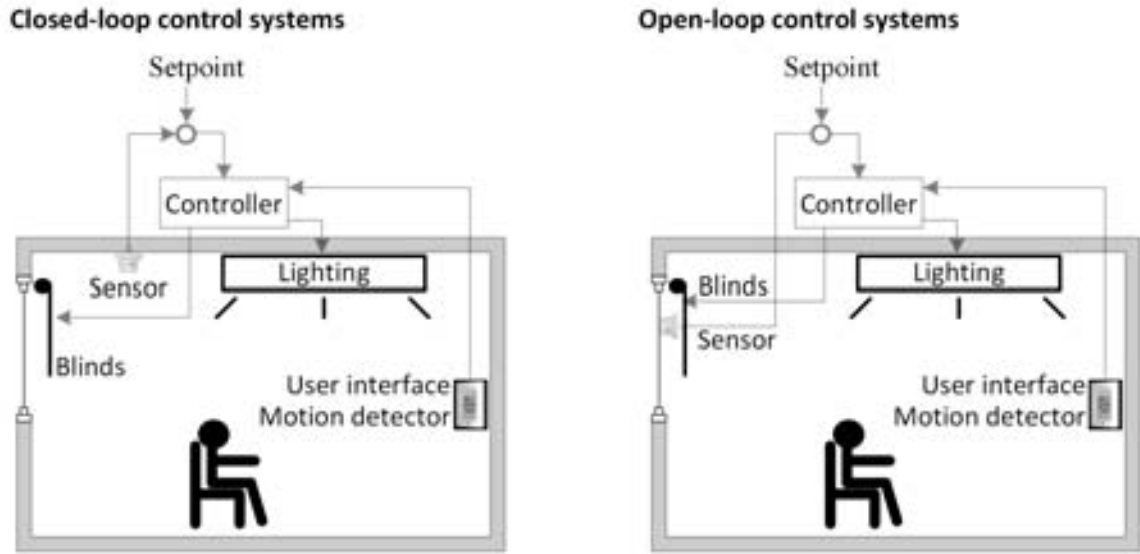


Figure 6.1: Two common configurations for daylight-integrated blinds and lighting automation with examples from the literature: (a) Closed-loop systems [193–196] and (b) Open-loop systems [45, 46, 197, 198]

and the interior design can deem outdoor conditions unrepresentative to execute control decisions. While the closed-loop systems can be used for lighting and blinds automation together or lighting automation alone, the open-loop systems – as they neglect the blind occlusion rate – should not be used for lighting automation only. The research on exterior window shades are not within the scope of this study.

An important gap in the reviewed literature is that the key findings of the occupant behaviour research on blinds and lighting use were not fully integrated during the design of the automation systems.

Occupants prefer to maintain their view and connection to outdoors [202]. In a case study Reinhart and Voss [45] found out that occupants rejected 88% of the blinds closing decisions by the automation. Furthermore, occupants are active in closing their blinds upon feeling discomfort [16]. Therefore, the automation in offices should not close the blinds during occupied hours. The blinds closing actions should

be exclusively left for the occupants. On the contrary, the blinds automation research has been focussing on predicting the conditions leading to discomfort or glare [37, 62]. Simply put, if there is discomfort due to glare, occupants can close their blinds.

Despite their activeness in closing their blinds upon discomfort glare, occupants can be tardier in opening their blinds. It might take days or even weeks before occupants consider reopening their blinds [39, 48]. This leads to a substantial reduction in the daylighting potential. The automation in offices should reopen blinds after ensuring that doing so would not cause glare.

Occupants in offices turn on their lights when they think that it is too dark. There is no need to automatically switch on the electric lighting. Two studies [48, 203] demonstrated that automation systems that switch on lighting automatically with occupancy and daylight can use more electricity than simple manually controlled lighting systems in perimeter office spaces. Furthermore, in some cases, occupants can cover their occupancy sensors with the expectation to stop their lights to turn on automatically [204]. On the contrary, the default setting for some of the most common lighting controllers is to turn on the lights automatically with occupancy detection [48] – regardless of the daylight availability.

Occupants tend to switch off their lights only prior to long absence periods (e.g., at the last departure of a day). They often forget to switch off their lights during intermediate breaks [38, 118]. In presence of motion sensors that switch off electric lighting (15 to 60 min after a departure) and dimmable lights, the tendency to leave their lights on during intermediate breaks further increases [37]. More importantly, regardless of the daylight availability, occupants almost never turn off their lights during intermediate occupancy [22, 45, 205]. During occupancy, the automation should turn off the lighting after ensuring that there is adequate daylight.

Occupants' visual comfort preferences are very diverse [53, 63, 205]. In addition,

occupants' position and orientation in the office [58] and the sensor positioning and control approach [196] play an important role in their perceived comfort. It is therefore challenging to find a set of fixed and prescribed automation setpoints that can satisfy everyone [48]. The setpoints in each office space needs to be tailored for the specific conditions and the individual user expectations. Regardless, the research has been trying to standardize illuminance levels for lighting controls [48].

Occupants want to be able to control their indoor environment. The automation systems that exclude occupants from the control-loop infuriate the occupants [50, 52]. Therefore, occupants should be given a temporary override option, so that they can reject automation decisions with which they are not satisfied.

6.2 Objectives and scope

The objective of this chapter is to develop a lighting and blinds control algorithm. It is intended to use the information recursively derived by the learning algorithms presented in Chapters 2 and 3. Its design intent is to better use daylight in an effort to reduce electric lighting use during occupied hours. It is shaped with the insights gathered in studying occupants' blinds and lighting use behaviour. Similar to the other algorithms presented in this thesis, it is intended to be embedded inside a typical commercial building controller. It is designed to work in a closed-loop system either only for lighting or for integrated lighting and blinds control.

This chapter presents the structure of this lighting and blinds control algorithm and its relationship with the learning algorithms (see Chapter 2 and 3). The laboratory and field implementation of the control algorithm in tandem with the learning algorithms are presented. Energy and comfort related implications of the control algorithms were analyzed through the implementation results. Alternative control

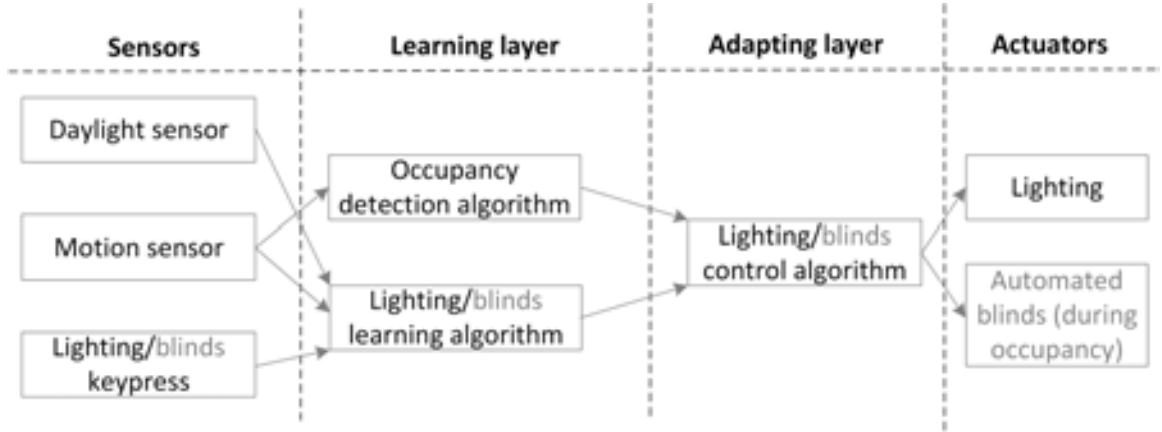


Figure 6.2: The relationship of the lighting and blinds control algorithm with the learning algorithms.

scenarios (and their energy and comfort-related implications) were studied through daylighting simulations conducted in Daysim – a Radiance-based daylighting analysis software [68].

6.3 Lighting and Blinds Control Algorithm

The lighting and blinds control algorithm inputs the outputs of the occupancy detection, and lighting and blinds learning algorithms. Recall that the occupancy detection algorithm (Chapter 2) outputs whether or not the zone is occupied. The lighting and blinds learning algorithm (Chapter 3) outputs two parameters: (1) the illuminance at darkest expected conditions triggering a blinds closing action ($E_{light,sp}$) and (2) the illuminance at brightest expected conditions triggering a light switch-on action ($E_{blind,sp}$). Figure 6.2 illustrates the relationship of the lighting and blinds control algorithm with these learning algorithms.

The algorithm has two simple tasks: (1) switch off lights when doing so does not bring the sensor reading below $E_{light,sp}$ or during vacancy and (2) open blinds when

doing so does not bring the sensor reading above $E_{blind,sp}$ during occupancy. If the blinds are not automated or not integrated to the BAS, the algorithm's task reduces to item (1) only. Recall that $E_{light,sp}$ is estimated recursively based on the user's light switch-on behaviour and $E_{blind,sp}$ is estimated recursively based on the user's blinds closing behaviour. The occupants are exclusively responsible to turn on their lights and to close their blinds. After a manual control action for blinds or lighting, the algorithm remains on hold for two hours – as long as the office remains occupied. During this temporary override period, occupants become fully responsible for the lighting and blinds automation.

In order to ensure that switching off the lights does not cause the sensor reading to fall below $E_{light,sp}$, the effect of switching off the lights on the sensors' reading can be inspected during the commissioning of the system. The effect of opening the blinds can be inspected in a similar fashion – to ensure that opening the blinds do not cause the sensor reading to rise above $E_{blind,sp}$. For example, Figure 6.3 illustrates the cumulative distribution of the change in the sensor readings in response to switching on the lights in ten offices with identical geometry, electric lighting, shading, and overall furniture characteristics – neglecting minor changes in the decoration of the rooms. In these offices, it is logical that the automation turns off the lighting, when the sensor reading is above $E_{light,sp}$ plus a deadband of 50 lux. Figure 6.4 illustrates the ratio of a ceiling-mounted photodiode sensor's readings before and after closing the blinds. In this office, by looking at the lower whisker, the automation is set to open the blinds, when the sensor reading falls below $E_{blind,sp}$ times 0.14. This conservatively implies that opening the blinds would increase the ceiling photodiode sensor's reading by about seven times ($1/0.14$). This manual commissioning process can be automated by changing the blinds position and lighting state automatically at




Figure 6.3: The cumulative distribution of the increase in the ceiling-mounted photodiode sensors' readings in response to switching on the lights in ten identical West-facing offices.




Figure 6.4: The ratio of a ceiling-mounted photodiode sensor's readings before and after closing the blinds in a Southwest-facing office space. The whiskers indicate the 10th and 90th percentiles.

different instances prior to occupancy. Future work should look into developing auto-commissioning routines to expedite and to reduce the cost of the daylight equipment commissioning process.

6.4 Implementation results and comfort implications

The lighting and blinds control algorithm – like other algorithms developed in this thesis – is programmed as a standard BACnet [93] program in the scripting language GCL+. It is implemented in a shared-office space (the controls laboratory presented in Chapter 1) and in five private offices.

6.4.1 Laboratory implementation

Recall that the controls laboratory is a Southwest-facing shared office space. It is a shallow perimeter space with equal window-to-wall and window-to-floor ratios (0.5). The windows have a visible transmittance of 0.7. The top-down interior roller blinds have a visible transmittance of 0.14. The floor surface is covered by a carpet (reflectance of 0.2). The white ceiling tiles have a reflectance of 0.8. The interior walls have a reflectance of 0.5. The reflectance values were measured by a reflectance spectrophotometer. The interior roller blinds are automated by a commercial shade motor – that can be programmed to have two preferred positions: open or closed. The electric lighting is provided by six recessed T8 fluorescent tubes in absence of a dimmable ballast (i.e., two state lighting on or off). Both lighting and automated-shades are integrated into a standard BACnet controls network. Recall that the laboratory has a standalone controls network – which permitted us to conduct rigorous verification tests without interfering with the rest of the building's network. Figure 6.5 presents the layout of the controls laboratory, and the sensors and actuators employed in testing the lighting and blinds control algorithm.

As a first step, the scripts of the learning algorithms presented in Chapters 2

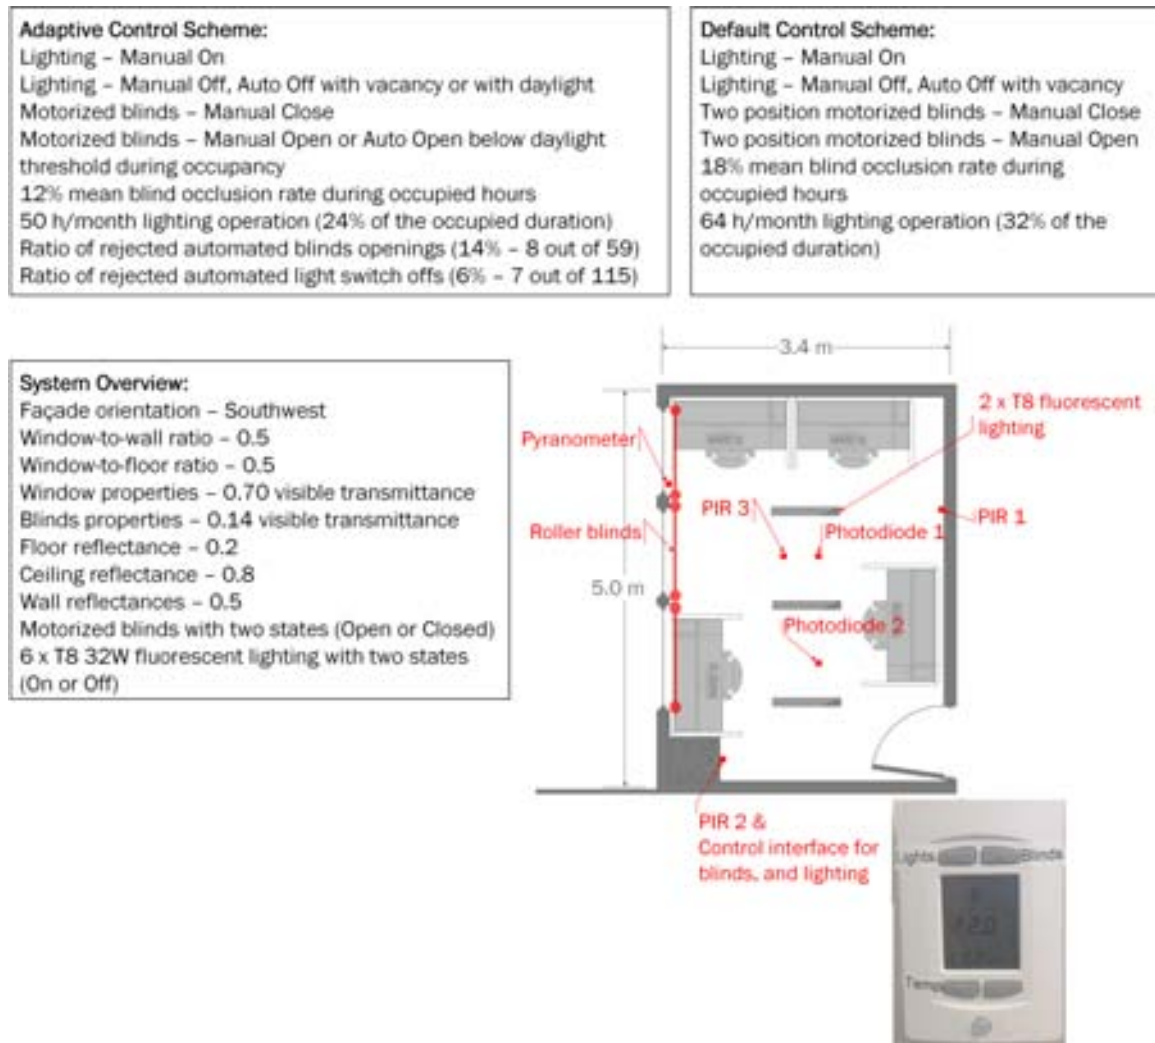


Figure 6.5: An overview of the controls laboratory. The sensors and actuators employed in testing the lighting and blinds control algorithm. A summary of the implementation results.

and 3 were verified. To this end, the algorithms were also developed inside a more robust programming environment Matlab. The sensory data gathered inside the network were downloaded and employed offline within the Matlab version of the algorithms [119, 186]. The parameters learned by the algorithms inside the controller were verified against those estimated offline. Figure 6.6.a illustrates the comparison between the recursively learned blinds closing model inside the controller and the discrete likelihoods estimated from the batch data. Figure 6.6.b contrasts the recursively developed discrete-time Markov light switch on model with the discrete likelihood weights computed offline. Both lighting and blinds models predict the likelihood of an action in the next 30 min.

In the controls laboratory, the algorithms continuously learned from the occupants' lighting and blinds control behaviour. However, two different blinds and lighting control algorithms were employed at alternating weeks: (1) the self-adaptive occupant learning blinds and lighting control algorithm developed in this thesis and (2) a more traditional control method whereby the blinds and lighting were operated by the occupants. In this control method, the automation's sole task was to switch off the lights once a period of 30 min elapses in absence of any movement detections. The laboratory experiments were conducted over a year – while each week one of the two lighting and blinds control algorithms followed the other. At the end of the year, it is confirmed that the distributions of the solar irradiance during the operation of the two control algorithms were nearly identical.

Figure 6.5 provides a summary of the key performance metrics of these two algorithms. With the automated blinds openings introduced by the adaptive controls method, the mean blind occlusion rate during occupied hours reduced from 18% to 12%. The electric lighting use reduction was 22% – from 64 to 50 h/month. In

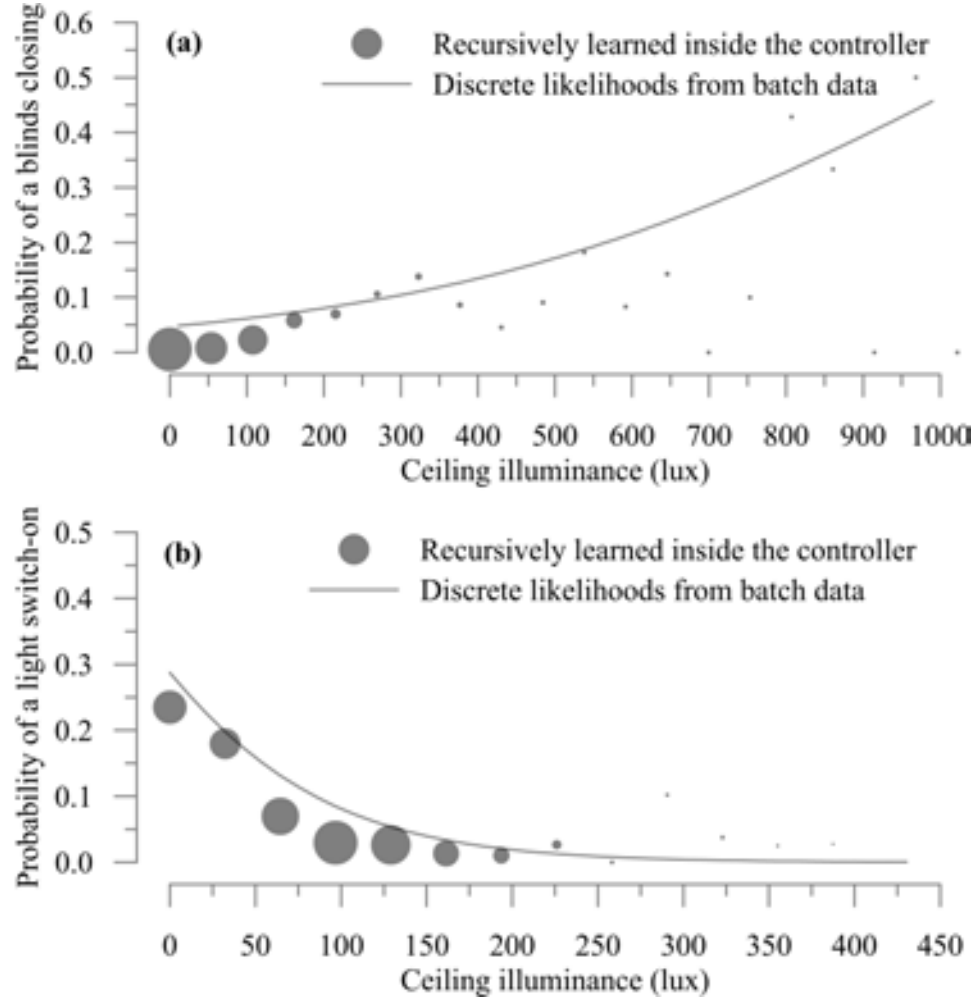


Figure 6.6: Verification of the recursive learning algorithms implemented inside the controllers against those estimated offline using batch datasets. discrete-time Markov models predicting (a) a blinds closing action and (b) a light switch on action in the next 30 min.

other words, the lights were on 32% of the occupied duration with the default control algorithm. This was 24% of the occupied duration during the period the adaptive controller was operational. The reduction in the lighting use can be attributed to the reduction of blind occlusion rate during occupied hours and the automated light switch off events. Figure 6.7 contrasts the ratio of occupied duration the lights were on with the two control algorithms. Results indicate that the lighting use with the two algorithms was comparable when the incident solar irradiance was less than 25 W.m^{-2} . The difference between the two algorithms became more evident as the incident solar irradiance was more than 50 W.m^{-2} . In contrast to those analyzed in Chapter 3, the blind occlusion rate during occupied hours in this room was inherently low at 18% with the default control algorithm. This can be attributed to the fact that the blinds were motorized. Even with the default control mode, the opening and closing actions could be easily executed with a push button. In line with this, Sutter et al. [206] observed that motorized blinds are used three times more frequently than manual blinds. Another interpretation of this can be the quality of the view of the seated occupants [207]. Arguably, the blinds automation could result in higher lighting electricity savings in offices with higher mean blind occlusion rate.

An important metric is the acceptance of the automated control decisions. After an automated blinds opening action if an occupant chose to close the blinds within 15 min, this was interpreted as the automation decision was inappropriate. Similarly, after an automated light switch off action if an occupant chose to switch on the lights within 15 min, this automation decision was deemed inappropriate. As shown in Figure 6.5, the ratio of the rejected light switch off decisions was 6% (7 out of 115). This can be interpreted as the light switch off decisions were overall appropriate. On the other hand, 14% of the blinds closing attempts were overridden by the occupants (8




Figure 6.7: Ratio of occupied duration the lights were on with respect to the incident solar irradiance on the facade.

out of 59). In two examples from the literature the ratios of overridden automated-blinds use actions were 25% in Lee et al. [62] and Reinhart and Voss 45% [45]. In the studied Southwest-facing office space, all of the overrides occurred during the heating season in the afternoon when the solar altitude was lowest. Although the incident solar irradiance on the facade was low in intensity (less than 30 W.m^{-2}), if it beams right onto an occupant, led to user overrides. Conditions leading to light switch on behaviour are related with the lack of diffuse light – which affects the four cubicles rather homogenously. However, conditions leading to blinds closing actions are related with the direct light – which affects the four cubicles unevenly. As a result, particularly in shared office spaces, it is challenging to learn the conditions that lead to blinds closing through a single daylight sensor. A question that was sought to answer was whether or not an open loop blinds controller with a pyranometer sensor would have been more appropriate in predicting the glare conditions. To this end, a silicon pyranometer integrated to the laboratory's controls network was installed on the window pane between the window and the roller blinds facing the incoming

transmitted solar irradiance. The data gathered from the ceiling photodiode and the silicon pyranometer were employed to develop two discrete-time Markov models predicting the blinds closing actions in the next 30 min (see Figure 6.8). Results indicate that the incoming transmitted solar irradiance is not a better predictor for the blinds closing behaviour than the ceiling-mounted photodiode sensor. The main source of uncertainty was whether or not the occupants at the more glare susceptible cubicles were present or not. The occupants, however, expressed their satisfaction when the blinds automatically open as soon as they walk in their offices in the morning. A simple and robust way of automating motorized blinds during occupancy can be simply programming them to open with the first arrival of the day in the morning.

6.4.2 Field implementation

Between September 2015 and April 2016, the adaptive lighting control algorithm was implemented in five private offices (see Chapter 1 and Figure 1.8). Recall that these offices had 4 m² West-facing double-glazed exterior windows with 70% visible transmittance. The window-to-wall ratio on the facade was 32% and the window-to-floor ratio was 0.24%. Each office was equipped with manually controlled top-down roller blinds. They had 14% visible transmittance. The reflectances of the floor, walls, and the ceiling were measured using a spectrophotometer as 0.2, 0.5, and 0.8, respectively. Each office was equipped with a ceiling-mounted photodiode sensor located 3 m away from the exterior window and above the workplane. The geometry of the room and location of the sensors were shown in Figure 1.9. Each office was equipped with four recessed T8 fluorescent lighting. Different from the controls laboratory, the lights were dimmable with five intermediate stages. Because the blinds were not automated, the control algorithm's role was reduced to switching




Figure 6.8: Comparison of the discrete-time Markov models predicting the blinds closing actions in the next 30 min by using a ceiling-mounted photodiode and a pyranometer sensor measuring the transmitted solar irradiance on the interior surface of the window. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bE_{lux} \text{ or } S_{rad})}}$. The properties of the regression parameters were annotated in the figure.

off the electric lighting when there was adequate daylight or when the office was vacated.

As shown in Figure 6.9.a, the algorithm mostly eliminated the lighting use during occupied hours when the ceiling illuminance exceeded 250 lux. In absence of automated blinds, the daylight potential was not fully utilized. As shown in Figure 6.9.b, the lights were on about 40% of the occupied duration even when the incident solar irradiance was more than 300 W.m^{-2} . This can be attributed to the fact that occupants kept their blinds partly or fully closed. Regardless, it is reasonable to expect that switching off the lights when there was enough daylight would have reduced the lighting use. However, the small sample size in this field implementation and the absence of a reference case restrained our ability to draw any conclusions on the energy savings potential from this field implementation. To complement this deficiency, a comprehensive simulation-based investigation was conducted using a daylight model of these offices (see Section 6.5).

Perhaps more importantly, the field implementation results provided invaluable insights regarding the user satisfaction with this daylight integrated adaptive lighting control strategy. For the eight months the lighting control algorithm was tested, the light switch off and dimming behaviours were analyzed. On average, occupants undertook 53 manual switch off actions over 100 h occupancy. Of them, 30 were to dim the lights and 23 were to fully switch them off. The dimming behaviour prior to departure can be explained as the occupants pressed the light switch off button at least once, but they did not necessarily ensure that the lights were fully switched off. This underlines the importance of the user interface design in user's behaviour. Alternatively, the switch off button could be programmed to turn off the lights in one or two stages – in lieu of five. 55% of the full switch off events were triggered

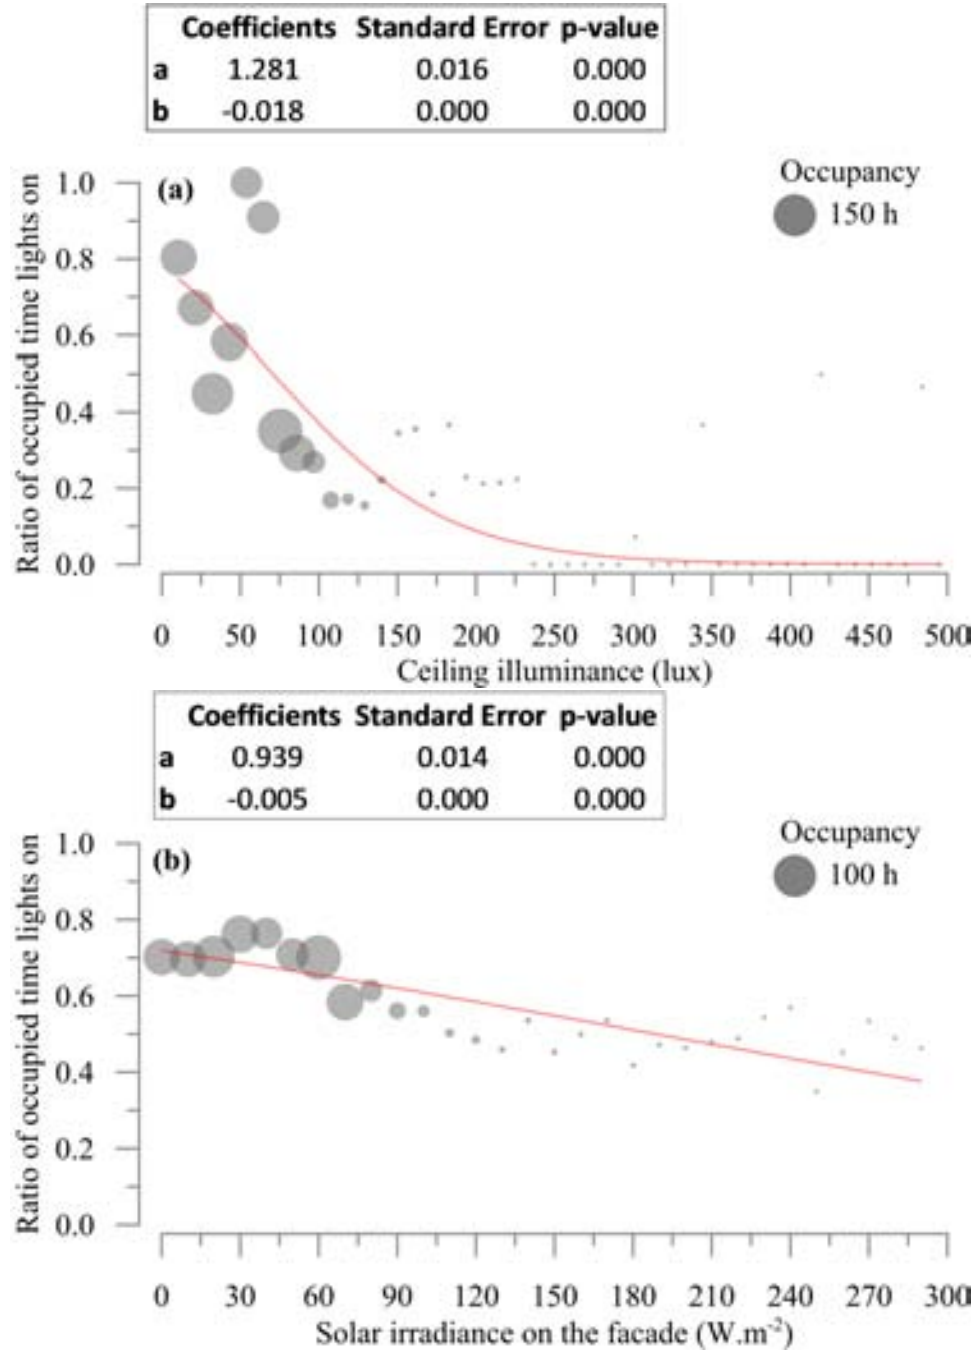


Figure 6.9: Field implementation results. The ratio of the occupied duration the lights were on as a function of (a) the ceiling-mounted photodiode sensor readings and (b) the solar irradiance on the facade. The univariate logistic regression models are in the following form: $p = \frac{1}{1+e^{-(a+bE_{lux} \text{ or } S_{rad})}}$. The properties of the regression parameters were annotated in the figure.

by the automation. On average, 65 automated switch off actions were executed over a 100 hour occupied period. Of them, 9 was due to the detection of vacancy and 56 was due to the detection of adequate daylight. After an automated light switch off decision, if an occupant chose to turn on the lights within 15 min, this event was interpreted as a rejected light switch off decision. It is observed that on average 8% of the automated light switch off decisions were rejected by the user. This ratio was 12% for the automated light switch offs executed based on the detection of vacancy, and it was 7% for those decided based on the detection of adequate daylight. Table 6.1 presents a summary of the breakdown of these metrics amongst the five offices. In the reviewed literature, the ratios of rejected automation-based light switch decisions to the total number of automation-based light switch decisions have not been reported. Only one lighting and blinds automation example from the literature reported the distribution of the unsolicited complaints during a post-occupancy survey; and it indicates that the electric lighting automation related issues can be the cause of a considerable portion of the complaints (19%) [62]. In addition, the light switch off setpoints ($E_{light,sp}$) learned in these offices were shown in Table 6.1. The $E_{light,sp}$ values varied by a factor of four between 108 and 430 lux on the ceiling in the studied offices. Given that these offices had identical geometry and material properties, and identically positioned sensors, this variation can be attributed to the differences in users' visual comfort preferences and the necessity of the adaptive setpoints tailored for individual offices.

6.5 Estimating energy-savings potential

A daylighting model of the studied rooms was developed in Daysim – a Radiance-based daylighting analysis software [68]. An overview of the model parameters is

Table 6.1: Summary of the performance metrics of the self-adaptive lighting learning control algorithm during the field implementation.

Events per 100 h occupancy	Room 1	Room 2	Room 3	Room 4	Room 5
Manual off	5.1	0.8	32.7	1.4	77.1
Manual dimming	11.1	10.0	31.8	25.4	69.1
Auto off	27.9	92.6	5.9	85.4	112.0
Auto off for vacancy	5.4	7.7	0.8	10.6	20.4
Auto off for daylight	22.5	85.0	5.0	74.8	91.7
Rejected auto off for daylight	1.5	3.8	0.0	4.2	9.5
Rejected auto off for vacancy	1.7	1.1	0.7	0.7	1.2
Light switch off setpoint (lux)	269	398	108	108	430

shown in Figure 6.10. The daylight simulations were conducted by inputting the standard weather year data CWC for Ottawa, Canada [69]. By repeating the simulations five times, the illuminance values at the sensor locations were computed for five discrete blinds position: open, quarter closed, half open, quarter open, closed. The illuminance values at the sensor location for each blind position and the incident solar irradiance on the facade were exported to an external programming environment (Matlab). In this programming environment, the occupants' blinds and lighting use behaviour were modelled based on the data collected from the ten monitored offices. The occupant behaviours and presence were systematically randomized to study different lighting and illuminance control strategies with a wide range of user types.

Page et al. [103]'s discrete-time non-homogeneous Markov model was employed to generate the annual presence and absence time-series at 15 min timesteps stochastically. This model inputs two parameters: the weekly occupancy schedule and the parameter of mobility. The mean occupancy schedule in the ten monitored offices is shown Figure 6.11. The individual occupancy schedules were generated by multiplying the mean weekday and the mean weekend schedule values by two random numbers generated from two distinct Weibull distributions (see Figure 6.11). The




Figure 6.10: The daylighting model geometry and material properties.

Weibull distribution's scale and shape parameters were selected to approximate the mean, median, and variance of the parent distribution. The purpose of this exercise was to generate realistically diverse occupancy schedules. Figure 6.11 illustrates 300 randomly generated weekly occupancy schedules – representing 300 different simulated occupants. The parameter of mobility represents the ratio of timesteps with an arrival or a departure event. They were determined from the 16 occupancy data records used in Chapter 2. The mobility parameters were randomly generated and were associated with the individual occupancy schedules. The probability of arrival and departure events at each timestep were computed as follows:

$$\begin{aligned}
 p_{arr}(t) &= \left(\frac{\mu-1}{\mu+1} \right) p_{occ}(t) + p_{occ}(t + \Delta t) \\
 p_{dpt}(t) &= 1 - \left(\left(\frac{p_{occ}(t)-1}{p_{occ}(t)} \right) p_{arr}(t) + \frac{p_{occ}(t+\Delta t)}{p_{occ}(t)} \right)
 \end{aligned} \tag{6.1}$$

where μ is the mobility parameter, p_{occ} is the occupancy schedule value, t is the time of week in hours, and Δt is the timestep (15 min). Based on these likelihood densities for arrival and departure, individual presence/absence time-series in 15 min timesteps were generated (e.g., see Figure 6.11). Note that the random sampling was repeated 300 times for each of the simulated occupants. This value was selected after verifying that it was adequate to generate occupancy time-series with repeatable characteristics. The reader can refer to Page et al. [103] for the model and Gunay et al. [102] for the implementation related details. As shown in Figure 6.12, the occupancy simulation resulted in annual individual occupied durations that range from 300 to 2100 h – with a mean of 1050 h and a standard deviation of 351 h. Note that the mean and the standard deviation of the total occupied duration in the parent distribution were 962 and 450 h, respectively.

The light switch and the blinds use behaviour of the simulated occupants were modelled as follows:

1. At arrival timesteps, the likelihood of a light switch on action was modelled as a discrete-event Markov model shown in Figure 3.20.b.
 2. During intermediate occupancy, the likelihood of a light switch on action in the next 15 min was modelled as shown in Figure 3.20.c.
 3. A light switch off action at departure was modelled as shown in Figure 3.21.
 4. The likelihood of a blinds closing action in the next 60 min was modelled as shown in Figure 3.16.
 5. The blinds opening was modelled as a survival function as shown in Figure 3.17.
- Upon a blinds closing action, the next opening action instance was sampled from this Weibull distribution. After the time elapses, the simulated occupant undertook a blinds opening action at the first occupied instance when the incident solar irradiance was less than 100 W.m^{-2} .

6. The blinds opening and closing increments were sampled from the empirically derived histograms shown in Figure 3.18.

In order to represent the diversity in the simulated population, the light switch models were randomized using Haldi [208]'s approach. This approach assumes that the parameters defining the occupants' behaviours across a population are normally distributed and independent from each other. Note that the occupant behaviour models employed the illuminance readings and the incident solar irradiance on the facade which were computed by the daylighting model. Table 6.2 presents an overview of the parameters of the adaptive behaviour models. For further information on adaptive occupant behaviour modelling methodologies used in this thesis, the reader can refer to [22, 102, 113, 118].

As tabulated in Table 6.3, five different lighting and blinds control scenarios were studied: (1) The lights switch on with the occupancy and switch off with the vacancy, and the blinds are manually controlled by the occupants; (2) The blinds and lighting are exclusively controlled by the occupants; (3) The lights switch off automatically based on the vacancy, the occupants are exclusively responsible for the light switch on and the blinds opening/closing; (4) The lights switch off automatically based on vacancy and the daylight availability, and the blinds are exclusively controlled by the occupants; (5) The lights switch off automatically based on the vacancy and the daylight availability, the blinds automatically open after ensuring that doing so will not cause a discomfort glare, the occupants are exclusively responsible for the light switch on and the blinds closing actions.

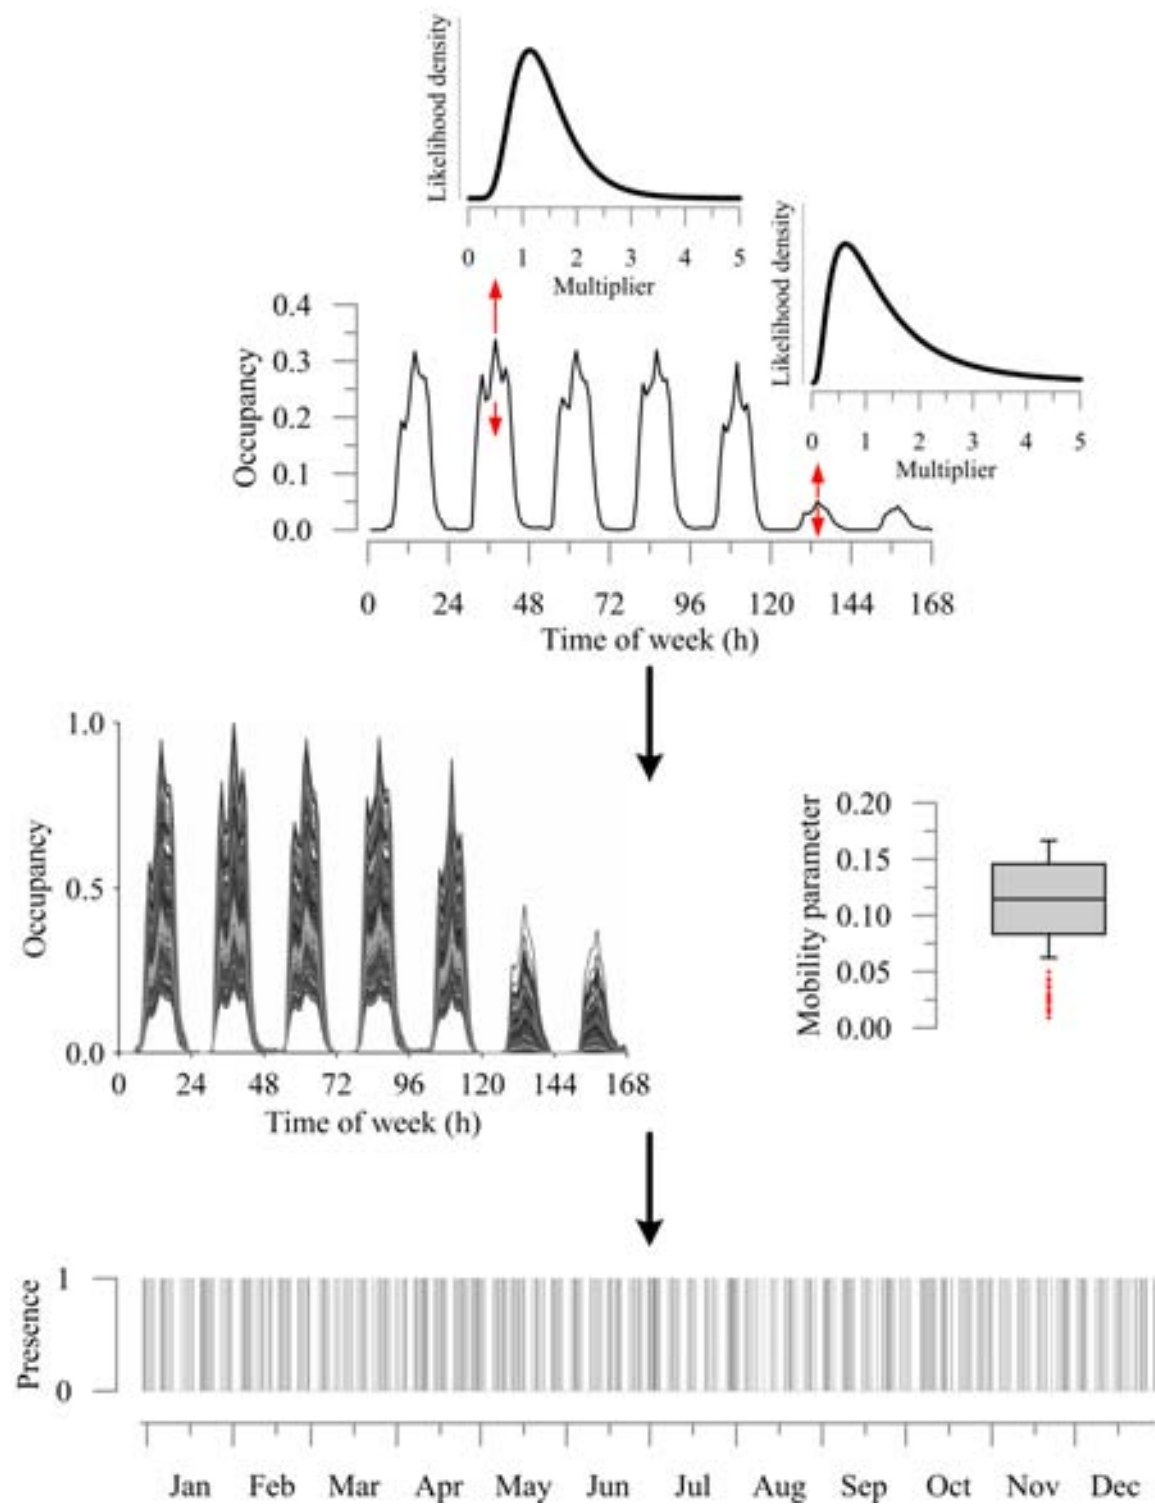


Figure 6.11: Stochastic occupancy time-series generation process.

Table 6.2: Lighting and blinds use model parameters. $N(\mu, \sigma)$ stands for normal and u stands for uniform distributions.

Behaviour	Calling Point	Functional Form	Parameter Coefficients	Predictor
Light switch-on	At arrival	$p_{\text{on, arr}} = \frac{1}{1+e^{-(a+bE_{lux})}}$	$a=N(-0.042, 0.840)$ $b=N(-0.023, 0.014)$	Ceiling illuminance
Light switch-on	15 min timestep	$p_{\text{on, int}} = \frac{1}{1+e^{-(a+bE_{lux})}}$	$a=N(-3.176, 0.400)$ $b=N(-0.025, 0.030)$	Ceiling illuminance
Light switch-off	At departure	$p_{\text{off, dpt}} = \frac{1}{1+e^{-(a+bD_{abs})}}$	$a=N(-1.205, 0.112)$ $b=N(1.177, 0.110)$	Duration of absence
Blinds closing	60 min timesteps	$p_{\text{close}} = \frac{1}{1+e^{-(a+bS_{rad})}}$	$a=N(-3.470, 0.300)$ $b=N(0.003, 0.001)$ $r = u(0, 1)$	Incident solar irradiation
Duration until blinds opening	At blinds closing	$\Delta t = (-e^{(1-r)})^{(1/b)} a$	$a=N(5.720, 1.161)$ $b=N(0.616, 0.053)$	–




Figure 6.12: Duration of occupancy for a synthetic population of 300 occupants. The log-normal model is in following form: $p(D_{occ}) = \frac{1}{D_{occ} scale \sqrt{2\pi}} e^{-\frac{(\ln D_{occ} - location)^2}{2scale^2}}$. The regression parameters are annotated in the figure.

Table 6.3: Studied control scenarios.

Scenarios	Light switch-on	Light switch-off	Blinds closing	Blinds opening
C1	Occupancy	Vacancy	Manual	Manual
C2	Manual	Manual	Manual	Manual
C3	Manual	Vacancy	Manual	Manual
C4	Manual	Vacancy or Daylight $> E_{light,sp}$	Manual	Manual
C5	Manual	Vacancy or Daylight $> E_{light,sp}$	Manual	Manual or Daylight $< E_{blind,sp}$

Note that the scenario 4 incorporates the adaptive occupant-learning lighting control algorithm; and the scenario 5 includes the adaptive occupant-learning lighting and blinds control algorithm. The first three scenarios represent traditional methods to control lighting in office buildings. In compliance with ASHRAE Standard 90.1 [209], electric lighting in most of the recently constructed office buildings has been controlled either with scenario 1 or 3. ASHRAE Standard 90.1 [209] prescribes that using occupancy sensors for lighting automation result in a 10% reduction in lighting electricity use – without dealing with how the occupancy sensors are utilized in light switch on and off decisions. The green building standard LEED credits the use of occupancy sensors for lighting controls – without specifying how they are used again [210]. For example, in the monitored building where the observational data used in this study were gathered, identical offices on different floors programmed by different technicians had these two different control approaches (scenario 1 and 3). The scenario 2 (manual control) is more common in older buildings – whereby BAS-integrated lighting and occupancy sensors are not readily available.

Figure 6.13 presents the distribution of lighting electricity use by different simulated occupants with each of the five control scenarios. A comparison between scenario 1 and 3 (Figures 6.13.a and c) suggests that the way we utilize the occupancy sensors makes a major difference in lighting electricity use patterns. For the modelled offices, scenario 1 (auto on/off based on occupancy) appeared to increase lighting electricity use by 25% in contrast to scenario 2 (manual control). This can be interpreted as the electricity saved by automatically turning the lights off at departure was less than the additional electricity consumed by keeping the lights on during the entire occupied duration. This is because turning the lights on with presence ignores the daylight availability, and it assumes that the occupants would need electric lighting as long as they are present in their office. This was not the case for these perimeter offices.

These observations are in line with the literature [48, 203] – i.e., automated on/off lighting controls with occupancy can consume more electricity than simple manual on/off controls in perimeter spaces. The standards should prescribe that in perimeter spaces occupancy sensor based lighting automation should be configured to turn off the lights with the detection of vacancy – not to turn them on with the detection of occupancy. Properly configuring the occupancy sensor based lighting control systems (auto off with occupancy and manual on) can be a low cost energy saving measure – e.g., a reduction of 720 h of the annual average lighting operation in these intermittently used perimeter offices. For older buildings with manual on/off light switches, installation of properly configured occupancy sensor based lighting systems can be a low cost energy efficiency upgrade as well – e.g., a reduction of 512 h/yr of lighting operation in the monitored offices. Given that the lighting power intensity in these offices was about 10 W.m^{-2} , these lighting operating hour reductions can be estimated as 7.4 and 5.3 kWh.m^{-2} , respectively.

Recall that the ten monitored offices had properly configured vacancy based auto off only lighting controls (see Figure 6.13.c). Of them, seven had their lights on less than 400 h per year. In line with this, 70% of the simulated occupants had less than 400 h lighting usage. Based on the simulation results, with the addition of the adaptive occupant-learning lighting control algorithm, the lighting use duration reduced by 30% (from an average of 309 h/yr in Figure 6.13.c to an average of 202 h/yr in Figure 6.13.d). It is estimated that the lighting use savings would have been about 40% (from an average of 309 h/yr in Figure 6.13.c to an average of 178 h/yr in Figure 6.13.e), if we were to incorporate the adaptive occupant learning blinds and lighting control algorithm. These reductions in the duration of operation of the lights result in only modest reductions in the electricity use intensity – 0.9 kWh.m^{-2}

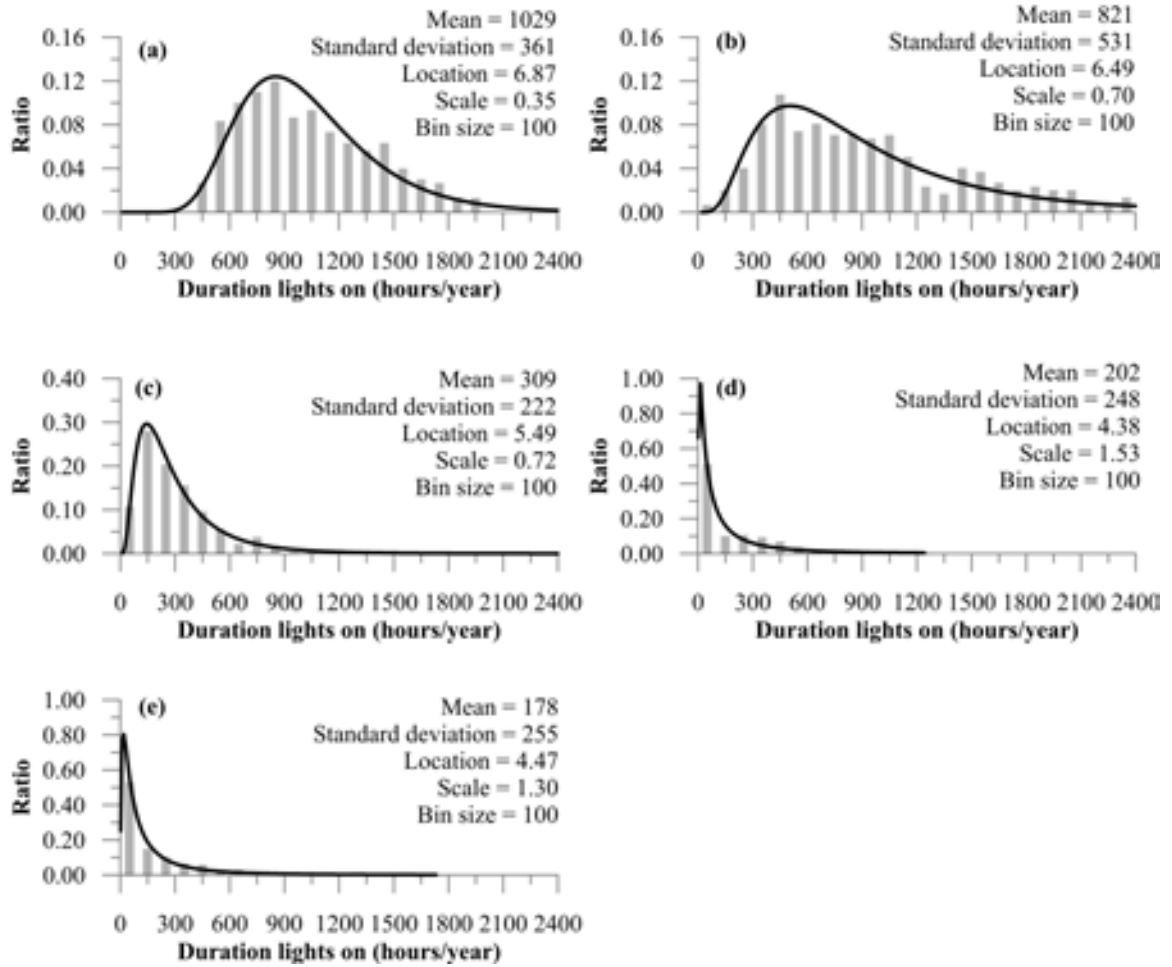


Figure 6.13: The distribution of duration of lighting use by the simulated occupants with the control scenarios (1) to (5) listed in Table 6.3 are shown in (a) to (e), respectively. The log-normal models are in following form: $p(D_{occ}) = \frac{1}{D_{occ} scale \sqrt{2\pi}} e^{-\frac{(\ln D_{occ} - location)^2}{2 scale^2}}$. The regression parameters are annotated in the figure.




Figure 6.14: Simulated mean blind occlusion rate during occupied hours with the manual blinds control (scenarios 1 to 4 in Table 6.3) and with the adaptive occupant-learning blinds control (scenario 5 in Table 6.3). The whiskers indicate 10th and 90th percentile in the distribution.

with adaptive lighting controls and 1.2 kWh.m⁻² with adaptive lighting and blinds controls. Arguably, the lighting load reduction potential would have been larger, if the occupancy in this academic office building was more regular. Note that the professors in the studied offices had flexible occupancy patterns – most of them spent less than 4 h per workday in their offices.

The distribution of the mean blind occlusion rate during occupancy for different simulated individuals are shown in Figure 6.14. When the blinds were controlled manually (C1 to C4 in Table 6.3), the mean blind occlusion rate during occupancy for 90% of the simulated occupants were between 40 and 72%. The mean blind occlusion rates were between 15 and 36% with the addition of the adaptive blinds control algorithm (C5 in Table 6.3) – which opens the blinds automatically after ensuring that doing so will not cause discomfort glare. Improving the view and connection to the outdoors can be considered as a secondary benefit of the adaptive blinds control algorithm.

Switching off the lights automatically (based on occupancy and/or daylight) is expected to increase the frequency of the light switch on actions in contrast to the manual lighting control. This is because occupants occasionally forget to turn off

their lights at departure and they often do not turn off their lights during intermediate occupancy when adequate daylight becomes available. Similarly, opening the blinds automatically (based on daylight) is expected to increase the frequency of the blinds closing actions in contrast to the manual blinds control. This is because occupants can forget to reopen their blinds for a long time. However, a properly configured control algorithm should not cause a substantial increase in the frequency of the occupants' natural behaviour patterns [50, 199]. Figure 6.15 shows that adaptive occupant-learning lighting control algorithm (C4 in Table 6.3) does not make a significant difference in the frequency of light switch on actions in contrast to the default lighting control algorithm in the monitored offices (C3 in Table 6.3). This can be interpreted as the daylight-based light switch off decisions were often accepted by the simulated occupants. For the median occupant, it is estimated that the frequency of the light switch on actions remained the same (7 per 100 hours of occupancy) and the frequency of the blinds closing actions would have increased by 3 per 100 hours of occupancy with the integration of the adaptive blinds control algorithm (C5 in Table 6.3). In brief, the total number of light switch on and blinds closing actions would be about 8 to 12 actions per 100 hours of occupancy with both C4 and C5. Therefore, the addition of the adaptive lighting and/or blinds control algorithms (C4 and C5) do not cause a substantial change in the users' simulated behaviour patterns.

It is worth noting that the primary purpose of this simulation exercise was to better understand the performance of the control algorithms with different user types. To this end, a synthetic population of 300 occupants were randomly generated by regularizing the occupant behaviour and presence models in Chapter 3. In this simulation-based investigation, although the state-of-the-art occupant behaviour and presence modelling methodologies [102, 118, 119] were used, the readers should be cautious in




Figure 6.15: Simulated frequency of (a) light switch on and (b) blinds closing actions by different control algorithms.

extending the simulation results to other contexts due to the limited size of the parent occupant distribution.

6.6 Summary

This chapter presents a closed-loop lighting and blinds control algorithm which employs the occupant learning algorithms introduced in Chapters 2 and 3. The algorithm has two simple tasks: (1) turn off lighting when doing so would not cause inadequate task lighting and (2) open blinds when doing so would not cause discomfort due to glare. The illuminance setpoint range that will likely not cause inadequate task lighting and discomfort due to glare is learned from user behaviour patterns in individual offices – as described in Chapter 3. The cases in which BAS-integrated motorized blinds are not available, the role of the control algorithm simply reduces to item (1).

The adaptive lighting and blinds control algorithm was implemented in the controls laboratory. The laboratory experiments were conducted such that the self-adaptive lighting and blinds control algorithm was employed only on the alternating weeks. After a year of data acquisition, it is observed that the total duration the lights remained on with the adaptive control algorithm was 22% less than it was

with the default control algorithm (auto-off with vacancy detection and manual on lighting, and manually controlled motorized blinds). In addition, with the adaptive control algorithm, the mean blind occlusion rate during occupancy reduced from 18% to 12%. The users accepted 94% of the automated light switch off decisions and 86% of the automated blinds opening decisions.

Between September 2015 and April 2016, the adaptive lighting learning controls were implemented in five private offices. Because the blinds were not automated, the control algorithm's role reduced to switching off the electric lighting when there was adequate daylight or when the office was vacated. The algorithm mostly eliminated the lighting use during occupied hours when the ceiling illuminance exceeded 250 lux. The occupants accepted 93% of the automated light switch off decisions upon the detection of adequate daylight.

In an effort to estimate the energy savings potential of the adaptive lighting and blinds control algorithm in this field implementation, a simulation-based investigation was conducted with the physical daylighting model of these offices and with the stochastic occupant behaviour and presence models – derived using the data gathered from these offices. Based on the simulation results, the addition of the adaptive occupant-learning lighting control algorithm could reduce the lighting use duration by 30%. It is estimated that the lighting use savings would have been about 40%, if we were to incorporate the adaptive occupant-learning blinds and lighting control algorithm. However, in these intermittently used office spaces (with an average duration per weekday spent less 4 h), these reductions in the duration of operation of the lights resulted in only modest reductions in the electricity use intensity – 0.9 kWh.m⁻² with adaptive lighting controls and 1.2 kWh.m⁻² with adaptive lighting and blinds controls.

Chapter 7

Conclusions and future work

Within the scope of this thesis, the occupant behaviour and presence patterns in an academic office building were explored. The existing controls infrastructure was employed to capture occupants' interactions with their light switches and thermostats. Concurrently, the indoor climate at individual room level was monitored through a range of building sensors.

From this data analysis, an adaptive indoor climate control tool was developed. The tool contained a set of novel learning algorithms: (1) Presence detection, (2) Presence prediction, (3) Thermostat learning, (4) Lighting and blinds learning, and (5) Building's temperature response learning. Each algorithm was developed after the survey of the relevant literature and studying a range of candidate methods. The algorithm selection was based on not only the predictive performance but also the transferability to the existing commercial building stock.

Motion detectors are the most common sensor type used in office buildings to detect human presence. Given that movements are discrete-events, in practice a delay value (e.g., 15 to 60 min) is heuristically selected to avoid incorrect vacancy detections during immobility. After each movement detection the space is assumed occupied for this delay period. In the development of the presence detection algorithm, the error in detecting presence by using motion detectors was demonstrated, and the potential of

ameliorating the motion detectors by using complementary sensors (i.e., door contact and CO₂ sensors) was investigated. It was observed that the door contact and the CO₂ sensors offer only a marginal improvement relative to a motion detector with an appropriate time delay. Consequently, a recursive adaptive time delay selecting algorithm was formulated. The algorithm learns from the frequency of the movement detections in real-time and adapts the delay period accordingly. As such, the cases in which the occupants tend to generate detectable movements more frequently, vacancy can be detected earlier than the cases in which the occupants tend to remain immobile for an extended period of time. The algorithm was implemented and verified in the controls laboratory.

In recognition of the transient nature of the heat transfer within building fabric, real-time occupancy detections are not adequate to control heating and cooling systems. In order to adapt the operation of heating and cooling systems, the likelihood of occupancy needs to be predicted well before the space is occupied. However, more than one fourth of the North American workforce reports having flexible work schedules. As the personal preferences and job specific restrictions translate into occupancy profiles, the occupancy in different rooms has been becoming increasingly diverse; and the operators have been challenged to reduce the temperature setback periods to accommodate this diversity. To this end, the recurring occupancy patterns in 16 private offices were examined, and six zone level occupancy parameters with utmost importance for temperature controls were proposed: (1) the earliest expected arrival time, (2) the latest expected arrival time, (3) the latest expected departure time, (4) the longest expected duration of intermediate vacancy period, (5) the longest expected duration of uninterrupted intermediate presence, (6) the ratio of the absent workdays. The presence learning algorithm was formulated to recursively identify the

six occupancy parameters at each zone controller. The algorithm was implemented and verified in the controls laboratory and in eight private office spaces.

The occupants' interaction patterns with their thermostats render an untapped opportunity to retrieve information regarding their temperature preferences. Thus, occupants' thermostat use behaviour patterns in 20 offices were investigated with concurrent occupancy and indoor temperature data. Their natural pace of interacting with their thermostats at varying indoor temperatures was examined. An algorithm was formulated to develop a statistical model that predicts the likelihood of a thermostat keypress action in real-time at the individual zone level. The model was employed to provide two key operational metrics: (1) The coldest expected indoor temperature to observe a setpoint decrease and (2) The warmest expected indoor temperature to observe a setpoint increase. For a given thermal zone, the former represents the operating temperature setpoint during the cooling season and the latter is the operating setpoint during the heating season. The algorithm was implemented and tested in the controls laboratory and in eight private offices.

Occupants' light switch and blinds use behaviours were analyzed with concurrent solar irradiance, ceiling illuminance, and occupancy data in ten private offices. The occupants' characteristics in using their electric lighting and blinds were studied, and their validity outside the ten studied occupants were discussed in the light of the occupant behaviour literature. Similar to the thermostat learning algorithm, an algorithm was formulated to develop statistical models that predict the light switch on and blinds closing behaviours recursively inside a local controller. The models were designed to provide two key operational metrics: (1) the darkest expected indoor illuminance levels to cause a blinds closing action and (2) the brightest expected indoor illuminance levels to cause a light switch on action.

Each office space with its unique casual and solar gains, air-transmission and

distribution, thermal capacitance and envelope conductance, and terminal HVAC characteristics has a signature temperature response. In an effort to filter out the warming/cooling capacity of the terminal HVAC units subject to these environment and occupant-driven loads, an inverse modelling study was conducted based on the data gathered from three perimeter office spaces. This information is needed in calculating the near-optimal setback-to-setpoint transition time in each zone prior to the occupants' arrival. Twelve different greybox models that predict the temperature response of a perimeter office space were developed at varying numbers of model inputs and parameters. Two different recursive parameter estimation methodologies were employed for each of the twelve models: the Extended Kalman Filter and the Particle Filter. A systematic way of selecting one of these models and recursive parameter estimation methodologies was introduced by using the sensory data collected from three perimeter office spaces. The Extended Kalman Filter was found to be an efficient recursive parameter estimation method adequate in zone level building controls applications. The simplest feasible models that can capture the timing and magnitude of the local temperature extrema were the model with 5 parameters and 6 sensory inputs (i.e., indoor and outdoor temperature sensors, an indoor light intensity sensor or a pyranometer, a motion sensor, a radiant panel heater's state, a VAV unit's discharge temperature and pressure). A recursive algorithm was formulated to identify the unknown parameters of this model inside a local controller. The algorithm was implemented and tested in the controls laboratory and in eight office spaces.

A temperature and an indoor illuminance control algorithm were introduced. These control algorithms employed the learning algorithms. Their setpoints, schedules, and equipment capacity information were autonomously tailored for an individual zone's occupancy and adaptive occupant behaviour patterns, and terminal HVAC unit characteristics.

The adaptive temperature control algorithm was first implemented in the controls laboratory. The laboratory experiments were conducted such that the control algorithm was employed only on the alternating weeks. One year after the implementation, the adaptive learning control algorithm extended the temperature setback periods by 14%. In addition, with information gathered by the thermostat learning algorithm, the heating setpoint could be lowered by 2°C (from 22 to 20°C) and the cooling setpoint could be increased by 3°C (from 22 to 25°C). As a result, the air used by the zone for cooling and ventilation reduced by 25% (from an average of 4.6 to 3.4 air changes per hour). Also, the duration the radiant panel heaters remained on decreased by 60% (from 625 to 254 h/year).

The adaptive temperature control algorithm is then implemented inside four commercial controllers serving eight private office spaces. A storey above these eight offices, there were eight other offices with identical geometry, orientation, construction, HVAC equipment and occupancy characteristics. They were used as a benchmark to assess the performance of the algorithms. One year after the implementation, by employing the recursively learned occupancy and temperature response information in individual offices, the setback periods were increased by 50% for the heating and 24% for the cooling season. In addition, the individually selected and learned thermostat setpoints lowered the heating setpoint by 2°C and increased the cooling setpoints by 2°C from the default 22°C setpoint. As a result, the duration the radiant panels were operational and the volume of chilled-air use in the offices controlled by the adaptive algorithm were 75% lower than the offices controlled by the default temperature control algorithm. In an effort to gather unbiased information about the user comfort, in lieu of conducting surveys, the institutional historical complaint logs, the frequency of the thermostat keypress actions, the indoor air temperature and the CO_2 concentration during the occupied hours were analyzed. However, no evidence was found

that the user comfort was adversely affected in the offices controlled by the adaptive temperature control algorithm.

In an effort to analyze the influence of a range of potential technician decisions in the controls programming for the default and the adaptive learning temperature control algorithms, a simulation-based investigation was conducted by employing a physical model of the eight offices. It is demonstrated that simple operational decisions (e.g., setpoint/setback temperature choices and schedules) that are made with little guidance account for great uncertainty in heating/cooling load intensities and user comfort. It was found that the adaptive temperature control algorithm not only tends to reduce the expected heating and cooling loads but also it reduces the risk of poor operational decisions leading to discomfort conditions and/or extravagance in heating and cooling end uses.

The indoor illuminance control algorithm was first implemented in the controls laboratory. The laboratory experiments were conducted such that the self-adaptive lighting and blinds control algorithm was employed only on the alternating weeks. After a year of data acquisition, it was observed that the total duration the lights remained on with the adaptive control algorithm was 22% less than it was with the default control algorithm (auto-off with vacancy detection and manual on lighting, and manually controlled motorized blinds). In addition, with the adaptive control algorithm, the mean blind occlusion rate during occupancy reduced by 6% – from 18% to 12%. The users accepted 94% of the automated light switch off decisions and 86% of the automated blinds opening decisions.

The adaptive lighting learning controls algorithm was then implemented in five private offices. Because the blinds were not automated, the control algorithm's role reduced to switching off the electric lighting when there was adequate daylight or when the office was vacated. The algorithm mostly eliminated the lighting use during

the occupied hours when the ceiling illuminance exceeded 250 lux. The occupants accepted 93% of the automated light switch off decisions upon the detection of adequate daylight.

In an effort to estimate the energy savings potential of the adaptive lighting and blinds control algorithm in this field implementation, a simulation-based investigation was conducted with the physical daylighting model of these offices and with the stochastic occupant behaviour and presence models – derived using the data gathered from these offices. Based on the simulation results, the addition of the adaptive occupant-learning lighting control algorithm could reduce the lighting use duration by 30%. It was estimated that the lighting use savings would have been about 40%, if we were to incorporate the adaptive occupant-learning blinds and lighting control algorithm.

Note that the intellectual property rights for the algorithms in Matlab, GCL+, and EnergyPlus EMS languages were exclusively assigned to Delta Controls for the first six months following the publication of this thesis. The researcher can be contacted for academic use.

Contributions:

The adaptive learning indoor temperature and illuminance control tool developed, implemented, and tested in this thesis is a promising method to substantially improve energy efficiency and user comfort in commercial buildings. This is a major contribution towards the commonly stated Canadian objective for all new buildings to be net-zero energy by 2030. Because the adaptive indoor climate control tool will largely use the existing sensor network and the BAS, retrofits using it will be relatively non-invasive and low-cost, and could be used to retrofit many of Canada's commercial buildings.

In addition, the concept of adaptive comfort and its relationship with adaptive

behaviours have been elaborated and practically exploited in a building's operation. The findings of the occupant behaviour and comfort research have been acknowledged in the design of the indoor climate control algorithms.

Established data analytics methodologies such as the recursive Bayesian filters were introduced to the building science community and were blended with the building physics – e.g., CO₂ and heat transfer and generation. Computational and analytical limitations of the building controllers were acknowledged, and systematic ways to select parsimonious algorithms were introduced.

For the first time, a self-adaptive learning model-based predictive control algorithm was developed and a model-based control algorithm was implemented inside a commercial controller. This is a very fundamental step for the buildings research community; as the MPC has long been criticised for being impractical outside the research community.

7.1 Future work

The future work recommendations were categorized in two groups: incremental and fundamental research steps.

Incremental research steps and unresolved issues:

1. Occupants' adaptive behaviours and presence were analyzed longitudinally from a small sample size – only ten to twenty occupants. Consequently, the diversity aspects could not be thoroughly characterized. Future research should study occupants' behaviour and presence with independent datasets from different buildings. However, it is worth noting that the sample size was not atypical in the state-of-the-art occupant behaviour research [22].

2. Occupants' adaptive behaviours were used to deduce the conditions disliked by

the occupants and adapt the operating setpoints accordingly. The link between the timing/frequency of adaptive behaviours and the user's comfort needs to be further investigated by using a larger dataset.

3. It was anecdotally demonstrated that adaptive controls by tailoring the setpoints and schedules to individual preferences diversify the heating and cooling demand profiles in individual zones. However, the impact on the aggregate loads of the system and plant level equipment were not studied. Future work should involve large scale implementation campaigns to better understand the role of adaptive controls on the load diversity (and peak load reduction and shifting). Different building types and orientations should be selected for this large scale implementation.

4. The implementation should be accompanied with high resolution metering of the end-uses so that, energy savings (e.g., savings in pumps, fans, chillers, and boilers) can be quantified accurately – along with the heating and cooling load reduction potential. An issue in using a portion of a building to test a temperature control algorithm's performance is the inevitable heat exchange between the test rooms and the rest of the building. Isolated test floors can be formed to minimize the interactions between the building and the test floors.

5. In this study, by employing distributed artificial intelligence in zone controllers, personal information such as the occupancy patterns has been collected. Although the project was conducted after receiving approval from the research ethics board, real-life applications may bring about privacy concerns. Future research should investigate the user perception about the collection and use of personal data inside zone controllers.

6. This study was conducted by observing the behaviour and occupancy patterns mainly in private office spaces (except that the controls laboratory was a shared office space). Future research should study the applicability of adaptive controls in multi-occupant spaces.

7. The adaptive controls framework presented in this thesis puts forward a bottom-up approach in which the terminal equipment utilization efficiency was improved. However, the terminal equipment use is subject to externalities pertaining to the system and plant level equipment availability schedules and capacity. These externalities were not explicitly considered in this study.

8. Conducting a survey, despite being a potential risk to cause a deviation in behaviour patterns during the experiments, could provide useful information with regards to the occupants' satisfaction with the overall adaptive control scheme upon the termination of the field experiments. Future work is planned to conduct a field survey following the completion of the next set of experiments.

Fundamental research steps:

1. The inverse modelling methods introduced in this thesis can be extended to identify the signature relative humidity and CO₂ response in individual thermal zones. These inverse modelling methods can be used to develop adaptive fault detection and diagnostics algorithms. The inverse models can estimate immeasurable or unmeasured performance indicators (e.g., effective thermal resistance, air infiltration rate). Similar to the algorithms shown in Chapter 4, the fault detection and diagnostics algorithms can learn the normal response of the building, and when an anomaly (or suboptimality) is observed, it can notify the operator and/or self-correct the faults.

2. The adaptive controls algorithms can be considered as a leap forward in the vision of autonomous buildings. A logical step in the same direction would be to develop auto-commissioning routines to inspect various building sensors and actuators systematically at regular intervals. For example, in a closed-loop lighting controls system, the effect of switching on/off the lights on the indoor illuminance sensor can be identified by changing the lighting state automatically at different instances prior to occupancy.

3. Contextual factors are occupants' habits, attitudes, and beliefs that affect the way we interact with building components. These factors may influence occupants' behaviours substantially [43] – which can inhibit our ability to infer their comfort preferences by looking at their behaviours. For example, the seasonal changes in occupants' clothing – and thus thermal comfort – preferences can be affected by the default temperature setpoints imposed by the automation. For example, during the summer occupants may choose to wear heavier clothing than they would normally prefer knowing that it will be cold in their office. Similarly, in private office spaces, occupants may have more flexibility to alter their clothing assembly than in shared office spaces. Laboratory scale comfort and behaviour studies can be conducted by carefully changing these contextual factors.

4. In lieu of using the learned occupancy parameters inside an adaptive controller, they can be used inside an innovative operator interface. Similarly, the parameters of the building temperature response learning algorithm can be used to inform the operator about the expected length of the setback-to-setpoint transition period. The parameters of the thermostat learning algorithm can be used to assist operators in selecting appropriate operating setpoints. As such, the traditional decision making process of the operators can be supported. Future research should investigate the potential of developing better operator control interfaces. Furthermore, ease of access and usability of the control interfaces play an important role in occupants' behaviour patterns as well. Therefore, the viability of using the information recursively derived by the learning algorithms in user interfaces should be studied.

5. The methods developed in this thesis target inverse modelling of the occupant preferences and the building physics using low-cost sensing technologies in the existing office buildings. The advances in new integrated sensing technologies will likely make some of the algorithms obsolete. For example, adaptive occupancy detection

algorithm for passive-infrared sensors can become outdated – with the development and widespread use of computer vision-based occupancy sensors. The adaptive indoor climate control research should constantly explore ways to incorporate the advances in the sensing and metering technologies into the betterment of buildings' operation.

List of References

- [1] EPA, “An office building occupants guide to indoor air quality,” tech. rep., United States Environmental Protection Agency, 1997.
- [2] OECD, “OECD green growth studies: Energy,” tech. rep., Organisation for Economic Co-operation and Development, 2011.
- [3] NRCan, “Energy use data handbook 1990 to 2010,” tech. rep., Natural Resources Canada, 2013.
- [4] DOE, “Buildings energy data book,” tech. rep., United States Department of Energy, 2011.
- [5] S. Katipamula and M. R. Brambley, “Review article: Methods for fault detection, diagnostics, and prognostics for building systems – a review, Part II,” *HVAC & R Research*, vol. 11, no. 2, pp. 169–187, 2005.
- [6] E. H. Mathews, C. P. Botha, D. C. Arndt, and A. Malan, “HVAC control strategies to enhance comfort and minimise energy usage,” *Energy and Buildings*, vol. 33, no. 8, pp. 853–863, 2001.
- [7] B. Bordass, A. Leaman, and P. Ruyssevelt, “Assessing building performance in use 5: conclusions and implications,” *Building Research & Information*, vol. 29, no. 2, pp. 144–157, 2001.
- [8] ASHRAE, *Handbook of Fundamentals*. ASHRAE, Atlanta, GA, 2013.
- [9] M. R. Brambley, N. Fernandez, W. Wang, K. A. Cort, H. Cho, H. Ngo, and J. Goddard, “Self-correcting controls for VAV system faults,” 2011.
- [10] L. Golden, “Flexible work schedules which workers get them?,” *American Behavioral Scientist*, vol. 44, no. 7, pp. 1157–1178, 2001.
- [11] T. M. McMenamin, “Time to work: recent trends in shift work and flexible schedules,” *Monthly Lab. Rev.*, vol. 130, p. 3, 2007.

- [12] I. U. Zeytinoglu, B. C. Gordon, and S. L. Mann, "Flexibility: Whose choice is it anyway?," *Relations Industrielles / Industrial Relations*, vol. 64, no. 4, pp. 555–574, 2009.
- [13] A. Mahdavi, A. Mohammadi, E. Kabir, and L. Lambeva, "Occupants' operation of lighting and shading systems in office buildings," *Journal of Building Performance Simulation*, vol. 1, no. 1, pp. 57–65, 2008.
- [14] J. A. D. III and D. W. Nutter, "Occupancy diversity factors for common university building types," *Energy and Buildings*, vol. 42, no. 9, pp. 1543 – 1551, 2010.
- [15] K. Sun, D. Yan, T. Hong, and S. Guo, "Stochastic modeling of overtime occupancy and its application in building energy simulation and calibration," *Building and Environment*, vol. 79, pp. 1 – 12, 2014.
- [16] H. Gunay, W. O'Brien, I. Beausoleil-Morrison, and J. Bursill, "Implementation of an adaptive occupancy and building learning temperature setback algorithm," *ASHRAE Transactions*, vol. 122, no. 1, 2016.
- [17] S. Schiavon and K. H. Lee, "Dynamic predictive clothing insulation models based on outdoor air and indoor operative temperatures," *Building and Environment*, vol. 59, pp. 250 – 260, 2013.
- [18] R. J. de Dear and G. S. Brager, "Thermal comfort in naturally ventilated buildings: revisions to ASHRAE standard 55," *Energy and Buildings*, vol. 34, no. 6, pp. 549 – 561, 2002. Special Issue on Thermal Comfort Standards.
- [19] C. Federspiel, R. Martin, and H. Yan, "Thermal comfort models and complaint frequencies."
- [20] ASHRAE, *Standard 55-2013 – Thermal Environmental Conditions for Human Occupancy*. ASHRAE, Atlanta, GA, 2013.
- [21] P. O. Fanger, *Thermal comfort. Analysis and applications in environmental engineering*. Denmark: Krieger, 1970.
- [22] H. B. Gunay, W. O'Brien, and I. Beausoleil-Morrison, "A critical review of observation studies, modeling, and simulation of adaptive occupant behaviors in offices," *Building and Environment*, vol. 70, pp. 31–47, 2013.

- [23] M. A. Humphreys and J. F. Nicol, "Understanding the adaptive approach to thermal comfort," *Ashrae Transactions*, vol. 104, p. 991, 1998.
- [24] A. Leaman and B. Bordass, "Productivity in buildings: the killer variables," *Building Research & Information*, vol. 27, no. 1, pp. 4–19, 1999.
- [25] L. Preller, T. Zweers, B. Brunekreef, and J. Boleij, "Sick leave due to work-related health complaints among office workers in Netherlands," 1990.
- [26] A. Leaman and B. Bordass, "Keeping occupants satisfied," *Energy and Environmental Management*, pp. 23–27, 2000.
- [27] M. Fountain, G. Brager, and R. de Dear, "Expectations of indoor climate control," *Energy and Buildings*, vol. 24, no. 3, pp. 179 – 182, 1996.
- [28] D. P. Wyon, "Enhancing productivity while reducing energy use in buildings," in *Proceedings of the E-Vision 2000 Conference*, pp. 11–13.
- [29] S. Karjalainen, "Usability guidelines for room temperature controls," *Intelligent Buildings International*, vol. 2, no. 2, pp. 85–97, 2010.
- [30] S. Karjalainen, "Thermal comfort and use of thermostats in Finnish homes and offices," *Building and Environment*, vol. 44, no. 6, pp. 1237 – 1245, 2009.
- [31] W. Kempton, D. Feuermann, and A. E. McGarity, "I always turn it on super: user decisions about when and how to operate room air conditioners," *Energy and Buildings*, vol. 18, no. 34, pp. 177 – 191, 1992.
- [32] T. Salsbury, "A survey of control technologies in the building automation industry," in *16th IFAC World Congress*, pp. 1396–1407, 2005.
- [33] T. I. Salsbury, *The Smart Building*, pp. 1079–1093. Springer, 2009.
- [34] H. B. Gunay, W. O'Brien, and I. Beausoleil-Morrison, "Development of an occupancy-learning algorithm for terminal heating and cooling units," *Building and Environment*, vol. 93, pp. 71–85, 2015.
- [35] J. A. Veitch, D. W. Hine, and R. Gifford, "End users knowledge, beliefs, and preferences for lighting," *Journal of Interior Design*, vol. 19, no. 2, pp. 15–26, 1993.
- [36] J. Heerwagen and D. Heerwagen, "Lighting and psychological comfort," *Lighting Design and Application*, vol. 16, no. 4, pp. 47–51, 1986.

- [37] C. F. Reinhart, "Lightswitch-2002: a model for manual and automated control of electric lighting and blinds," *Solar Energy*, vol. 77, no. 1, pp. 15–28, 2004.
- [38] S. Pigg, M. Eilers, and J. Reed, "Behavioral aspects of lighting and occupancy sensors in private offices: a case study of a university office building," *ACEEE 1996 Summer Study on Energy Efficiency in Buildings*, 1996.
- [39] W. O'Brien, K. Kapsis, and A. K. Athienitis, "Manually-operated window shade patterns in office buildings: A critical review," *Building and Environment*, vol. 60, pp. 319–338, 2013.
- [40] F. Haldi and D. Robinson, "Adaptive actions on shading devices in response to local visual stimuli," *Journal of Building Performance Simulation*, vol. 3, no. 2, pp. 135–153, 2010.
- [41] G. Newsham, "Manual control of window blinds and electric lighting: implications for comfort and energy consumption," *Indoor and Built Environment*, vol. 3, no. 3, pp. 135–144, 1994.
- [42] B. Bordass, R. Cohen, M. Standeven, and A. Leaman, "Assessing building performance in use 2: technical performance of the probe buildings," *Building Research & Information*, vol. 29, no. 2, pp. 103–113, 2001.
- [43] W. O'Brien and H. B. Gunay, "The contextual factors contributing to occupants' adaptive comfort behaviors in offices—a review and proposed modeling framework," *Building and Environment*, vol. 77, pp. 77–87, 2014.
- [44] J. Heerwagen, "Adaptations and coping: Occupant response to discomfort in energy efficient buildings," 1992.
- [45] C. F. Reinhart and K. Voss, "Monitoring manual control of electric lighting and blinds," *Lighting Research and Technology*, vol. 35, no. 3, pp. 243–258, 2003.
- [46] E. Vine, E. Lee, R. Clear, D. DiBartolomeo, and S. Selkowitz, "Office worker response to an automated venetian blind and electric lighting system: a pilot study," *Energy and buildings*, vol. 28, no. 2, pp. 205–218, 1998.
- [47] M. Lee, C. M. Carswell, W. Seidelman, and M. Sublette, "Green expectations the story of a customizable lighting control panel designed to reduce energy use," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 57, pp. 1353–1357, SAGE Publications, 2013.

- [48] H. B. Gunay, W. O'Brien, I. Beausoleil-Morrison, and B. Huchuk, "On adaptive occupant-learning window blind and lighting controls," *Building Research & Information*, vol. 42, no. 6, pp. 739–756, 2014.
- [49] M. J. O'Neill, "Work space adjustability, storage, and enclosure as predictors of employee reactions and performance," *Environment and Behavior*, vol. 26, no. 4, pp. 504–526, 1994.
- [50] A. Leaman and B. Bordass, "Assessing building performance in use 4: the probe occupant surveys and their implications," *Building Research & Information*, vol. 29, no. 2, pp. 129–143, 2001.
- [51] P. Vaidya, T. McDougall, J. Steinbock, J. Douglas, and D. Eijadi, "What's wrong with daylighting? Where it goes wrong and how users respond to failure," *Proceedings of ACEEE 2004 Summer Study*, 2004.
- [52] R. J. Cole and Z. Brown, "Reconciling human and automated intelligence in the provision of occupant comfort," *Intelligent Buildings International*, vol. 1, no. 1, pp. 39–55, 2009.
- [53] A. D. Galasiu and J. A. Veitch, "Occupant preferences and satisfaction with the luminous environment and control systems in daylit offices: a literature review," *Energy and Buildings*, vol. 38, no. 7, pp. 728–742, 2006.
- [54] M. Paciuk, *The role of personal control of the environment in thermal comfort and satisfaction at the workplace*. PhD thesis, University of Wisconsin-Milwaukee, 1989.
- [55] M. A. ul Haq, M. Y. Hassan, H. Abdullah, H. A. Rahman, M. P. Abdullah, F. Hussin, and D. M. Said, "A review on lighting control technologies in commercial buildings, their performance and affecting factors," *Renewable and Sustainable Energy Reviews*, vol. 33, pp. 268–279, 2014.
- [56] D. Maniccia, B. Rutledge, M. S. Rea, and W. Morrow, "Occupant use of manual lighting controls in private offices," *Journal of the Illuminating Engineering Society*, vol. 28, no. 2, pp. 42–56, 1999.
- [57] K. Van Den Wymelenberg and M. Inanici, "A critical investigation of common lighting design metrics for predicting human visual comfort in offices with daylight," *Leukos*, vol. 10, no. 3, pp. 145–164, 2014.

- [58] J. A. Jakubiec and C. F. Reinhart, "The adaptive zone—a concept for assessing discomfort glare throughout daylit spaces," *Lighting Research and Technology*, vol. 44, no. 2, pp. 149–170, 2012.
- [59] J. Veitch, E. Dikel, G. Burns, and S. Mancini, "Individual control over light source spectrum: effects on perception and cognition," *27th Session of the Commission Internationale de l'Éclairage*, pp. 213–218, 2011.
- [60] S. Gilani, W. O'Brien, H. B. Gunay, and J. S. Carrizo, "Use of dynamic occupant behavior models in the building design and code compliance processes," *Energy and Buildings*, 2015.
- [61] M. Chiogna, R. Albatici, and A. Frattari, "Electric lighting at the workplace in offices: Efficiency improvement margins of automation systems," *Lighting Research and Technology*, vol. 45, no. 5, pp. 550–567, 2013.
- [62] E. S. Lee, L. L. Fernandes, B. Coffey, A. McNeil, R. Clear, T. Webster, F. Bauman, D. Dickeroff, D. Heinzerling, and T. Hoyt, "A post-occupancy monitored evaluation of the dimmable lighting, automated shading, and underfloor air distribution system in the new york times building," *Berkeley National Laboratory*, 2013.
- [63] T. Moore, D. Carter, and A. Slater, "Long-term patterns of use of occupant controlled office lighting," *Lighting Research and Technology*, vol. 35, no. 1, pp. 43–57, 2003.
- [64] L. Roche, "Summertime performance of an automated lighting and blinds control system," *Lighting Research and Technology*, vol. 34, no. 1, pp. 11–25, 2002.
- [65] J. Mardaljevic, L. Heschong, and E. Lee, "Daylight metrics and energy savings," *Lighting Research and Technology*, vol. 41, no. 3, pp. 261–283, 2009.
- [66] T. Hilliard, L. Swan, M. Kavacic, Z. Qin, and S. DeWolf, "Using building simulation software to quantify the savings of a model predictive control implementation," in *eSim 2016. Hamilton: IBPSA Canada*, pp. 13–21, IBPSA Canada, 2016.
- [67] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, *et al.*,

- “Energyplus: creating a new-generation building energy simulation program,” *Energy and buildings*, vol. 33, no. 4, pp. 319–331, 2001.
- [68] R. Lamberts and J. Hensen, *Building performance simulation for design and operation*. New York;London;: Spon Press, 2011.
- [69] C. W. Energy, “Engineering data sets (cweeds); cwec weather files,” *CWEEDS Appendix D*.
- [70] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, “Occupancy-driven energy management for smart building automation,” in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pp. 1–6, ACM, 2010.
- [71] B. Roisin, M. Bodart, A. Deneyer, and P. Dherdt, “Lighting energy savings in offices using different control systems and their real consumption,” *Energy and Buildings*, vol. 40, no. 4, pp. 514–523, 2008.
- [72] R. H. Dodier, G. P. Henze, D. K. Tiller, and X. Guo, “Building occupancy detection through sensor belief networks,” *Energy and buildings*, vol. 38, no. 9, pp. 1033–1043, 2006.
- [73] Z. Nagy, F. Y. Yong, M. Frei, and A. Schlueter, “Occupant centered lighting control for comfort and energy efficient building operation,” *Energy and Buildings*, vol. 94, pp. 100–108, 2015.
- [74] H. B. Gunay, A. Fuller, W. OBrien, and I. Beausoleil-Morrison, “Detecting occupants' presence in office spaces: a case study,” in *eSim 2016. Hamilton: IBPSA Canada*, pp. 753–763, IBPSA Canada, 2016.
- [75] D. Maniccia and X. Luan, “Methods for assessing the maintained and initial detection performance of occupancy sensors,” *Journal of the Illuminating Engineering Society*, vol. 23, no. 2, pp. 108–115, 1994.
- [76] Maniccia, “Occupancy sensors: Motion sensors for lighting control,” tech. rep., Rensselaer Polytechnic Institute, 1997.
- [77] ASHRAE, “Standard 62.1-2013 – ventilation for acceptable indoor air quality,” 2013.
- [78] H. A. Aglan, “Predictive model for CO₂ generation and decay in building envelopes,” *Journal of applied physics*, vol. 93, no. 2, pp. 1287–1290, 2003.

- [79] S. Wang, J. Burnett, and H. Chong, “Experimental validation of CO₂-based occupancy detection for demand-controlled ventilation,” *Indoor and Built Environment*, vol. 8, no. 6, pp. 377–391, 1999.
- [80] A. Arora, M. Amayri, V. Badarla, S. Ploix, and S. Bandyopadhyay, “Occupancy estimation using non intrusive sensors in energy efficient buildings,” in *14th Conference of International Building Performance Simulation Association, Hyderabad, India*, pp. 1441–1448, International Building Performance Simulation Association, 2015.
- [81] B. Dong and B. Andrews, “Sensor-based occupancy behavioral pattern recognition for energy and comfort management in intelligent buildings,” in *Proceedings of Building Simulation*, pp. 1444–1451, 2009.
- [82] K. P. Lam, M. Höynck, B. Dong, B. Andrews, Y.-S. Chiou, R. Zhang, D. Benitez, J. Choi, *et al.*, “Occupancy detection through an extensive environmental sensor network in an open-plan office building,” *IBPSA Building Simulation*, vol. 145, pp. 1452–1459, 2009.
- [83] B. Dong, B. Andrews, K. P. Lam, M. Höynck, R. Zhang, Y.-S. Chiou, and D. Benitez, “An information technology enabled sustainability test-bed (itest) for occupancy detection through an environmental sensing network,” *Energy and Buildings*, vol. 42, no. 7, pp. 1038–1046, 2010.
- [84] T. A. Nguyen and M. Aiello, “Beyond indoor presence monitoring with simple sensors,” in *PECCS*, pp. 5–14, 2012.
- [85] K. Padmanabh, A. Malikarjuna V, S. Sen, S. P. Katru, A. Kumar, S. K. Vupala, S. Paul, *et al.*, “isense: a wireless sensor network based conference room management system,” in *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pp. 37–42, ACM, 2009.
- [86] V. L. Erickson, M. Carreira-Perpinan, and A. E. Cerpa, “Observe: Occupancy-based system for efficient reduction of HVAC energy,” in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pp. 258–269, IEEE, 2011.
- [87] L. J. Brackney, A. R. Florita, A. C. Swindler, L. G. Polese, and G. A. Brunnemann, “Design and performance of an image processing occupancy sensor,” in *Proceedings: The Second International Conference on Building Energy and*

Environment 2012987 Topic 10. Intelligent buildings and advanced control techniques, Citeseer, 2012.

- [88] H. Ramoser, T. Schlogl, C. Beleznai, M. Winter, and H. Bischof, “Shape-based detection of humans for video surveillance applications,” in *Image Processing, 2003. ICIIP 2003. Proceedings. 2003 International Conference on*, vol. 3, pp. III–1013, IEEE, 2003.
- [89] E. Hailemariam, R. Goldstein, R. Attar, and A. Khan, “Real-time occupancy detection using decision trees with multiple sensor types,” in *Proceedings of the 2011 Symposium on Simulation for Architecture and Urban Design*, pp. 141–148, Society for Computer Simulation International, 2011.
- [90] N. Li and B. Becerik-Gerber, “Performance-based evaluation of rfid-based indoor location sensing solutions for the built environment,” *Advanced Engineering Informatics*, vol. 25, no. 3, pp. 535–546, 2011.
- [91] M. Philipose, S. Consolvo, T. Choudhury, K. Fishkin, M. Perkowitz, I. Fox, H. Kautz, and D. Patterson, “Fast, detailed inference of diverse daily human activities,” *Demonstrations at Ubicomp*, vol. 7, p. 2004, 2004.
- [92] M. Milenkovic and O. Amft, “Recognizing energy-related activities using sensors commonly installed in office buildings,” *Procedia Computer Science*, vol. 19, pp. 669–677, 2013.
- [93] H. B. Gunay, “Supplemental: GCL+ scripts for the control algorithms available at <https://goo.gl/asnm9i>,” 2016.
- [94] A. Mahdavi and F. Tahmasebi, “The deployment-dependence of occupancy-related models in building performance simulation,” *Energy and Buildings*, vol. 117, pp. 313 – 320, 2016.
- [95] B. Bordass, R. Cohen, M. Standeven, and A. Leaman, “Assessing building performance in use 2: technical performance of the probe buildings,” *Building Research & Information*, vol. 29, no. 2, pp. 103–113, 2001.
- [96] D. o. L. US, “Workers on flexible and shift schedules in May 2004,” tech. rep., Bureau of Labor Statistics, 2005.
- [97] E. L. Kelly, P. Moen, and E. Tranby, “Changing workplaces to reduce work-family conflict schedule control in a white-collar organization,” *American Sociological Review*, vol. 76, no. 2, pp. 265–290, 2011.

- [98] Z. Yang and B. Becerik-Gerber, "Modeling personalized occupancy profiles for representing long term patterns by using ambient context," *Building and Environment*, vol. 78, pp. 23–35, 2014.
- [99] Z. Yang and B. Becerik-Gerber, "The coupled effects of personalized occupancy profile based HVAC schedules and room reassignment on building energy use," *Energy and Buildings*, vol. 78, pp. 113–122, 2014.
- [100] A. Mahdavi, "Patterns and implications of user control actions in buildings," *Indoor and Built Environment*, 2009.
- [101] F. Haldi and D. Robinson, "The impact of occupants' behaviour on building energy demand," *Journal of Building Performance Simulation*, vol. 4, no. 4, pp. 323–338, 2011.
- [102] H. B. Gunay, W. O'Brien, and I. Beausoleil-Morrison, "Implementation and comparison of existing occupant behaviour models in energyplus," *Journal of Building Performance Simulation*, pp. 1–46, 2015.
- [103] J. Page, D. Robinson, N. Morel, and J.-L. Scartezzini, "A generalised stochastic model for the simulation of occupant presence," *Energy and buildings*, vol. 40, no. 2, pp. 83–98, 2008.
- [104] C. Wang, D. Yan, and Y. Jiang, "A novel approach for building occupancy simulation," in *Building simulation*, vol. 4, pp. 149–167, Springer, 2011.
- [105] W. Parys, D. Saelens, and H. Hens, "Coupling of dynamic building simulation with stochastic modelling of occupant behaviour in offices—a review-based integrated methodology," *Journal of Building Performance Simulation*, vol. 4, no. 4, pp. 339–358, 2011.
- [106] D. Wang, C. C. Federspiel, and F. Rubinstein, "Modeling occupancy in single person offices," *Energy and buildings*, vol. 37, no. 2, pp. 121–126, 2005.
- [107] L. J. Lo and A. Novoselac, "Localized air-conditioning with occupancy control in an open office," *Energy and Buildings*, vol. 42, no. 7, pp. 1120–1128, 2010.
- [108] B. Dong, K. P. Lam, and C. Neuman, "Integrated building control based on occupant behavior pattern detection and local weather forecasting," in *Twelfth International IBPSA Conference. Sydney: IBPSA Australia*, pp. 14–17, Cite-seer, 2011.

- [109] J. S. Bendat and A. G. Piersol, *Random data: analysis and measurement procedures*, vol. 729. John Wiley & Sons, 2011.
- [110] J. F. Nicol and M. A. Humphreys, “A stochastic approach to thermal comfort–occupant behavior and energy use in buildings,” *ASHRAE transactions*, vol. 110, no. 2, 2004.
- [111] A. Guillemin and N. Morel, “An innovative lighting controller integrated in a self-adaptive building control system,” *Energy and buildings*, vol. 33, no. 5, pp. 477–487, 2001.
- [112] F. Haldi and D. Robinson, “Interactions with window openings by office occupants,” *Building and Environment*, vol. 44, no. 12, pp. 2378–2395, 2009.
- [113] D. Yan, W. OBrien, T. Hong, X. Feng, H. B. Gunay, F. Tahmasebi, and A. Mahdavi, “Occupant behavior modeling for building performance simulation: current state and future challenges,” *Energy and Buildings*, vol. 107, pp. 264–278, 2015.
- [114] R. V. Andersen, B. W. Olesen, and J. Toftum, “Modelling occupants heating set-point preferences,” *Proceedings of Building Simulation, Sydney*, 2011.
- [115] S. DOca, V. Fabi, S. P. Corgnati, and R. K. Andersen, “Effect of thermostat and window opening occupant behavior models on energy use in homes,” in *Building Simulation*, vol. 7, pp. 683–694, Springer, 2014.
- [116] P. Hoes, J. Hensen, M. Loomans, B. De Vries, and D. Bourgeois, “User behavior in whole building simulation,” *Energy and Buildings*, vol. 41, no. 3, pp. 295–302, 2009.
- [117] F. Haldi and D. Robinson, “On the behaviour and adaptation of office occupants,” *Building and environment*, vol. 43, no. 12, pp. 2163–2177, 2008.
- [118] H. B. Gunay, W. OBrien, I. Beausoleil-Morrison, S. DOca, and S. P. Corgnati, “On modelling and simulation of occupant models,” in *14th Conference of International Building Performance Simulation Association, Hyderabad, India*, pp. 1501–1508, International Building Performance Simulation Association, 2015.
- [119] H. B. Gunay, W. OBrien, and I. Beausoleil-Morrison, “A toolkit for developing data-driven occupant behaviour and presence models,” in *eSim 2016. Hamilton: IBPSA Canada*, pp. 743–752, IBPSA Canada, 2016.

- [120] H. B. Gunay, W. O'Brien, I. Beausoleil-Morrison, R. Goldstein, S. Breslav, and A. Khan, "Coupling stochastic occupant models to building performance simulation using the discrete event system specification formalism," *Journal of Building Performance Simulation*, vol. 7, no. 6, pp. 457–478, 2014.
- [121] H. B. Rijal, P. Tuohy, F. Nicol, M. A. Humphreys, A. Samuel, and J. Clarke, "Development of an adaptive window-opening algorithm to predict the thermal comfort, energy use and overheating in buildings," *Journal of Building Performance Simulation*, vol. 1, no. 1, pp. 17–30, 2008.
- [122] V. Inkarojrit and G. Paliaga, "Indoor climatic influences on the operation of windows in a naturally ventilated building," in *Proceedings of the 21st international conference on passive and low energy architecture*, pp. 19–22, 2004.
- [123] M. Foster and T. Oreszczyn, "Occupant control of passive systems: the use of venetian blinds," *Building and Environment*, vol. 36, no. 2, pp. 149–155, 2001.
- [124] P. Warren and L. Parkins, "Window-opening behaviour in office buildings," *Building Services Engineering Research and Technology*, vol. 5, no. 3, pp. 89–101, 1984.
- [125] J. A. Clarke, I. Macdonald, and J. F. Nicol, "Predicting adaptive responses-simulating occupied environments," 2006.
- [126] F. Haldi and D. Robinson, "Modelling occupants personal characteristics for thermal comfort prediction," *International journal of biometeorology*, vol. 55, no. 5, pp. 681–694, 2011.
- [127] V. Inkarojrit, "Monitoring and modelling of manually-controlled venetian blinds in private offices: a pilot study," *Journal of Building Performance Simulation*, vol. 1, no. 2, pp. 75–89, 2008.
- [128] J. F. Nicol, "Characterising occupant behaviour in buildings: towards a stochastic model of occupant use of windows, lights, blinds, heaters and fans," in *Proceedings of the seventh international IBPSA conference, Rio*, vol. 2, pp. 1073–1078, 2001.
- [129] H. B. Rijal, P. Tuohy, M. A. Humphreys, J. F. Nicol, A. Samuel, and J. Clarke, "Using results from field surveys to predict the effect of open windows on thermal comfort and energy use in buildings," *Energy and buildings*, vol. 39, no. 7, pp. 823–836, 2007.

- [130] Y. Zhang and P. Barrett, “Factors influencing the occupants window opening behaviour in a naturally ventilated office building,” *Building and Environment*, vol. 50, pp. 125–134, 2012.
- [131] S. Theodoridis, *Machine learning: a Bayesian and optimization perspective*. Academic Press, 2015.
- [132] S. S. Haykin, S. S. Haykin, S. S. Haykin, and S. S. Haykin, *Neural networks and learning machines*, vol. 3. Pearson Education Upper Saddle River, 2009.
- [133] Z. Wang, D. Yi, X. Duan, J. Yao, and D. Gu, *Measurement data modeling and parameter estimation*. CRC Press, 2011.
- [134] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, “Use of model predictive control and weather forecasts for energy efficient building climate control,” *Energy and Buildings*, vol. 45, pp. 15–27, 2012.
- [135] A. I. Dounis and C. Caraiscos, “Advanced control systems engineering for energy and comfort management in a building environmenta review,” *Renewable and Sustainable Energy Reviews*, vol. 13, no. 6, pp. 1246–1261, 2009.
- [136] S. Katipamula and M. R. Brambley, “Review article: methods for fault detection, diagnostics, and prognostics for building systemsa review, part i,” *HVAC&R Research*, vol. 11, no. 1, pp. 3–25, 2005.
- [137] A. Saberi Derakhtenjani, J. A. Candanedo, Y. Chen, V. R. Dehkordi, and A. K. Athienitis, “Modeling approaches for the characterization of building thermal dynamics and model-based control: A case study,” *Science and Technology for the Built Environment*, vol. 21, no. 6, pp. 824–836, 2015.
- [138] A. Afram and F. Janabi-Sharifi, “Theory and applications of HVAC control systems—a review of model predictive control (MPC),” *Building and Environment*, vol. 72, pp. 343–355, 2014.
- [139] J. A. Candanedo, V. R. Dehkordi, A. Saberi-Derakhtenjani, and A. K. Athienitis, “Near-optimal transition between temperature setpoints for peak load reduction in small buildings,” *Energy and Buildings*, vol. 87, pp. 123–133, 2015.
- [140] J. A. Clarke, *Energy simulation in building design*. Routledge, 2001.

- [141] S. F. Fux, A. Ashouri, M. J. Benz, and L. Guzzella, “EKF based self-adaptive thermal model for a passive house,” *Energy and Buildings*, vol. 68, pp. 811–817, 2014.
- [142] P. Bacher and H. Madsen, “Identifying suitable models for the heat dynamics of buildings,” *Energy and Buildings*, vol. 43, no. 7, pp. 1511–1522, 2011.
- [143] M. Kummert, P. Andr textsfle, and J. Nicolas, “Optimal heating control in a passive solar commercial building,” *Solar Energy*, vol. 69, pp. 103–116, 2001.
- [144] J. Penman, “Second order system identification in the thermal response of a working school,” *Building and Environment*, vol. 25, no. 2, pp. 105–110, 1990.
- [145] H. B. Gunay, W. O'Brien, I. Beausoleil-Morrison, P. Bisailon, and Z. Shi, “Development and implementation of control-oriented models for terminal heating and cooling units,” *Energy and Buildings*, vol. 121, pp. 78–91, 2016.
- [146] B. Lehmann, V. Dorer, and M. Koschenz, “Application range of thermally activated building systems tabs,” *Energy and buildings*, vol. 39, no. 5, pp. 593–598, 2007.
- [147] A. Martincevic, A. Starcic, and M. Vasak, “Parameter estimation for low-order models of complex buildings,” in *Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), 2014 IEEE PES*, pp. 1–6, IEEE, 2014.
- [148] M. Maasoumy, M. Razmara, M. Shahbakhti, and A. S. Vincentelli, “Handling model uncertainty in model predictive control for energy efficient buildings,” *Energy and Buildings*, vol. 77, pp. 377–392, 2014.
- [149] J. E. Braun and N. Chaturvedi, “An inverse gray-box model for transient building load prediction,” *HVAC&R Research*, vol. 8, no. 1, pp. 73–99, 2002.
- [150] T. Chen and A. Athienitis, “Investigation of practical issues in building thermal parameter estimation,” *Building and Environment*, vol. 38, no. 8, pp. 1027–1038, 2003.
- [151] B. Huchuk, H. B. Gunay, W. O'Brien, and C. A. Cruickshank, “Model-based predictive control of office window shades,” *Building Research & Information*, pp. 1–13, 2015.

- [152] H. B. Gunay, J. Bursill, B. Huchuk, W. O'Brien, and I. Beausoleil-Morrison, "Shortest-prediction-horizon model-based predictive control for individual offices," *Building and Environment*, vol. 82, pp. 408–419, 2014.
- [153] D. Coley and J. Penman, "Second order system identification in the thermal response of real buildings. paper ii: recursive formulation for on-line building energy management and control," *Building and Environment*, vol. 27, no. 3, pp. 269–277, 1992.
- [154] B. Huchuk, "Model-based predictive control of window shades," 2014.
- [155] H. B. Gunay, W. O'Brien, and I. Beausoleil-Morrison, "Control-oriented inverse modeling of the thermal characteristics in an office," *Science and Technology for the Built Environment*, 2016.
- [156] H. B. Gunay, "Supplemental: Verification of the sensor calibration at <https://goo.gl/uko37w>," 2016.
- [157] H. B. Gunay, W. OBrien, I. Beausoleil-Morrison, and S. Gilani, "Modeling plug-in equipment load patterns in private office spaces," *Energy and Buildings*, 2016.
- [158] N. Bergman, "Recursive bayesian estimation," *Department of Electrical Engineering, Linköping University, Linköping Studies in Science and Technology. Doctoral dissertation*, vol. 579, 1999.
- [159] J. V. Candy, "Bootstrap particle filtering," *IEEE Signal Processing Magazine*, vol. 24, pp. 73–85, July 2007.
- [160] A. S. Stordal, "Sequential monte carlo methods for bayesian filtering," 2008.
- [161] P. Bisailon, *Bayesian Model Selection and Parameter Estimation for Strongly Nonlinear Dynamical Systems*. PhD thesis, Carleton University Ottawa, 2013.
- [162] M. Khalil, A. Sarkar, and S. Adhikari, "Nonlinear filters for chaotic oscillatory systems," *Nonlinear Dynamics*, vol. 55, no. 1-2, pp. 113–137, 2009.
- [163] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *Signal Processing, IEEE Transactions on*, vol. 50, no. 2, pp. 174–188, 2002.
- [164] D. Gyalistras and O. Team, "Final report: Use of weather and occupancy forecasts for optimal building climate control (opticontrol)," *Zürich, Switzerland: ETH Zürich*, 2010.

- [165] A. Menezes, A. Cripps, R. A. Buswell, J. Wright, and D. Bouchlaghem, “Estimating the energy consumption and power demand of small power equipment in office buildings,” *Energy and Buildings*, vol. 75, pp. 199–209, 2014.
- [166] T. E. Kuhn, C. Bühler, and W. J. Platzer, “Evaluation of overheating protection with sun-shading systems,” *Solar Energy*, vol. 69, pp. 59–74, 2001.
- [167] N. G. Polson, J. R. Stroud, and P. Müller, “Practical filtering with sequential parameter learning,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 70, no. 2, pp. 413–428, 2008.
- [168] ASHRAE, “Standard 135: Bacnet a data communication protocol for building automation and control networks,” 2012.
- [169] J. Ma, J. Qin, T. Salsbury, and P. Xu, “Demand reduction in building energy systems based on economic model predictive control,” *Chemical Engineering Science*, vol. 67, no. 1, pp. 92–100, 2012.
- [170] P.-D. Moroşan, R. Bourdais, D. Dumur, and J. Buisson, “Building temperature regulation using a distributed model predictive control,” *Energy and Buildings*, vol. 42, no. 9, pp. 1445–1452, 2010.
- [171] J. A. Candanedo and A. K. Athienitis, “Predictive control of radiant floor heating and solar-source heat pump operation in a solar house,” *HVAC&R Research*, vol. 17, no. 3, pp. 235–256, 2011.
- [172] G. Huang, “Model predictive control of VAV zone thermal systems concerning bi-linearity and gain nonlinearity,” *Control Engineering Practice*, vol. 19, no. 7, pp. 700–710, 2011.
- [173] Q. Zou, J. Ji, S. Zhang, M. Shi, and Y. Luo, “Model predictive control based on particle swarm optimization of greenhouse climate for saving energy consumption,” in *World Automation Congress (WAC), 2010*, pp. 123–128, IEEE, 2010.
- [174] K. F. Fong, V. I. Hanby, and T. T. Chow, “A robust evolutionary algorithm for HVAC engineering optimization,” *HVAC&R Research*, vol. 14, no. 5, pp. 683–705, 2008.
- [175] L. dos Santos Coelho and V. C. Mariani, “Improved firefly algorithm approach applied to chiller loading for energy conservation,” *Energy and Buildings*, vol. 59, pp. 273–278, 2013.

- [176] T. Hilliard, M. Kavgic, and L. Swan, “Model predictive control for commercial buildings: trends and opportunities,” *Advances in Building Energy Research*, pp. 1–19, 2015.
- [177] F. Oldewurtel, D. Gyalistras, M. Gwerder, C. Jones, A. Parisio, V. Stauch, B. Lehmann, and M. Morari, “Increasing energy efficiency in building climate control using weather forecasts and model predictive control,” in *Clima-RHEVA World Congress*, no. EPFL-CONF-169735, 2010.
- [178] A. R. Florita and G. P. Henze, “Comparison of short-term weather forecasting models for model predictive control,” *HVAC&R Research*, vol. 15, no. 5, pp. 835–853, 2009.
- [179] E. Canada, “<http://dd.weather.gc.ca/forecasting/>,” 2016.
- [180] R. A. Tanner and G. P. Henze, “Stochastic control optimization for a mixed mode building considering occupant window opening behaviour,” *Journal of Building Performance Simulation*, vol. 7, no. 6, pp. 427–444, 2014.
- [181] F. Oldewurtel, D. Sturzenegger, and M. Morari, “Importance of occupancy information for building climate control,” *Applied Energy*, vol. 101, pp. 521–532, 2013.
- [182] L. Wang, *Model predictive control system design and implementation using MATLAB®*. Springer Science & Business Media, 2009.
- [183] K. Antonopoulos and C. Tzivanidis, “Finite-difference prediction of transient indoor temperature and related correlation based on the building time constant,” *International journal of energy research*, vol. 20, no. 6, pp. 507–520, 1996.
- [184] Z. Shi, A. Abdelalim, W. O’Brien, R. Attar, P. Akiki, K. Graham, B. Van Waarden, S. Fai, A. Tessier, and A. Khan, “Digital campus innovation project: integration of building information modelling with building performance simulation and building diagnostics,” in *Proceedings of the Symposium on Simulation for Architecture & Urban Design*, pp. 51–58, Society for Computer Simulation International, 2015.
- [185] H. B. Gunay, “Supplemental: Matlab scripts used in inverse modelling available at <https://goo.gl/r60jau>,” 2016.

- [186] H. B. Gunay, “Supplemental: Matlab scripts used in developing occupant models available at <https://goo.gl/ozcfqc>,” 2016.
- [187] S. Karjalainen *et al.*, *The characteristics of usable room temperature control*. VTT Technical Research Centre of Finland, 2008.
- [188] S. Karjalainen, “Gender differences in thermal comfort and use of thermostats in everyday thermal environments,” *Building and environment*, vol. 42, no. 4, pp. 1594–1603, 2007.
- [189] U. Satish, M. J. Mendell, K. Shekhar, T. Hotchi, D. Sullivan, S. Streufert, and W. J. Fisk, “Is co2 an indoor pollutant? direct effects of low-to-moderate co2 concentrations on human decision-making performance,” *Environmental health perspectives*, vol. 120, no. 12, p. 1671, 2012.
- [190] ASHRAE, “Standard 74,” 1988.
- [191] K. Gowri, D. W. Winiarski, and R. E. Jarnagin, *Infiltration modeling guidelines for commercial building energy analysis*. Pacific Northwest National Laboratory USA, 2009.
- [192] M. Barclay, N. Holcroft, and A. Shea, “Methods to determine whole building hygrothermal performance of hemp–lime buildings,” *Building and environment*, vol. 80, pp. 204–212, 2014.
- [193] F. Rubinstein, G. Ward, and R. Verderber, “Improving the performance of photo-electrically controlled lighting systems,” *Journal of the Illuminating Engineering Society*, vol. 18, no. 1, pp. 70–94, 1989.
- [194] R. Mistrick, C.-H. Chen, A. Bierman, and D. Felts, “A comparison of photosensor-controlled electronic dimming systems in a small office,” *Journal of the Illuminating Engineering Society*, vol. 29, no. 1, pp. 66–80, 2000.
- [195] J. D. Jennings, F. M. Rubinstein, D. DiBartolomeo, and S. L. Blanc, “Comparison of control options in private offices in an advanced lighting controls testbed,” *Journal of the Illuminating Engineering Society*, vol. 29, no. 2, pp. 39–60, 2000.
- [196] S. Mukherjee, D. Birru, D. Cavalcanti, E. Shen, M. Patel, Y.-J. Wen, and S. Das, “Closed loop integrated lighting and daylighting control for low energy buildings,” *Proceedings of the 2010 ACEEE*, pp. 252–269, 2010.

- [197] M. Skelly and M. Wilkinson, "The evolution of interactive facades: improving automated blind control," *Whole Life Performance of Facades*, pp. 129–142, 2001.
- [198] J.-H. Kim, Y.-J. Park, M.-S. Yeo, and K.-W. Kim, "An experimental study on the environmental performance of the automated blind in summer," *Building and Environment*, vol. 44, no. 7, pp. 1517–1527, 2009.
- [199] E. S. Lee and S. E. Selkowitz, "The new york times headquarters daylighting mockup: Monitored performance of the daylighting control system," *Energy and buildings*, vol. 38, no. 7, pp. 914–929, 2006.
- [200] E. Lee, R. Clear, G. Ward, and L. Fernandes, "Commissioning and verification procedures for the automated roller shade system at the new york times headquarters, new york, new york," *Berkeley, CA: Lawrence Berkeley National Laboratory*, 2007.
- [201] S.-T. Claros and A. Soler, "Indoor daylight climate—influence of light shelf and model reflectance on light shelf performance in madrid for hours with unit sunshine fraction," *Building and Environment*, vol. 37, no. 6, pp. 587–598, 2002.
- [202] T. Inoue, T. Kawase, T. Ibamoto, S. Takakusa, and Y. Matsuo, "The development of an optimal control system for window shading devices based on investigations in office buildings," *ASHRAE transactions*, vol. 94, pp. 1034–1049, 1988.
- [203] A. Tzempelikos, "The impact of manual light switching on lighting energy consumption for a typical office building," 2010.
- [204] J. Heerwagen, "Adaptations and coping: Occupant response to discomfort in energy efficient buildings," 1992.
- [205] P. Boyce, J. Veitch, G. Newsham, C. Jones, J. Heerwagen, M. Myer, and C. Hunter, "Occupant use of switching and dimming controls in offices," *Lighting Research and Technology*, vol. 38, no. 4, pp. 358–376, 2006.
- [206] Y. Sutter, D. Dumortier, and M. Fontoyntont, "The use of shading systems in vdu task offices: A pilot study," *Energy and Buildings*, vol. 38, no. 7, pp. 780–789, 2006.

- [207] M. B. Aries, J. A. Veitch, and G. R. Newsham, “Windows, view, and office characteristics predict physical and psychological discomfort,” *Journal of Environmental Psychology*, vol. 30, no. 4, pp. 533–541, 2010.
- [208] F. Haldi, “A probabilistic model to predict building occupants diversity towards their interactions with the building envelope,” in *Proceedings of the international IBPSA conference, Chambery, France*, p. 147, 2013.
- [209] ASHRAE, *ANSI/ASHRAE/IES Standard 90.1-2013 – Energy Standard for Buildings Except Low-Rise Residential Buildings*. ASHRAE, 2013.
- [210] A. Cooperman and J. Dieckmann, “Control systems & leed,” *ASHRAE Journal*, vol. 54, no. 6, p. 96, 2012.