

**An approach to data-driven sensing and predictive  
supervisory control for commercial buildings with in-  
situ evaluation**

by

**Jayson Bursill**

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

**Doctor of Philosophy**  
in  
**Mechanical Engineering**

Department of Mechanical and Aerospace Engineering

Carleton University

Ottawa, Ontario

April 2020

©Copyright

Jayson Bursill, 2020

## **Abstract**

Modern building automation systems offer virtually limitless programmability, sensing, and connectivity, but often do not achieve the energy use or occupant comfort goals of building stakeholders. One major source of the shortfall of modern building automation systems is room-level and zone-level control. Given the lack of perceived importance of zone-level controllers compared to system-level and building-level controllers, comprehensive and efficient zone-level programming is overlooked at the cost of quality of implementation. Accordingly, both sensing capabilities and the use of advanced control algorithms (such as model-predictive control) have been largely unexploited at the room and zone level. This thesis addresses these obstacles to improve comfort and energy performance by developing and testing a suite of interrelated approaches and technologies. Field testing was conducted in 27 highly instrumented (but typical) offices in the Health Sciences Building at Carleton University.

One approach to reducing sensor infrastructure (and associated installation and maintenance costs) for advanced control applications is to apply proxy sensing using data-driven modelling. By comparing indirect sensors for energy use estimation (i.e., proxy sensors) to direct energy use estimation sensors, this thesis develops an approach to rank all physically reasonable sensors at the zone level for energy use estimation usefulness. From this approach, it was found that outdoor air temperature, hydronic flow, and hydronic temperature sensors were the most useful for proxy sensing.

The performance of novel supervisory control approaches can be quantified with the ability to accurately estimate room-level energy use. This thesis developed a simplified model-based predictive control approach for the heating season to improve zone-level

control. This approach involved testing detailed model-based predictive control (MPC) in a single office and using classification learning (i.e., sorting predicted outputs into bins based on shared input qualities) to simplify the algorithm. The detailed MPC and simplified MPC (sMPC) saved an estimated 23% and 21% of energy use respectively when compared to a reactive control baseline over two weeks of testing for MPC and one week for sMPC. The sMPC achieved 94% of the energy savings of MPC, but the implementation in a typical building automation system (BAS) controller was much simpler. The sMPC still contained a model for indoor air temperature prediction at the next time-step, and because only one office was tested variation in ambient conditions produced uncertain results. Further, this test office did not have direct energy measurement sensors and relied on equipment run-time and air volume measurements for energy use estimates.

The simplified MPC approach was modified to remove the underlying indoor air temperature model and scaled to operate in both the heating and cooling seasons in 27 rooms of an institutional building at Carleton University. This 27-room study is the largest in-situ zone-level rule extraction (RE) control study to date. The MPC and RE control reduced energy use by 42% and 27% in the cooling season and 18% and 33% in the heating season when compared over four weeks of testing for each algorithm and season. Thus, the RE approach was experimentally verified to have comparable energy savings to detailed MPC while using only a few dozen lines of BAS controller-embedded code.

Future work should expand and package the proxy sensing and rule extraction approaches developed in this thesis into commercial BAS products to take advantage of the potential for zone-level energy savings and improve the overall operation of commercial buildings.

## Preface

This integrated article thesis consists of three journal papers, either published or in review.

The articles integrated into this thesis include the following:

1. Jayson Bursill, William O'Brien, Ian Beausoleil-Morrison. Proxy zone-level energy use estimation in a commercial building with a variable air volume system. *Manuscript submitted for publication.*
2. Jayson Bursill, William O'Brien, Ian Beausoleil-Morrison. 2019. Experimental application of classification learning to generate simplified model predictive controls for a shared office heating system. *Science and Technology for the Built Environment 25 (5), 615-628.*
3. Jayson Bursill, William O'Brien, Ian Beausoleil-Morrison. Multi-zone field study of rule extraction control to simplify implementation of predictive control to reduce building energy use. *Energy and Buildings, Accepted for publication.*

Some modifications have been made to the articles. These modifications include the removal of the abstract sections and the merging of the conclusion sections in Chapter 7 of this thesis.

The use of copyrighted material from the previously mentioned articles in this thesis is acknowledged to adhere to the publisher's permission guidelines with respect to the author's rights provided in Appendix A.

In the co-authored articles, Jayson Bursill was the principal contributor to the research methodology, concept and design, experiment set up and conduct, data acquisition, data analysis and interpretation, and preparing and writing the material

presented in the articles under the supervision of William O'Brien and Ian Beausoleil-Morrison.

## **Acknowledgments**

I would like to begin by thanking my supervisors Dr. Liam O'Brien and Dr. Ian Beausoleil-Morrison. Without their guidance and support, I would not have been able to conduct this research over the last four years.

I would like to thank my colleagues from Carleton University. These people include but are not limited to Burak Gunay, Shawn Shi, Chris Baldwin, Connor Brackley, Adam Wills, Dan Lowcay, Mohammad Derakhti, Brock Conley, Sebastien Brideau, Brent Huchuk, Phillip Drouillard, Arty David, and Garrett Lecoq.

I would like to acknowledge the support of Delta Controls, the NSERC and Carleton University Ontario Graduate Scholarships, ASHRAE Grant-in-Aid, and the IBPSA-Canada travel award. These were essential to financially support this research project. I would also like to acknowledge the support of Carleton Facilities Management and Planning for allowing the research team to be an integral part of the Health Sciences Building design.

I would like to thank Katherine Kenny, without whom I may not have had the will to complete this thesis.

Finally, I would like to thank my family. Without my Mother Shelley and Father Barry, I could not be here today. Without my Uncle Darryl and Aunt Francine, I would not have been enlightened to the potential of the building industry. Without the rest of my family, I would not have had the additional support needed to proceed.

# Table of Contents

<b>Abstract</b> .....	<b>ii</b>
<b>Preface</b> .....	<b>iv</b>
<b>Acknowledgments</b> .....	<b>vi</b>
<b>Table of Contents</b> .....	<b>vii</b>
<b>List of Tables</b> .....	<b>xi</b>
<b>List of Figures</b> .....	<b>xiii</b>
<b>List of Appendices</b> .....	<b>xx</b>
<b>Nomenclature</b> .....	<b>xxi</b>
<b>Chapter 1: Introduction</b> .....	<b>1</b>
1.1    Research Objectives and Questions.....	5
1.2    Theoretical Framework .....	6
1.3    Thesis Outline.....	12
<b>Chapter 2: A Brief History of Commercial Building Controls and Modern Trends</b> .....	<b>14</b>
2.1    Antiquity to Pneumatics .....	14
2.2    Pneumatics to Digital Control .....	16
2.3    The Data Expansion: Modern Trends and Future Directions .....	18
<b>Chapter 3: Data-driven Applications for Commercial Building Supervisory Control</b> .....	<b>27</b>
3.1    Data-driven Approach Background.....	27
3.2    Artificial Neural Network Models.....	30
3.3    State-space Models .....	30

3.4	Classification Trees .....	31
3.5	Support Vector Machines .....	32
3.6	Reinforcement Learning .....	33
3.7	Data-driven Approaches Summary .....	34
<b>Chapter 4: A Sensor Feasibility Study.....</b>		<b>36</b>
4.1	Introduction .....	37
4.2	Literature Review .....	40
4.3	Methodology.....	43
4.3.1	Sensor Introduction .....	44
4.3.2	Energy Quantification .....	50
4.3.3	Uncertainty Analysis.....	52
4.3.4	Sensor Proxy Approach.....	54
4.3.5	Ranking Sensors.....	62
4.4	Results and Discussion .....	63
4.4.1	Energy End Uses .....	64
4.4.2	Sensor Proxy Results.....	70
4.4.3	Hierarchy of Commercial Building Sensors .....	75
4.4.4	Limitations .....	77
4.4.5	Anecdotal Observations .....	79
4.5	Summary.....	83
<b>Chapter 5: Lab Implementation of Classification Learning-based sMPC.....</b>		<b>86</b>

5.1	Introduction .....	87
5.2	Literature Review .....	89
5.3	Methodology.....	93
5.3.1	Experimental Office Characteristics and BAS infrastructure .....	94
5.3.2	Experimental Approach .....	96
5.3.2.1	Modeling and Control.....	104
5.3.2.2	Classification Learning for Rule Extraction.....	107
5.3.2.3	Energy and Comfort Quantification .....	114
5.4	Results .....	116
5.4.1	Weather Forecast Accuracy .....	116
5.4.2	Reactive Control Baseline.....	117
5.4.3	Online MPC .....	119
5.4.4	Online sMPC.....	121
5.4.5	Seasonal Energy Savings and Thermal Comfort.....	123
5.5	Summary.....	128
<b>Chapter 6: Field Implementation of Rule Extraction Supervisory Control .....</b>		<b>131</b>
6.1	Introduction .....	132
6.2	Methodology.....	135
6.2.1	Description of Facility.....	136
6.2.2	Uncertainty Analysis.....	141
6.2.3	Detailed Model Predictive Control Development.....	142

6.2.3.1	Test Plan .....	143
6.2.3.2	Model Development .....	144
6.2.3.3	Control Strategy .....	149
6.2.4	Rule Extraction Predictive Control Development.....	153
6.2.5	Energy Use Analysis .....	160
6.3	Results and Discussion .....	161
6.3.1	Control Mode Operational Comparison.....	162
6.3.2	Energy Use and Indoor Air Temperature .....	172
6.3.3	Limitations of Study.....	177
6.4	Summary.....	179
<b>Chapter 7: Conclusions and Future Work .....</b>		<b>182</b>
7.1	Contributions .....	186
7.2	Future Work.....	187
<b>Appendices.....</b>		<b>190</b>
Appendix A : Manuscript Publishing Agreements.....		190
A.1	Elsevier (Journal of Building Engineering and Energy and Buildings) – Personal Use .....	190
A.2	Taylor & Francis (Science and Technology for the Built Environment) - Free Materials Definition .....	191
A.3	Taylor & Francis (Science and Technology for the Built Environment) - Usage rights for Free Materials .....	191
<b>References .....</b>		<b>193</b>

## List of Tables

Table 1-1: Sensor schedule for the rooms studied in Chapter 4 and Chapter 6 with sensor type and manufacturer-provided instrumentation error .....	9
Table 2-1: Aggregated cost information and energy savings estimates for HVAC sensors .....	20
Table 2-2: Aggregated cost information and energy savings estimates for plug load and lighting sensors .....	21
Table 2-3: Aggregated cost information and energy savings estimates for window shading technology.....	21
Table 4-1: Sensor schedule with simulation abbreviation, sensor type, and instrumentation error. Common sensors in commercial buildings are indicated in italics. ....	46
Table 4-2: Energy use breakdown for measured rooms for the cooling and heating seasons. “*” indicates missing data replaced by the average value, “***” indicates rooms with high window to wall ratios.....	66
Table 4-3: Energy use prediction score for each input with respect to each energy end use studied for the cooling season (top) and heating season (bottom).....	73
Table 4-4: Sensor hierarchy using aggregated input scores from all end uses for both the heating and cooling season for the measured building. The sensors have been divided into tiers based on distinct gaps in total score.....	76
Table 5-1: Input sensors for Reduced Order Models.....	96

Table 5-2: Weekly algorithm schedule, average OAT, and average occupant fraction ...	97
Table 5-3: Input average value over the testing period and manufacturer-provided instrument errors .....	115
Table 5-4: Weekly energy consumption, average solar proxy, average OAT, and average occupancy values for each algorithm.....	127
Table 6-1: Sensor schedule with sensor type and manufacturer-provided instrumentation error.....	139
Table 6-2: Sensor error summary table.....	142
Table 6-3: Floor-level weekly algorithm schedule, average OAT, and average occupant fraction .....	144
Table 6-4: Summary of weekly energy use of the floors and their corresponding algorithms. Average floor-level occupancy ratios and IAT setpoint deviation are also provided.....	175

## List of Figures

Figure 1-1: The simulation framework for inverse modelling used in this thesis. Further details of the individual steps within this framework can be found in (Bursill, O’Brien et al. 2019) .....	7
Figure 1-2: Building façade viewed from North (top) and diagram of conditioned space with zone-level HVAC system and select sensors (bottom).....	8
Figure 2-1: Typical digital control BAS network topology.....	18
Figure 2-2: BACnet commercial building controls market share of new equipment from 2012 to 2017 .....	24
Figure 2-3: BACnet vendor counts from 1999 to 2019 .....	24
Figure 2-4: BACnet vendor growth at five-year intervals from 1999 to 2019 .....	25
Figure 4-1: Return air temperature sensor access location marked with a red box (view from below). Return air is drawn back to the AHU via the duct at the grated opening in the bottom-left of the figure.....	47
Figure 4-2: Sensors, their inferred conditions, and the sensors that can be utilized to infer such conditions (“*” indicates sensors intended for direct measurement) .....	49
Figure 4-3: Input count per zone sensor from the inferred conditions of Figure 4-2, where each appearance of a zone sensor in an inferred condition increments the input count of that zone sensor by one (e.g., indoor air temperature appeared in 10 inferred conditions). .....	50
Figure 4-4: ANN structure of proxy sensor selection script for a SISO case with one input	

on the 1x1 hidden neuron iteration (the minimum complexity iteration for the SISO case)  
..... 58

Figure 4-5: ANN structure of proxy sensor selection script for a MISO case with two inputs  
on the 5x5 hidden neuron iteration (the maximum complexity iteration for the two input  
single output case)..... 58

Figure 4-6: Aggregate energy prediction over the two-month validation period for the  
validation data using individual sensor proxy models for office-level radiant heater energy  
using only OAT (left) and IAT (right). Each dot represents the aggregate measured and  
modelled energy use for one room over the validation period. .... 59

Figure 4-7: Forward model MBE progression of ANNs for VAV reheat coil power (top)  
and AHU cooling power (bottom) in the cooling season. The input labels on the x-axis are  
provided in Table 4-1..... 61

Figure 4-8: Forward model MBE progression of ANNs for plug loads (top) and AHU  
heating power (bottom) in the heating season. The input labels on the x-axis are provided  
in Table 4-1. .... 62

Figure 4-9: Energy use breakdown and occupancy fraction for monitored offices for the  
cooling season (top) and heating season (bottom). Lighting energy use is at the top and not  
visible for most offices..... 67

Figure 4-10: Energy usage per end use by percent for the average of the rooms (left) and  
the median total energy room (right) for the cooling season (top) and heating season  
(bottom)..... 69

Figure 4-11: Suspected space heater usage in an office due to high plug load and the

corresponding effect on IAT.....	80
Figure 4-12: Suspected space heater temperatures and usage data.....	81
Figure 4-13: Heatmap comparison of indoor air temperature and return air temperature for all rooms.....	83
Figure 5-1: Location of building automation system sensors used for reduced order modelling within the measured office.....	96
Figure 5-2: Observed ceiling illuminance linear fit to CanMETEO solar irradiation predictions at the office wall orientation. ....	101
Figure 5-3: Occupancy profiles for the testing period where a 0 indicated the room was unoccupied and a 1 indicated the room was occupied. The minimum occupied time for the sensor was 15 minutes. ....	103
Figure 5-4: MPC override classification pseudocode for heating season. A 0 indicates the heating equipment should be overridden and a 1 indicates it should continue using reactive control. ....	105
Figure 5-5: MPC (top) and reactive control (bottom) block diagrams .....	106
Figure 5-6: sMPC classification pseudocode implemented in conditioned space for heating season.....	109
Figure 5-7: Sample classification tree for simulated predictive control output. A 0 indicates the heating equipment should be overridden and a 1 indicates it should continue using reactive control.....	112
Figure 5-8: Confusion matrix for simulated predictive control output training. The target	

class and output class correspond with the MPC and sMPC outputs, respectively. Classification error is indicated on the diagonal of the matrix. ....	113
Figure 5-9: Confusion matrix for simulated output validation. The target class and output class correspond with the MPC and sMPC outputs, respectively. Classification error is indicated on the diagonal of the matrix.....	113
Figure 5-10: Input prediction accuracy on April 3 <sup>rd</sup> , 2018 for OAT, illuminance, and occupancy. ....	117
Figure 5-11: Typical temperature predictions and control outputs for reactive control baseline on March 22 <sup>nd</sup> , 2018, where the OAT was on average 1.8°C .....	119
Figure 5-12: Typical temperature predictions and control outputs for MPC case on March 26 <sup>th</sup> , 2018, where the OAT was on average -2.0°C .....	120
Figure 5-13: Typical temperature predictions and control outputs for sMPC case on April 17 <sup>th</sup> , 2018, where the OAT was on average -1.2°C.....	123
Figure 5-14: Weekly heating and cooling energy use for the testing period. Average solar proxy, OAT, and occupancy for each week is also indicated. ....	126
Figure 6-1: High-level MPC and RE design and implementation workflow diagram ...	136
Figure 6-2: Diagram of a two-room zone with zone-level HVAC system and BAS-connected sensors.....	140
Figure 6-3: Mean occupancy and predicted occupancy schedule for room 1. The predicted occupancy schedule calculates the recursive mean arrival and departure times and applies 1.5 standard deviations of the respective times to encompass 90% of all arrivals and	

departures under a Gaussian assumption. This follows the same format as used in (Gunay, O'Brien et al. 2016). ..... 149

Figure 6-4: Flow chart for reactive control suppression from the MPC algorithm and setpoint data where  $t$  and  $n$  are the current time-step and predicted time-step respectively. When  $MPC = 0$  it indicates the HVAC equipment should be overridden, and when  $MPC = 1$  it indicates the HVAC equipment should continue using reactive control.  $\Delta U$  is the array of predicted energy required by the HVAC system at each timestep to condition the IAT to the IAT setpoint. .... 151

Figure 6-5: Detailed MPC control agent workflow. BACnet reactive controller input, BACnet reactive controller output, and CanMETEO weather prediction data are read into comma delimited files on a PC using Python scripts. The ACE then executes the MPC algorithm on the data, writes the predicted overrides into comma delimited files, and sends BACnet output override commands back to the BAS from the comma delimited files via a script. This is all executed on a at regular intervals using Windows Task Scheduler. ... 153

Figure 6-6: Node vs. iteration for the cooling season (top) and heating season (bottom) classification tree fitting processes for 1000 sample-randomized scenarios. Results were sorted such that node counts increment from left to right. One node represents one split where two unique decision tree outcomes can occur..... 156

Figure 6-7: Cooling season decision tree. A 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control. .... 158

Figure 6-8: Heating season decision tree. A 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control. .... 159

Figure 6-9: Reactive control operation for room 3 in the cooling season (top) and heating season (bottom)..... 163

Figure 6-10: MPC operation for room 3 in the cooling season (top) and heating season (bottom). For the MPC bit, a 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control..... 166

Figure 6-11: RE operation for room 3 in the cooling season (top) and heating season (bottom). For the RE bit, a 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control..... 167

Figure 6-12: Comparison of MPC and RE overrides (top), zone and MPC overrides (middle), and zone and RE overrides (bottom) in the cooling season. The period of condition compares the percent of operational time that the MPC override, RE override, and zone overrides meet the conditions in the legend of each bar graph. In ideal conditions the period of condition of the solid black “=” condition would be 100% for each room and each graph. Classification error and differences in room HVAC needs cause this to be less than 100%. ..... 170

Figure 6-13: Comparison of MPC and RE overrides (top), zone and MPC overrides (middle), and zone and RE overrides (bottom) in the heating season. The period of condition compares the percent of operational time that the MPC override, RE override, and zone overrides meet the conditions in the legend of each bar graph. In ideal conditions the period of condition of the solid black “=” condition would be 100% for each room and each graph. Classification error and differences in room HVAC needs cause this to be less than 100%. ..... 171

Figure 6-14: Zoning of rooms indicated by room numbers within the same box. Zones groupings varied in size in the measured building from one to four rooms per zone..... 172

Figure 6-15: Weekly energy use and average ambient conditions for cooling season (top) and heating season (bottom) ..... 176

## **List of Appendices**

Appendix A: Manuscript Publishing Agreements.....	190
---	-----

## Nomenclature

Abbreviation	Full name
ACE	Agnostic control engine
AHU	Air handling unit
ANN	Artificial neural network
ASHRAE	American Society of Heating, Refrigerating, and Air-conditioning Engineers
BACnet	Building automation control network protocol
BAS/BMS	Building automation system/Building management system
BSRIA	Building Services Research and Information Association
CART	Classification and regressions tree
EUI	Energy use intensity
GCL+	General Control Language
HMI	Human-machine interface
HVAC	Heating, ventilating, and air conditioning
IAT	Indoor air temperature
ID	Identification number
IoT	Internet of things
LED	Light-emitting diode
MAE	Mean absolute error
MBE	Mean bias error
MISO	Multiple input single output
MPC	Model-based predictive control
MS/TP	Master-slave token-passing protocol
OAT	Outdoor air temperature
PI	Proportional-integral
PID	Proportional-integral-derivative
RAT	Return air temperature
RE	Rule-extraction
RL	Reinforcement learning
RMSE	Root-mean-square error
SISO	Single input single output
sMPC	Simplified model-based predictive control
SVM	Support vector machine
VAV	Variable air volume
VDC	Volts direct current

<b>Symbol</b>	<b>Full name</b>
A	Linear time-invariant constant
airflow	VAV airflow
B	Linear time-invariant constant
blind	Blind position
$c_{pw}$	Specific heat of water
FWD	Forward selection model score
H	Hessian
hour1	Hour of day
i	Index
IAT	Indoor air temperature
IND	Individual input score
J	Cost constant
Ldim	Lighting dim level
LP	Lighting power
LS	Lighting state
lux1	Illuminance near window
lux2	Illuminance in center of the room
lux3	Illuminance near wall
M	Measured data
N	Number of samples
OAT	Outdoor air temperature
occ	Occupancy state
plug	Plug load
Q/U	Energy
$Q_{sa}$	Supply air heating energy
R	Matrix of penalty weights
$R^2$	Coefficient of determination
rad	Radiant heater valve position
rad_flow	Radiant heater flow rate
radt1	Radiant heater inlet temperature
radt2	Radiant heater outlet temperature
RAT	Return air temperature
rhc_flow	Reheat coil flow rate
rhct1	Reheat coil inlet temperature
rhct2	Reheat coil outlet temperature
s	Aggregate uncertainty
S	Simulated data
sat	VAV supply air temperature
SP	Temperature setpoint
T	Indoor air temperature

$t$	Time-step index
$T_{1w}$	Supply water temperature
$T_{2w}$	Return water temperature
$\Delta U$	Energy input from the heating and cooling system
$u$	Input array
$V_{AHU}$	Air handling unit air flow rate
$V_h$	Fluid flow rate
$V_r$	Room-level air flow rate
$win$	Window state
$x$	Input variable
$\rho_w$	Density of water
$\phi$	MPC parameter matrix

## Chapter 1: Introduction

Up to 70% of total building energy use can be attributed to heating and cooling (Afram and Janabi-Sharifi 2014, Dong, Prakash et al. 2019). One approach to reducing heating and cooling energy use, model-based predictive control (MPC), has demonstrated the potential provide significant energy savings through software modifications to building automation system (BAS) controls (Gunay, O'Brien et al. 2016, Gunay, O'Brien et al. 2016, Hilliard, Swan et al. 2016). Despite the potential and effort provided by the academic community to advance the use of predictive control in buildings (Anderson, Buehner et al. 2008, Oldewurtel, Parisio et al. 2012, Touretzky and Baldea 2014, Žáčková, Váňa et al. 2014, Wang and Ong 2017, Serale, Fiorentini et al. 2018), commercial building controls remain largely based on proportional-integral-derivative (PID) control without any predictive capabilities (Salsbury 2005, Wang and Ma 2008, Killian and Kozek 2016, Stłuka, Parthasarathy et al. 2018). Commercially implemented predictive control is uncommon at the building level, and virtually non-existent at the zone and room levels (Royapoor, Antony et al. 2018).

Perhaps the greatest fundamental challenge of MPC use is the underlying modelling approach. There is a trade-off between model accuracy and computation time. Detailed models can require thermo-physical data that is difficult to estimate through measurement (Killian and Kozek 2016). Data-driven methods for predictive control are often selected over physics-based methods for control applications due to the volume of data currently available from commercial buildings for inverse modelling and the lower computational resource requirements after the model is trained (Aswani, Master et al. 2012, Ma, Kelman

et al. 2012, Cigler, Siroky et al. 2013, Ryzhov, Ouerdane et al. 2019, Xuan, Xuehui et al. 2019, Maddalena, Lian et al. 2020). However, physics-based methods have had success when considered in the design stage, when used as a basis for simplified modelling (i.e., “meta-modelling”), or if intermediate parameters not available in data-driven approaches are useful (Zhao, Lam et al. 2013, Zhao, Lam et al. 2015, Salakij, Yu et al. 2016, Zhang and Lam 2017). There is still contention within the MPC community over which method will obtain commercial dominance.

The practical challenges of supervisory MPC are defined by Maddalena et al. as “the building control problem” using four points summarized as follows:

1. Large sampling periods: Sample intervals are usually 15 minutes or longer.
2. Low-quality measurements: Many sensors have manufacturer-provided measurement uncertainty of 2% full scale or greater (e.g., pressure transducers). In practice, this can become much higher when including instrumentation error and aggregating multiple sensors (and corresponding uncertainties) to estimate a value.
3. Operations-critical: It is difficult to obtain override control to test algorithms in-situ.
4. Difficult modelling: Dynamic data-driven modelling is difficult when compounded with low-quality measurements.

Four attributes of ideal supervisory algorithms are defined by Maddalena et al. as:

1. Lifetime adaptability: Models must refine and improve over the operation period.
2. Safety: Guarantees of safety must be made via manual supervision or automated fail-safes.

3. Incorporation of exogenous signals: Uncontrollable disturbances (e.g., weather, occupancy patterns) must be accounted for in some form.
4. Scalability: The operation of the predictive algorithm should not take an unreasonable amount of time, and be applicable at multiple levels of heating, ventilating, and air conditioning (HVAC) systems.

The core concept of the building control problem is not new and was clearly identified as early as the mid-1980s in the American Control Conference proceedings (Nesler 1985). As presented by Nesler, to address the building control problem, a “self-tuning” controller must:

1. Implement the PI controller algorithm (i.e., algorithm parameters).
2. Not require any prior knowledge of the process.
3. Be sufficiently robust to operate without supervision.

The term “self-tuning” can be interpreted in modern terminology as a supervisory algorithm. At the time of the research conducted in this thesis, supervisory control is rare in commercial buildings. While not explicitly supervisory control, ASHRAE Guideline 36 defines high-performance sequences of operation (but does not provide suggestions for when unexpected failures or inefficiencies occur within the programming) (ASHRAE 2018).

Anecdotally, the most common criticism of supervisory algorithms from BAS vendors is the level of expertise required for implementation. Thus, a simplified, justified, well-documented, and packaged approach reduces the required level of expertise to implement supervisory algorithms like MPC.

Inspired by the potential for energy savings cited in the literature, the research in this thesis studies the usage of indirect sensing (e.g., outdoor air temperature sensors to estimate room-level energy use) to reduce the cost of energy use estimation and simplified implementation of MPC. Energy use estimation is an essential element of MPC implementation to validate the actual energy savings. The results of this work are intended to address the building control problem by developing and testing a novel supervisory control approach. The results are intended to be useful to building owners and operators while remaining simplified into a format that can be implemented in modern building controllers without (or with limited) expert assistance.

This PhD thesis project addresses the building control problem using the following studies:

1. Identifying the most valuable sensors for energy use estimation at the room and zone control level and using them to quantify energy use at the room level. Building-integrated sensors are expensive to install, thus reducing the number of sensors installed at the zone-level can reduce costs significantly and provide a baseline for comparing supervisory control algorithms.
2. Initial implementation of classification-learning-based supervisory heating control in a single office using a data-driven zone air temperature model and a remote computer to generate a simplified MPC approach.
3. Expanding and improving the classification-learning-based approach from study 2 for both heating and cooling to 27 rooms in an occupied institutional building. The accompanying aggregate energy use and setpoint control were analyzed to quantify

energy savings (using approaches introduced in study 1) and occupant comfort. This third study combines the baselines and some of the sensors used in the first study with the approaches to supervisory control used in the second study.

The following sections of the introduction describe the research objectives, theoretical framework, and thesis outline.

## **1.1 Research Objectives and Questions**

The main objectives of this research are to improve the use of sensing, data storage, and the processing capabilities of modern building automation systems to reduce energy use, maintain or improve occupant comfort (as defined by the indoor air temperature and its setpoint), and simplify the implementation of advanced algorithms at the room level. The previously mentioned objectives of this research were conducted, and the findings are presented in a format that can be used more easily by commercial building control programmers than existing approaches. The experimental analysis in this thesis begins by estimating energy use because energy savings are difficult to quantify without a baseline. Baseline estimation is followed by gradual implementation of proxy (i.e., indirect) sensing and MPC algorithms to address the research questions as follows:

1. Can indirect sensors estimate room-level energy use with a level of accuracy that is accepted by common modelling standards (e.g., ASHRAE Guideline 14)?
2. Can simple rules be extracted from a simple supervisory MPC algorithm that can achieve similar energy savings to the original algorithms in-situ?
3. Can a rule-extraction approach to predictive control achieve in-situ energy savings on a larger scale than a single office?

## 1.2 Theoretical Framework

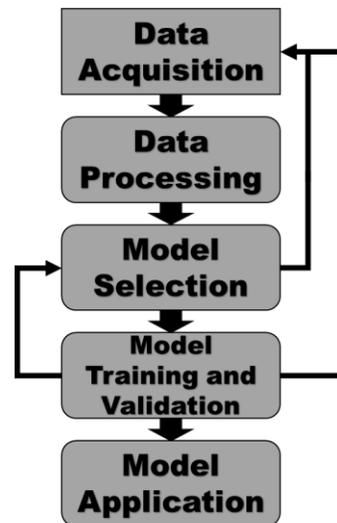
A theoretical framework based on zone-level implementation and fair comparison of algorithms was used to achieve the objectives of this research. Comprehensive sensing is a necessary foundation to achieve a fair comparison of algorithms by estimating the energy use of the rooms within the zones.

To facilitate comprehensive testing, two similar floors of an institutional building were used to evaluate models and algorithms with one test floor and one control (i.e., untested) floor that had near-identical ambient conditions. The test and control floors were rotated on a weekly basis to minimize the effect of occupant behavior on energy use results. With comprehensive sensing and data acquisition, longitudinal data were used to develop data-driven models for proxy sensing and MPC.

The research conducted in this thesis is composed primarily of in-situ evaluation of algorithms but has a significant basis in simulation in the form of data-driven modelling. Thus, the theoretical framework is split into simulation and experimental frameworks in this research. The specific experimental frameworks of the manuscript-based in-situ studies are described in detail as they pertain to their objective studies in Chapters 4 to 6.

The data-driven modeling in this thesis required a rigorous simulation framework. The simulation framework in this thesis relies on the iterative process shown in Figure 1-1. Data acquisition, data processing, model selection, model training and validation, and model application are defined as the minimally necessary steps for data-driven modelling (Bursill, O'Brien et al. 2019). Model cross-validation was used for inverse model validation. In this research half of the data is used for validation (and half for training), but

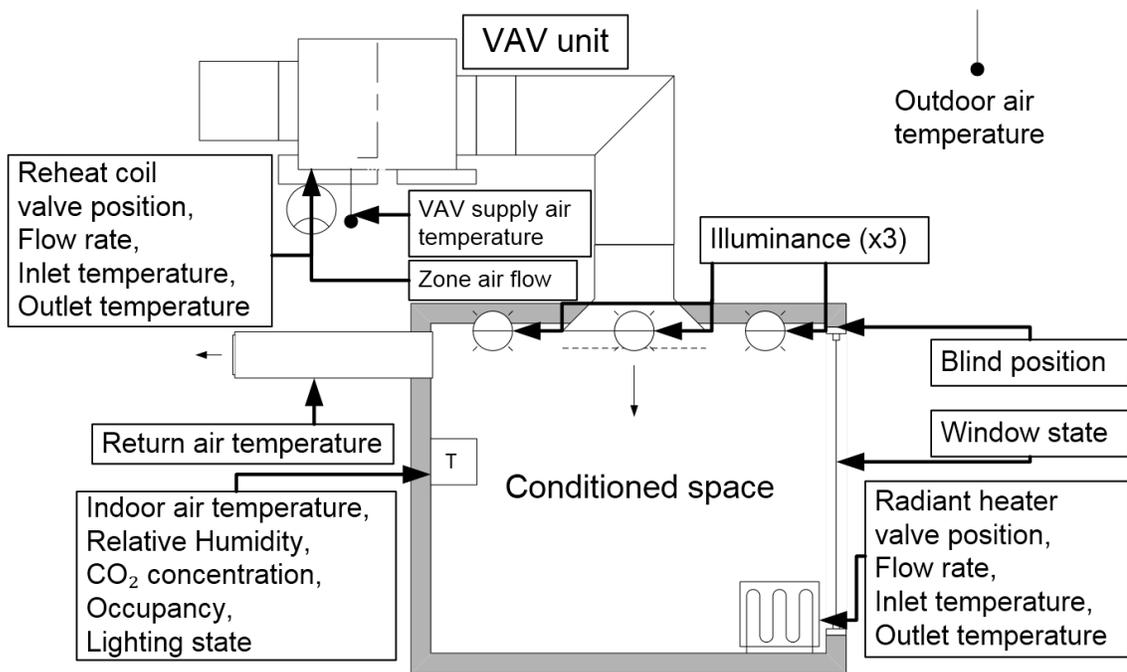
different distributions of training and validation data may be used for specialized applications (Shahin, Maier et al. 2004).



**Figure 1-1: The simulation framework for inverse modelling used in this thesis. Further details of the individual steps within this framework can be found in (Bursill, O’Brien et al. 2019)**

Data was collected from the Health Sciences Building at Carleton University in Ottawa, Canada from 27 rooms within 11 thermal zones. Most of the instrumented rooms (25) are single-occupant private offices, while the other two rooms are kitchenettes. Each office is the same size, shape, and has the same diffuser location and orientation. The primary façade for the rooms was oriented 30 degrees North of East and is shown in Figure 1-2. Four of the rooms are located on corners of the building (including the two kitchenettes). Two offices per floor have larger window to wall ratios than the others (0.8 compared to 0.4). In addition, the kitchenettes both have very high window to wall ratios ( $>0.9$ ) and have Southeast-facing primary facades. The sensor schedule for the measured rooms and a diagram of typical sensor locations is provided in Table 1-1 and Figure 1-2

respectively. The instrumentation error values in Table 1-1 correspond with manufacturer specified random error. Further details of the in-situ sensing and control capabilities of the Health Sciences Building are discussed in Chapter 4 and Chapter 6.



**Figure 1-2: Building façade viewed from North (top) and diagram of conditioned space with zone-level HVAC system and select sensors (bottom)**

**Table 1-1: Sensor schedule for the rooms studied in Chapter 4 and Chapter 6 with sensor type and manufacturer-provided instrumentation error**

Measured Quantity	Sensor	Instrumentation error
Blind position	Controller output feedback	
Hour of day	Internal clock	
Illuminance near window	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance in center of room	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance near wall	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance in center of room	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Indoor air carbon dioxide concentration	Dual-beam nondispersive infrared	$\pm(50 \text{ ppm} + 2\% \text{ of value})$
Indoor air relative humidity	Integrated relative humidity and temperature	$\pm 3\%$
Indoor air temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Lighting power	Current transducer	$\pm 2\%$ full scale (4.8 W)
Lighting state	Controller output feedback	
Occupancy state	Passive infrared	
Outdoor air temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Plug load	Current transducer	$\pm 0.5\%$ of value
Radiant heater flow rate	Single jet flow sensor	$\pm 5\%$ full scale (30 l/h)
Radiant heater inlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Radiant heater outlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Radiant heater valve position	Controller output feedback	
Reheat coil flow rate	Single jet flow sensor	$\pm 5\%$ full scale (30 l/h)
Reheat coil inlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Reheat coil outlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Reheat coil valve position	Controller output feedback	
Return air temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
VAV airflow	Differential pressure transducer	$\pm 1\%$ full scale (5 Pa)
VAV supply air temperature	4-20 mA Resistance temperature detector	$\pm 0.2^\circ\text{C}$
Window state	Wireless dry contact	

The Health Sciences Building has a variable air volume (VAV) zone-level mechanical ventilation system with air handling units (AHU) on each of the seven floors of the building. The hydronic system is configured such that the heating and cooling modes

cannot operate simultaneously, but the VAV reheat coils remain active for zone-level reheating of supply air in the cooling season. The zone-level reheat coil valve is modulated to heat the zone-level supply air temperature (after the reheat coil) to a temperature high enough to meet the room-level indoor air temperature setpoint of the room with the highest indoor air temperature setpoint in the zone in both the heating and cooling seasons.

Heating and cooling systems that use VAV equipment can have complicated control algorithms that use a combination of PI loops, scaling functions, and schedules. In the Health Sciences Building the AHU supply air temperature scales inversely between the minimum and maximum setpoint (15°C and 21°C, respectively) based on the OAT within the range of -10°C to 20°C using an OAT bias of 0°C. When the input of a control loop is equal to the setpoint, a proportional control loop output is equal to the bias (which is typically 50% within the total output range of 0-100%). The AHU outdoor air damper in the Health Sciences Building is either fully closed or fully open based on a static occupancy schedule (this may not be the case in other buildings). At the time of this study the occupancy schedule for the AHU is set to be always active (24 hours per day, 7 days per week). The VAV flow is controlled based on the heating or cooling requirement of the spaces within the zone based on the room-level PI loop with the highest demand (heating or cooling) during the occupied portion of the schedule. The room-level PI loops use the indoor air temperature as the input and indoor air temperature setpoint as the setpoint (and static P and I parameters).

Few mechanically ventilated facilities have control and measurement capability at the zone and room level comparable to the Health Sciences Building at Carleton

University. Herrick Labs at Purdue University has extensive instrumentation to monitor occupants and the air supply system, with a focus on occupant and indoor environment quality studies (Deng and Chen 2019). The National Renewable Energy Laboratory's Research Support Facility is another commercial building with in-situ research applications, largely in occupant comfort and net-zero energy (Lobato, Pless et al. 2011). The Canadian National Research Council's Canadian Centre for Housing Technology research facility contains several highly instrumented residential buildings and one small commercial building that can measure supplied energy from heating and cooling systems (Zirnhelt and Richman 2015). FLEXLAB at Lawrence Berkeley Laboratory can evaluate the energy efficiency of major building systems as an integrated system under real-world conditions (McNeil, Kohler et al. 2017). The Health Sciences Building at Carleton University combines some of the best features for in-situ testing that were introduced with the previously described facilities, namely:

- 25 near-identical offices.
- BAS-connected sensors that measure most energy flows for each room.
- Realistic occupants (not research participants).
- A well-labelled and research-accessible BAS and data acquisition interface.

The general experimental framework for this research relies on the use of 27 rooms with numerous BAS-connected sensors. These rooms were all used for the proxy sensor study (Chapter 4) and divided into test and control groups for the wide-scale supervisory control study (Chapter 6).

### 1.3 Thesis Outline

The body of this thesis consists of five main chapters: (Chapter 2) a brief history of building controls, (Chapter 3) an introduction of contemporary data-driven modelling approaches to building modelling, (Chapter 4) a methodology for the selection of zone-level sensors for energy use estimation, (Chapter 5) a pilot study using classification learning to emulate MPC, and (Chapter 6) a wide-scale study advancing the approach from the pilot study to 27 rooms with in-situ testing. Topic-specific literature reviews are contained within the three manuscript chapters (Chapters 4, 5, and 6). Chapters 2 and 3 are intended to provide supporting information that fortifies the approaches used in this thesis. The details of the experimental approach for the papers contained within this thesis are described in detail within their respective chapters.

**Chapter 2** introduces the progression of building controls from electrification to modernity. Market trends and the market share of communication protocols are also shown to justify the impact potential of the approach to supervisory control proposed in this thesis.

**Chapter 3** reviews common data-driven modelling approaches for building modelling and control applications from the literature. Most of these approaches are used in the following chapters with corresponding references. The additional approaches (support vector machines and reinforcement learning) are introduced and discussed as they pertain to building control.

**Chapter 4** presents a methodology to ranking zone-level sensors with respect to energy use estimation over the heating and cooling seasons. Hierarchies are generated using the proposed approach and data from the Health Sciences Building for end uses (e.g., lighting,

radiant heating) and individual sensors (e.g., indoor air thermistor, lighting current transducer). This chapter is under publication review as Article 1. With the theory and experimental set-up established in Chapter 4, it was possible to implement and quantify the performance of MPC.

**Chapter 5** presents a single-office study comparing the heating season energy use of unmodified reactive control, a basic MPC algorithm, and an emulation of the MPC algorithm generated using classification learning and an embedded temperature model called simplified model-based predictive control (sMPC). This chapter has been published as Article 2 (Bursill, O'Brien et al. 2019).

**Chapter 6** extends the study in the previous chapter to two thermally similar floors of offices in an institutional building with real occupants for both the heating and cooling seasons. The parallel testing used in this chapter reduced the disturbances introduced by ambient conditions and facilitated an aggregate analysis of multiple adjacent zones. This chapter has been published as Article 3 (Bursill, O'Brien et al. 2020).

**Chapter 7** concludes the previous chapters. The contributions of this thesis are highlighted, and recommendations for future work are presented.

## **Chapter 2: A Brief History of Commercial Building Controls and Modern Trends**

Commercial building controls have existed in some form since the electrification of buildings. Building control began with an interface as simple as an on and off switch without programmable logic. The evolution of hardware from the first commercial lighting system to modern BASs has been rapid due to the use of microcontrollers. However, user satisfaction and energy use efficiency have not advanced at the same pace as hardware. This stagnation could be in large part due to software and embedded logic limitations.

The purpose of this chapter is to demonstrate that despite over a century of work in building automation there are still major shortfalls in BAS performance. Improvements in technology and the adoption of fully programmable controllers increased the potential of supervisory algorithms but also presented additional risk potential when programming was poorly implemented. This chapter concludes with a brief analysis of modern trends and the potential direction of building controls that support the research conducted in the manuscripts presented later in this thesis.

### **2.1 Antiquity to Pneumatics**

As early as 1882, Thomas Edison began developing an incandescent electric lighting system. This was a problem the research community had been working on for over 40 years (Widagdo 2006). Edison built upon the work of many prior researchers while developing a scalable solution to electric lighting. Perhaps one of the most difficult challenges for the proliferation of electric lighting was the required change to physical and organizational

infrastructure (i.e., changing from gas to electric lighting and the accompanying supply chains) (Hargadon and Douglas 2001). Meeting the expectations of users and installation staff was important. This process took time, innovation, and the creation of several companies.

Around the same time as Edison invented the light bulb, Warren Johnson invented the thermostat based on a round, hand-adjusted, mercury-based thermal switch (Johnson 1883, Peksenar 2013). These thermostats were initially used to control non-residential building heating equipment. In the earliest implementation, the thermostat would trigger a light that would signal when the janitors would have to shovel more coal into the boilers serving the building (Epstein 2015). Over the next decade (the 1880s), mechanical compressed air systems were used to automate the heating systems (Epstein 2015). Thus, the commercial BAS was born.

Heating was only one half of what was desired from commercial HVAC systems. Building again on prior work, Willis H. Carrier developed the first modern air-conditioning system in 1902 to solve a production problem at the Sackett & Wilhelm's printing plant in Brooklyn, New York using mechanical cooling (Carrier 2019). Thermostatic control (previously developed by Johnson) allowed for temperature-based regulation of the vapor compression process (Carrier 1926). The proliferation of illuminated, cooled, and automated buildings brought humanity a level of comfort in the built environment that was previously unimaginable.

The commercialization of lighting, control, and cooling systems were perhaps the most significant non-electronic BAS developments to date. The next phase of commercial

building control advancement would be based on the transition from mechanical pneumatic systems to digital systems due to computing advancements in the digital revolution.

## **2.2 Pneumatics to Digital Control**

Leading up to the 1970's commercial building control had expanded using pneumatic sensors and actuators (Epstein 2015). The improvements in quality and price reductions of microprocessors in the digital revolution (spanning the 1950's to 1970's) drove the transition from pneumatic controls to digital controls. Digital controls offer better reliability, more robust programming environments, and accompanying energy savings by replacing pneumatics with electric actuators (Bengea, Kelman et al. 2012). One of the most notable programming enhancements that emerged from the switch to digital control was the programmable PID controller, which remains the basis of control in modern BASs for most processes (e.g., valve control, damper control, setpoint control).

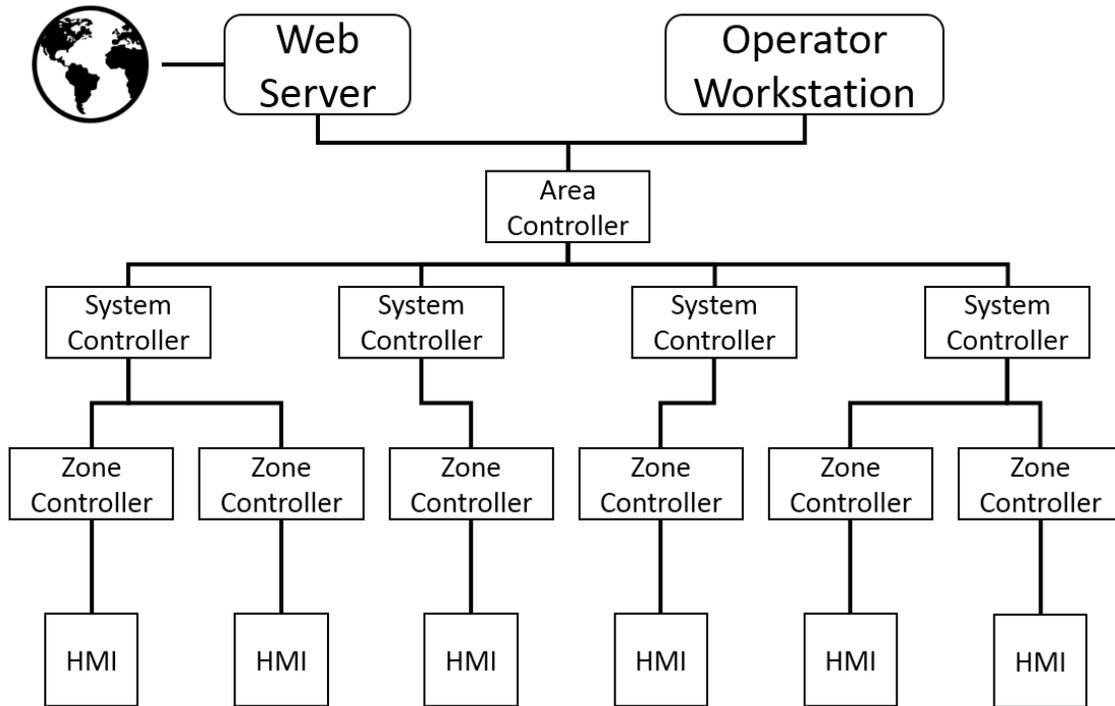
Digital control proliferation did not come without challenges. Previous communication driven by direct analog input signals, output signals, and manual interfaces was replaced with numerous digital communication protocols (with limited interoperability between different protocols). Some protocols have published standards that define data packet structures and database property value ranges that are intended to be vendor agnostic. Vendor-agnostic protocols that are commonly used for digital building controls are Modbus (MODBUS 2018), LonTalk (ISO 2012), and BACnet (ASHRAE 2017). LonTalk and BACnet are maintained and updated by the International Organization for Standardization and the American Society of Heating, Refrigerating and Air-Conditioning Engineers respectively. Modbus is an older protocol than the other two and is typically

used for lower-level devices that do not require complicated programming (e.g., meters), and is maintained by the Modbus Organization (MODBUS 2018). Modern wireless protocols such as Zigbee and EnOcean also add to the complexity of machine to machine communication but facilitate the use of additional low-cost sensors.

In addition to the previously discussed non-proprietary protocols, there are still many proprietary protocols in use. Inter-protocol communication is an ongoing problem in building control and has resulted in the creation of inter-vendor products. Protocol gateways are a solution that can provide mixed results. In many cases, lower levels of systems with different protocols cannot be accessed even with gateways. Thus, the transition from simple pneumatics to digital control has solved some problems but introduced others.

In addition to gateways, modern improvements to data storage and cloud computing access provide promising solutions to protocol and semantic issues within older BASs. Figure 2-1 shows the controller network topology of an ideal digital control-based BAS, where the operator workstation and a web server exist in parallel to enable both building operation and cloud data storage. The operator workstation may also be connected to the internet, but that is not a necessary requirement (particularly for older systems that predate high-speed internet). Sensors would connect in Figure 2-1 to the human-machine interface (HMI), zone controllers, or system controllers. HMIs can be colloquially called thermostats and in this context are used by building occupants, but the term HMI when describing network topology is more robust for terminal user interfaces in this system description. System controllers would typically be connected to large equipment (e.g., air handling

units). The emergence of low-cost data storage and web-enabled BASs in the early 2000's provided the foundation for modern building control research. This “data expansion” is the next topic of discussion.



**Figure 2-1: Typical digital control BAS network topology**

### **2.3 The Data Expansion: Modern Trends and Future Directions**

Over the last decade, the volume of sensors and controllers that can interface with BASs has expanded significantly due to research and commercial interest in the internet of things (IoT) (Kitchin 2014, Bajer 2018). Combining sensors over multiple protocols with remote, long-term data storage has generated a copious amount of data for researchers to analyze.

Sensing and computing technology have advanced significantly in the last decade, and the result has been lower controller costs and more data storage capacity and access. For example, controllers from the vendor for the buildings used in this research had controller memory grow 10 times and processing power grow 1000 times between the 2000's and 2010's. There is now far greater potential for supervisory control approaches when the new control hardware is coupled with web-enabled BAS servers.

Improvements in controller hardware have been concurrent with an increase in sensor density. On the campus where this research was undertaken each building has roughly 1,500 points (i.e., inputs and outputs), where roughly two-thirds are inputs. System controllers that are connected to large HVAC equipment (e.g., boilers, air handling units) will typically have dozens of points, while a zone or room-level controllers will have less than 10 (usually comprised of room temperature and terminal HVAC control and feedback). Newer buildings will often feature thermostats that include relative humidity and carbon dioxide sensors which can triple the point count for zone-level controllers.

While standalone sensor costs have decreased over the last decade, BAS integrated sensors remain costly with respect to the value of the energy that they measure due to labor costs associated with installation. Table 2-1 provides estimates of cost (excluding installation), energy savings, simple payback, and sensor life based on information from several sources (King 2017). Table 2-2 provides similar information for plug load and lighting sensors and Table 2-3 for window shading technology. Even without installation costs included, most sensors do not make an obvious business case for installation. When labor is accounted for, cost per point can reach upwards of \$900 for a fully programmable

BAS in North America (Zeitoun 2011). The high cost for BAS-connected sensors highlights the need for an approach to optimize sensor selection to only install the most useful physical sensors for the goals of the control system.

**Table 2-1: Aggregated cost information and energy savings estimates for HVAC sensors**

Components	Category	Technology	Cost	Estimated energy savings from reference	Simple payback	Measure life
Energy, temperature, flow, pressure, humidity sensors <sup>1</sup>	HVAC	Wired sensor	\$50–100/sensor + \$5.25/linear meter wiring	Not applicable	Not applicable	15–30 years
Energy, temperature, flow, pressure, humidity sensors <sup>2</sup>	HVAC	Wireless sensor	\$150–300/sensor	Not applicable	Not applicable	15–30 years
Smart thermostat <sup>3</sup>	HVAC	Smart thermostat	\$150–330/thermostat	5–10% HVAC	3–5 years	10 years
Door switches, occupancy sensors <sup>4</sup>	HVAC & lighting	Hotel guest room occupancy controls	\$100–500/guest room	12–24% HVAC, 16–22% lighting	2.5–3.0 years	10 years

---

<sup>1</sup> Source: Kintner-Meyer, M., M. R. Brambley, T. A. Carlon and N. N. Bauman (2002). Wireless sensors: technology and cost-savings for commercial buildings. Teaming for Efficiency: Proceedings of the 2002 ACEEE Summer Study on Energy Efficiency in Buildings.

<sup>2</sup> Source: Shoemaker, S. (2015). "Innovative, Lower Cost Sensors and Controls Yield Better Energy Efficiency." from [www.ornl.gov/news/innovative-lower-cost-sensors-and-controls-yield-better-energy-efficiency](http://www.ornl.gov/news/innovative-lower-cost-sensors-and-controls-yield-better-energy-efficiency).

<sup>3</sup> Source: Grant, C. K., P (2016). Smart Thermostats: An Alternative to Load Control Switches? Arlington, VA.

<sup>4</sup> Sources: Kintner-Meyer, M., M. R. Brambley, T. A. Carlon and N. N. Bauman (2002). Wireless sensors: technology and cost-savings for commercial buildings. Teaming for Efficiency: Proceedings of the 2002 ACEEE Summer Study on Energy Efficiency in Buildings, King, J. P., Christopher (2017). Smart Buildings: Using Smart Technology to Save Energy in Existing Buildings. Washington, DC, American Council for an Energy-Efficient Economy: 55.

**Table 2-2: Aggregated cost information and energy savings estimates for plug load and lighting sensors**

Components	Category	Technology	Cost	Estimated energy savings from reference	Simple payback	Measure life
Smart plug (120 V) <sup>5</sup>	Plug load	Smart plug	\$100 each	50–60%	4–12 months	9 years
Smart plug (220 V)			\$200 each			
Tier One advanced power strips <sup>6</sup>	Plug load	Advanced power strip	\$45–50 each	25–50%	8–18 months	10–20 years
Occupancy, daylighting, task tuning, lumen maintenance, dimming, daylighting <sup>7</sup>	Lighting	Advanced lighting controls	\$22–43/m <sup>2</sup>	45%	3–6 years	10–20 years
Software and hardware <sup>8</sup>	Lighting	Web-based lighting management system	\$12.4/m <sup>2</sup>	20–30% above controls savings	1–4 years	10–15 years

**Table 2-3: Aggregated cost information and energy savings estimates for window shading technology**

Components	Category	Technology	Cost	Estimated energy savings from reference	Simple payback	Measure life
Shades w/ automatic controls <sup>9</sup>	Window shading	Automated shade system	\$375 (motorized shades)	21–38%	4 years	10–20 years
Self-adhered film <sup>10</sup>	Window shading	Switchable film	\$160-220/m <sup>2</sup>	32–43%	2–3 years	10 years
Thermochromic smart glass <sup>11</sup>	Window shading	Smart glass	\$430/m <sup>2</sup>	20–30%	21 years	30 years
Electrochromic smart glass			\$660/m <sup>2</sup>		33 years	50 years

<sup>5</sup> Source: Boss (2016). Case Study: City of Pittsburgh Plug Load Energy Management. Pittsburgh, PA, Boss Controls.

<sup>6</sup> Source: GSA (2012). Advanced Power Strips. Washington, DC, General Services Administration.

<sup>7</sup> Sources: Gilliland, D. (2016). A Field Evaluation of Cloud-Based Predictive HVAC Control Software—The Next Step in Energy Efficiency, or Marketing Mumbo Jumbo? 2016 ACEEE Summer Study on Energy Efficiency in Buildings. Washington, DC: 1-13.

<sup>8</sup> Source: *ibid*.

<sup>9</sup> Source: Wagner, A., T. Lützkendorf, K. Voss, G. Spars, A. Maas and S. Herkel (2014). "Performance analysis of commercial buildings—results and experiences from the german demonstration program 'energy optimized building (EnOB)'." Energy and Buildings **68**: 634-638.

<sup>10</sup> Source: InvisiShade. (2016). "FAQ." Retrieved December 11, 2019, from [www.invisishade.com/](http://www.invisishade.com/).

<sup>11</sup> Sources: Lutron. (2014). "Lutron Energy Savings Claims." from [www.lutron.com/TechnicalDocumentLibrary/Lutron\\_Energy\\_Savings\\_Claims.pdf](http://www.lutron.com/TechnicalDocumentLibrary/Lutron_Energy_Savings_Claims.pdf), InvisiShade. (2016). "FAQ." Retrieved December 11, 2019, from [www.invisishade.com/](http://www.invisishade.com/), SageGlass. (2016). "SageGlass." Retrieved December 11, 2019, from [www.sageglass.com/products](http://www.sageglass.com/products).

Above the physical sensor level in a BAS network, one must consider the viability of the connected controllers and communication protocols for industry application. Despite the expansion of the BAS field due to IoT devices and gateways, the backbone control of most commercial buildings remains limited to a few standardized and proprietary protocols. Buildings are typically expected to last up to 60 years, and control system replacements are not common within an average building lifetime (Biswas, Shrestha et al. 2016). This places importance on using vendor-agnostic protocols that will have longevity. A standardized control networking approach includes the ability to interface with newer controllers added during renovations, repairs, or retrofits. The buildings studied in this thesis primarily used the BACnet protocol.

The market share of BACnet devices was examined by BACnet International using Building Services Research and Information Association (BSRIA) global market research report data. The resulting annual global and continental market shares of BACnet are shown in Figure 2-2 (BSRIA 2018). The trends in Figure 2-2 suggest that BACnet has a majority share of BAS protocols in North America and is growing in popularity worldwide. BACnet also makes up over half of the building communication protocol market share in Europe, where the BAS technology industry is more fragmented than in North America and LonTalk and KNX have significant market share (Dietrich, Bruckner et al. 2010). The aggregate market share information suggests that supervisory BAS technology should target BACnet devices at the commercial building level to have the maximum impact in an embedded implementation.

The growth of BACnet vendors and products were also assessed. Data was collected for active BACnet vendors recorded on the BACnet website from 1999 to 2019 to find the change in the number of BACnet vendors over the last 20 years (and the growth over the same period) using the Wayback machine (Eltgrowth 2009). The resulting data is shown in Figure 2-3 and Figure 2-4 where the categories have the following meanings as defined by the BACnet website (ASHRAE 1999):

- Controller: The vendor offers a device capable of performing digital control that uses BACnet as its "native protocol." This category also includes BACnet-capable sensors and actuators.
- Gateway: The vendor offers a device that translates messages between BACnet and some other protocol (usually the vendor's own proprietary protocol).
- Router: The vendor offers a device that allows the interconnection of different types of BACnet local area networks (e.g., Ethernet to MS/TP).
- Workstation: The vendor offers a workstation capable of directly communicating on one or more types of BACnet local area network; alternatively, the vendor may provide a web server that provides operator functions via a web browser. The listings currently do not distinguish between these two cases.
- Tools: The vendor offers a protocol stack, protocol analyzer, or some other form of technical support for BACnet implementers or users.

An additional category of Full Suite was added for vendors that offered all five categories. The sum of the vendors of the five categories will be greater than the total number of vendors because many vendors have more than one product category.

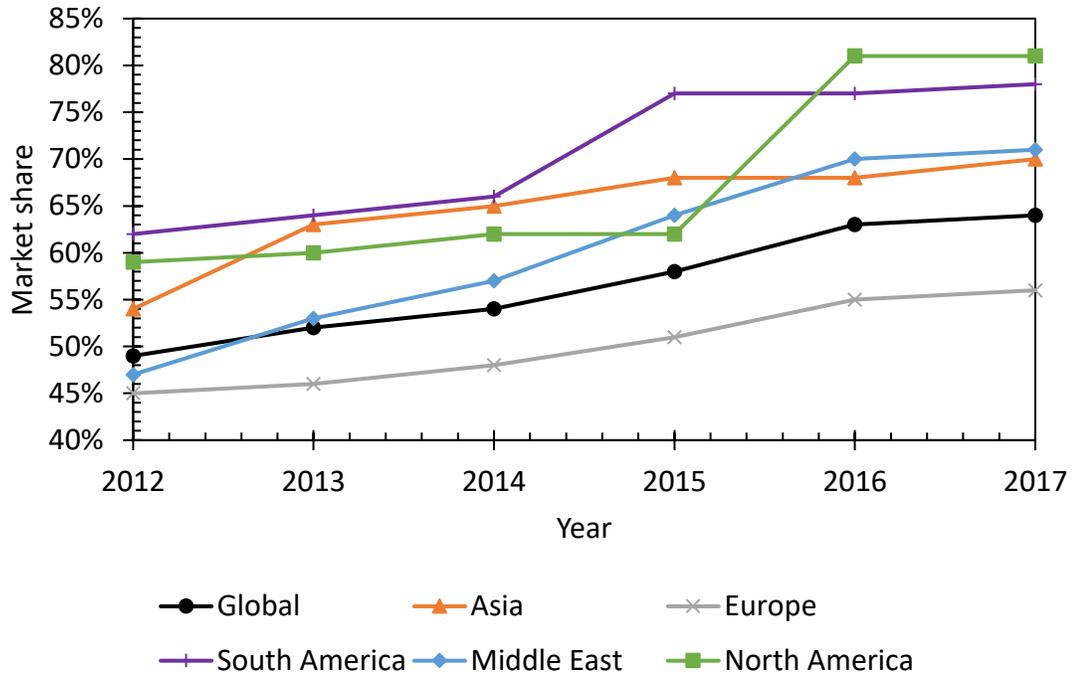


Figure 2-2: BACnet commercial building controls market share of new equipment from 2012 to 2017

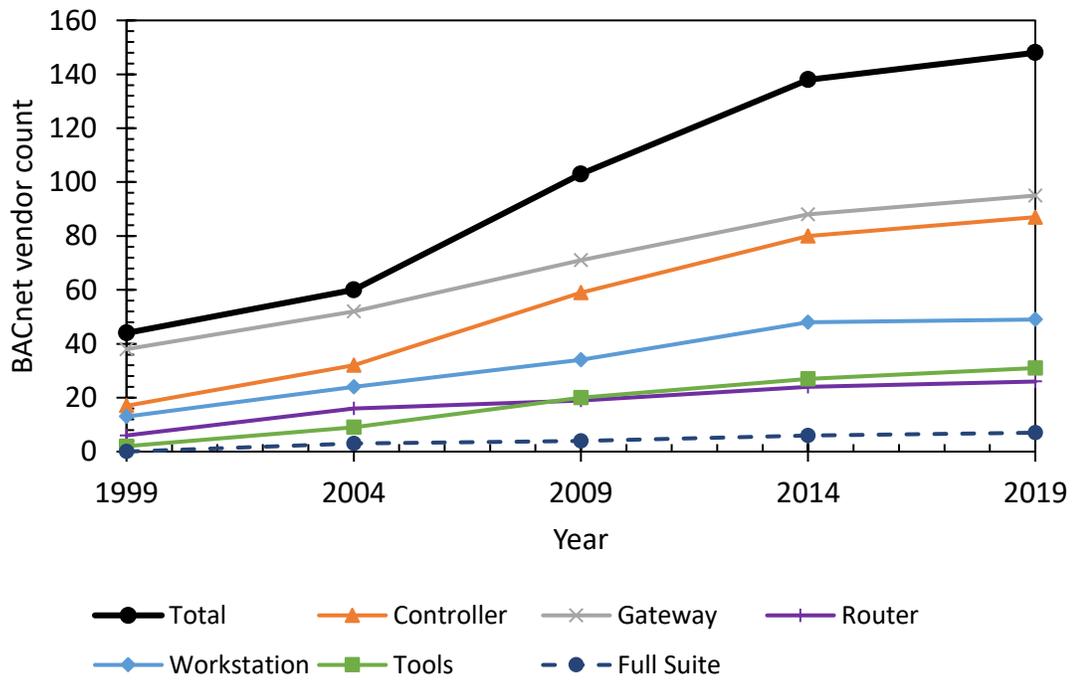
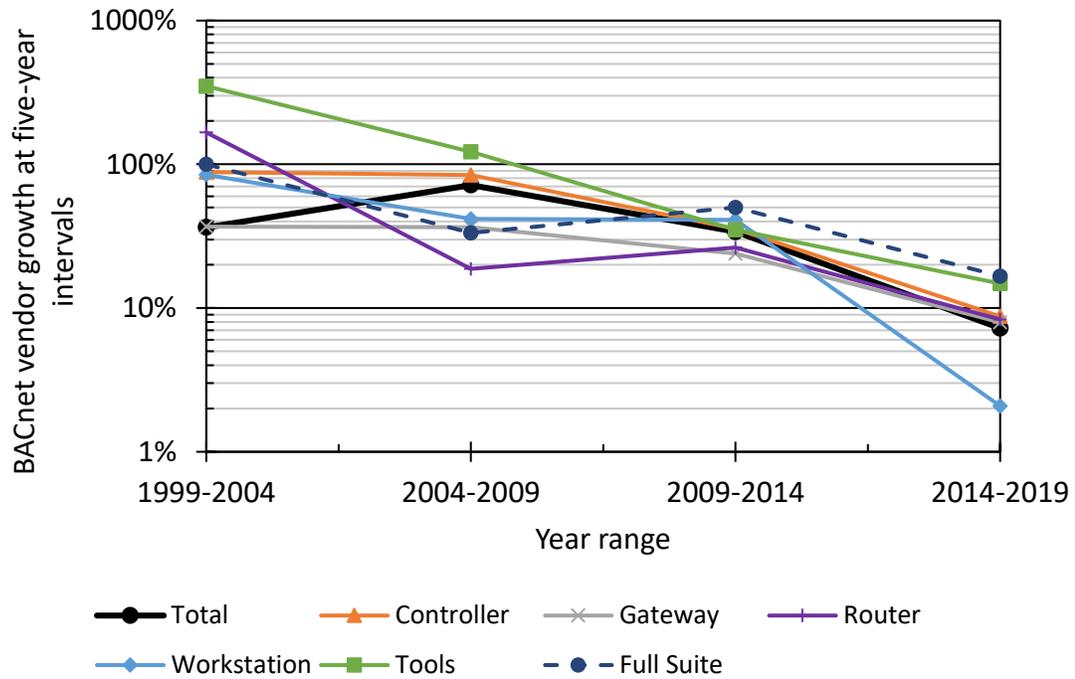


Figure 2-3: BACnet vendor counts from 1999 to 2019



**Figure 2-4: BACnet vendor growth at five-year intervals from 1999 to 2019**

There are currently over 1,100 vendor IDs (i.e., unique identifiers given to potential BACnet product developers), which compares to ~150 active vendors in 2019 (BACnet 2019). Most of the vendor IDs are abandoned, placeholders for BACnet usage, or for prospective or defunct vendors. The overall growth of BACnet vendors is slowing as shown in Figure 2-4, and as of 2019, a few vendors control large portions of the market with full suite offerings. These full-suite offerings are appealing to large commercial building projects (e.g., public-private-partnerships and design-builds). Figure 2-3 and Figure 2-4 visualize the consolidation among BAS manufacturers, where the industry is dominated by a few large players and there are significant barriers to entry.

Since 1999 when BACnet tools were not offered, tools have grown the fastest relative to controllers, gateways, routers, and workstations. However, this is also evidence that the high-level software side of BACnet is still in relative infancy compared to firmware and hardware. The previously discussed modern trends suggest that software-based control solutions are the most promising area of growth within BACnet controls (which occupies most of the North American and global markets) because they can be applied to multiple vendor's products. The proliferation of BAS connected sensors and high-volume data storage provides an opportunity to use large amounts of data to train models with machine learning for embedded BAS implementation if a vendor-agnostic approach is applied effectively.

Despite the advances over the last century, many modern BASs still have trouble reliably reading sensors and controlling spaces within the desired temperature ranges. BAS-connected sensors are expensive to install and have dubious reliability without frequent human monitoring. Significant energy waste can occur when lower level (i.e., zone-level) control is neglected (Gunay, O'Brien et al. 2016, Gunay, O'Brien et al. 2018, Dong, Prakash et al. 2019). These issues need to be addressed before larger scale integrations of protocols (e.g., IoT) and smart building initiatives can be successful. The following chapter discusses popular applications of machine learning to commercial buildings with an emphasis on the different applied techniques to address contemporary building control challenges.

## **Chapter 3: Data-driven Applications for Commercial Building**

### **Supervisory Control**

Data-driven approaches to modelling present practical applications of the abundance of data recorded in modern building automation systems, and facilitates the use of machine learning algorithms to discover non-intuitive relationships within the data (Cuzzocrea, Song et al. 2011, D'Oca and Hong 2014, Kitchin 2014, Afram and Janabi-Sharifi 2015, Jung, Kim et al. 2015, Amasyali and El-Gohary 2018, Wei, Zhang et al. 2018, Causone, Carlucci et al. 2019, Maddalena, Lian et al. 2020). This chapter briefly introduces and discusses several data-driven approaches to modelling and their applications to supervisory control. In addition to what is presented, many data-driven approaches exist that are not covered in this chapter. Consequently, this chapter is intended to provide a brief contrast of the approaches used in this thesis and introduce approaches that could be used in future work. The first section of this chapter provides background to the goals, approaches, and relative accuracy of five commonly used control-oriented data-driven modelling approaches. The sections that follow will reference where they are discussed in further detail in Chapters 4, 5, and 6 or present supplementary literature describing applied usage.

#### **3.1 Data-driven Approach Background**

Data-driven approaches are used frequently in this thesis to demonstrate practical applications of the abundance of BAS data available in buildings. Data-driven approaches are based on the concept that enough data is available to accurately model a system for specific applications when given input and output data from a system and limited

knowledge of the underlying physical behavior are available (Solomatine and Ostfeld 2008, Kramer, van Schijndel et al. 2012). Compared to detailed physics-based modelling, data-driven modelling is in its infancy due to the additional computational power required. The appropriate volume of input data varies depending on the problem. Approaches are taken in Chapter 4 of this thesis to reduce the volume of data required to estimate energy use by reducing the width (i.e., number of inputs) of the two-dimensional input data array.

Data-driven approaches require a rigorous modelling workflow and enough data to have some form of model validation. A generalized and simplified process for data-driven modelling is introduced in Bursill et al., where data cleaning (i.e., identifying, reconciling, and removing outlying and erroneous data) makes up a significant portion of the analysis effort (Bursill, O'Brien et al. 2019). Once data is cleaned and in a suitable format, an iterative process of model fitting and model selection must occur to produce an applied model. Model selection and model fitting are discussed in detail in Chapters 4, 5, and 6 of this thesis as it pertains to energy use estimation models and indoor air temperature models for individual rooms.

Models can generally be divided into stochastic and deterministic forms (Reddy 2011). Data-driven models rely on using stochastic models to manage accuracy and overfit (i.e., when a data-driven model corresponds too closely to one set of data and is not indicative of the overall trend of the data). While stochastic models are more robust by accounting for random error, they are also more computationally intensive due to the need for applied Bayesian approaches and underlying random number processing (Reddy 2011, Oldewurtel, Parisio et al. 2012). Once the application of random error is chosen, regression

and classification approaches for model fitting must be defined. Linear regression using least squares is the most common form of model regression and is frequently a suitable solution compared to more complex regression approaches (e.g., logistic, quadratic, weighted) (Thompson 1982, Zhao, Chen et al. 2001). Likewise, a classification approach must also be defined for data-driven models. One simple example is the classification and regression tree (CART) method used in classification trees (Tsanas and Xifara 2012). In some cases, the classification is entirely automated (e.g., classification tree random forests), and in others, a classification is defined and must be manually augmented (e.g., model form in state-space models). Applied model classification is described in further detail in the manuscript-adapted chapters of this thesis.

Building upon the previously described inverse model details, the following sections provide cursory information and examples of data-driven modelling approaches. Data-driven modelling procedures as they pertain to the studies in this thesis are described in detail in Sections 4.3.4 (artificial neural network modelling), 5.3.2.1 (state-space modelling), and 5.3.2.2 (classification learning). In addition, support vector machines and reinforcement learning are introduced in this chapter in Section 3.5 and 3.6 respectively due to their extensibility to anomaly detection and adaptation to changing system parameters. However, support vector machines and reinforcement learning approaches were not applied in the manuscript studies in this thesis and are discussed to provide additional context to the proposed future work in Chapter 0.

### **3.2 Artificial Neural Network Models**

Artificial neural network (ANN) models are based on neuron structures in the brain and can be applied to complicated and ill-defined problems (Kalogirou 2000). One such application in this thesis is in Chapter 4, where ANN models are used to estimate energy use using indirect input data. Input signals, excitation functions, inter-neuron relationships, and output functions form a highly nonlinear model between multiple inputs and a single output (Ahmad, Hassan et al. 2014). ANN model structures are flexible and defined by the modeller, which enables both flexibility and the risk of overfitting. Several forms of ANN models exist for various applications (e.g., convolution neural networks for image processing), and training approaches can differ depending on the data volume and model application (e.g., feed-forward, backpropagation) (Amasyali and El-Gohary 2018).

Exploratory modelling for proxy (i.e., indirect) sensing was conducted in Chapter 4 using feed-forward ANN models with structures defined by the number of inputs and multiple hidden layers. Hidden layers allow for additional secondary relationships to be modelled between inputs and outputs. Overfit is identified by applying cross-validation and reduced by adjusting the ANN model form. This section provides a brief introduction of ANN models as they relate to the analysis in Chapter 4. Further details on the background of ANNs and the relevant application to this research are provided in section 4.3.4.

### **3.3 State-space Models**

State-space models can be used to model multi-dimensional nonlinear systems in the time domain. Compared to physics-based models, state-space models are a relatively new

model form (Kalman 1960). State-space models originate from transitioning frequency domain control principles to the time domain to allow for multiple input and multiple output systems. Stochastic and nonlinear control problems can be approximated with state-space models analytically or using data-driven modelling (Koller and Friedman 2009). Inverse state-space models are of interest as it pertains to this thesis to reduce model complexity and post-training computation time. Many of the features of state-space models are not used in this thesis to reduce model complexity and maintain pseudo-grey-box model forms that can be qualitatively assessed with respect to physical model inputs (i.e., a large parameter for an outdoor air temperature input can indicate the model's sensitivity to outdoor air temperature). In this thesis, multiple input single output linear time-invariant state-space models were used to emulate single node grey-box indoor air temperature models in Chapters 5 and 6 (Gouda, Danaher et al. 2002).

Grey-box model forms can be adapted to state-space models with some algebraic rearrangement and the use of linear time-invariant parameters. For example, a lumped capacitance air temperature model (Athienitis 1994) can be used to formulate a similar control-oriented state-space model (Gunay, O'Brien et al. 2016, Bursill, O'Brien et al. 2019, Bursill, O'Brien et al. 2019). Details of the approach for applying state-space models as surrogates to traditional grey-box models are provided in the adapted manuscripts in Sections 5.3.2.1 and 6.2.3.

### **3.4 Classification Trees**

If target output data for a dataset is in a discrete form, classification learning can be effective at finding inputs and thresholds to emulate the underlying physical system (Van

Der Heijden, Duin et al. 2005, Wei, Zhang et al. 2018). This principle is applied in Chapters 5 and 6 to simplify MPC algorithms from having dozens of inputs to only a handful while maintaining similar energy use and occupant comfort to the original MPC. The output of a CART decision tree resembles a flow chart and can be easily abstracted by non-experts when visualized, and conversion to a Boolean code structure is clear (and an output of the MATLAB Classification Learning Toolbox functions). Classification trees are easier to implement and analyze than regression trees but are not applicable to continuous data. When using classification trees, confusion matrices can clearly visualize error in the form of correct predictions, false-positives, and false-negatives. These features can be used to adjust the training algorithm to reduce false-positives or false-negatives depending on which may have a more significant negative impact on the system (Van Der Heijden, Duin et al. 2005).

Classification trees are applied and described in detail in this thesis in Sections 5.3.2.2 and 6.2.4 where the discrete override commands from a detailed MPC algorithm to a reactive control commercial BAS controller are emulated in less than 30 lines of controller-embedded code.

### **3.5 Support Vector Machines**

Support vector machine (SVM) approaches were not used in the manuscripts within this thesis, but are a common alternative approach to the ANN models used in Chapter 4 and the classification trees used in Chapters 5 and 6. An SVM is a kernel-based machine learning algorithm, which can be used for both regression and classification (Wu, Kumar et al. 2008). The SVM approach is based on applying a non-linear map to a high-

dimensional form of the problem, followed by reducing the problem to a linear map from the non-linear map using a kernel function (Amasyali and El-Gohary 2018). SVMs have been shown to achieve similar energy use prediction accuracy to ANN models in recent studies from the literature (Ahmad, Hassan et al. 2014, Fu, Li et al. 2015, Jung, Kim et al. 2015, Zhang, Deb et al. 2016, Amasyali and El-Gohary 2018). The previously mentioned studies in the literature provide more thorough descriptions of the underlying theory and applications of SVM approaches than are relevant to the scope of this thesis.

One notable advantage that SVMs have over classification trees and ANN models is visual anomaly detection applications (Chen, Chen et al. 2011, Konar and Chattopadhyay 2011). If one of the objectives of analysis is anomaly detection (e.g., fault detection, data quality checking), then the clustering structure of SVMs can be used with expert knowledge to identify outliers in the data. Clusters within SVM fitting algorithms are intermediate parameters and can be used in conjunction with the model fit to describe the output data.

### **3.6 Reinforcement Learning**

Reinforcement learning (RL) approaches can provide a model-free (i.e., completely based on an action and reward structure with no explicit predictions of the next state) alternative to MPC. Reinforcement learning can also be applied to other data-driven approaches to tune models in model-based approaches (e.g., improving classification trees using prediction error feedback to adjust thresholds). The concept of supervisory RL nesting supervisory MPC (or simplified classification trees) is a topic that will be discussed in the future work of Chapter 0.

Reinforcement learning-based algorithms continuously learn from different operating conditions via physical interactions with system inputs. This allows RL algorithms to continuously improve over time (Yang, Nagy et al. 2015). RL is considered neither a supervised or unsupervised machine learning approach because it uses actions and perceived consequences (i.e., trial and error) to develop policies. These policies can be used in conjunction with supervised learning to achieve faster policy convergence. More thorough explanations of RL can be found in the literature (Szepesvári 2010, Sutton and Barto 2018).

Q-learning is the most common RL approach in buildings and is based on an update process. In a simulation environment, this process can be implemented in controller logic (Li and Xia 2015, Wei, Wang et al. 2017, Wei, Zhang et al. 2018). Yang et al. implemented RL control in a building in Switzerland and found that the RL algorithm outperformed rule-based control by 10% with respect to heating demand, operation temperature, and regulation of ground temperature in the ground-source heating system (Yang, Nagy et al. 2015). RL usage in live buildings is still relatively untested and is an area that will need to be expanded in the literature, particularly when coupling RL and supervised learning approaches.

### **3.7 Data-driven Approaches Summary**

The previous sections of Chapter 3 related three data-driven approaches to work conducted in the manuscript chapters of this thesis and introduced two data-driven approaches that could be useful in future work in supervisory control. This chapter is not intended to be exhaustive and is a summary of data-driven approaches identified as common in the

literature or promising to future applications.

Compared to contemporary commercial building sensor uncertainty (e.g., upwards of 30% for hydronic heating energy measurements at the zone level), the differences in accuracy achieved between most data-driven approaches is negligible. This claim is verified by Deng et al. in a comparison of linear regression, SVMs, ANNs, and random forest models for predicting US commercial building energy use intensities (EUI), where the root-mean-square error (RMSE) varied between 111.3 kWh/m<sup>2</sup> and 122.7 kWh/m<sup>2</sup> for total EUI for all data-driven modelling approaches tested (Deng, Fannon et al. 2018). This corresponds to an RMSE difference of 5.2% of the average EUI in the study between the most and least accurate data-driven modelling approach, an order of magnitude lower than instrumentation error.

Perhaps the largest barrier to advanced algorithms gaining market penetration is the complexity of building-wide implementation by conventional control technicians. Control technicians are not data scientists or machine learning experts, and need algorithms delivered in a format that can be easily and quickly implemented in the field. This issue is paramount to solving the building control problem.

The following three chapters are adapted from journal publications that individually address segments of the building control problem. All three papers use forms of data-driven modelling. The three papers from which the following three chapters are based were published (or have entered the publication process) at different times, and as such, have some differences in terminology and nomenclature.

## Chapter 4: A Sensor Feasibility Study

This chapter is based on the supplementary material of:

Jayson Bursill, William O'Brien, Ian Beausoleil-Morrison. (In review). Proxy zone-level energy use estimation in a commercial building with a variable air volume system. *Manuscript submitted for publication.*

The usefulness of zone level sensors is contentious in the context of commercial building controls. Few building owners are willing to increase capital costs with conflicting information about the most valuable sensors to owners and operators. Few studies have systematically explored this topic in a practical manner. This study compared zone-level sensors of 27 rooms of an institutional building for their ability to estimate energy use indirectly using a novel methodology. Notable sensors in this study include hydronic flow meters, room level lighting and plug load current transducers, window contacts, and ceiling illuminance at three interior depths. Data was collected at five-minute sample intervals over four months total. Sensors were scored and ranked by energy use using mean bias error to infer estimation accuracy over the validation period (half the measurement period) with artificial neural network models. A hierarchy of sensor value for each energy end use was generated. It was found that outdoor air temperature sensors provide the most usefulness for energy use estimation in the building studied (where heating and cooling loads dominated due to the high outdoor air supply ratio). Hydronic flow meter data was also useful for energy use estimation.

## 4.1 Introduction

Building management systems (BMS) and their accompanying sensors are essential components to achieving energy savings within the building stock by providing data to quantify energy use. Energy use is a widely understood metric to compare building operation decisions and parameters, and sensing is required to estimate energy use (Wagner, Lützkendorf et al. 2014, Minoli, Sohraby et al. 2017). There are many contemporary research areas within commercial building sensing including: typical sensor errors, systematic errors due to working environments inside buildings or systems, periodic calibration, in-situ calibration, sensor redundancy, position, number, additional sensors for new applications (Yoon and Yu 2017, Yoon and Yu 2017, Yoon and Yu 2018). These issues compound to generate significant uncertainty and confusion about which sensors are worth installing. Lifetime BMS integrated sensors costs (including equipment, installation, programming, and maintenance) are typically at least an order of magnitude greater than the value of the energy they are monitoring at the zone-level (Ali, Zanzinger et al. 2016, Kumar, Martani et al. 2016). The recommendations of useful sensors for energy use estimation that exist in the literature consist of high-level analysis and anecdotes, primarily focused on unquantified occupant related benefits (Peters 2014, Winston 2016, North 2018).

There are several motivations for installing sensors down to the office level: tracking energy and comfort, trying to reduce energy use and improve comfort, security and occupant monitoring, identifying and resolving faults, and research. Most work to date on analyzing commercial building sensors has been at the system or plant-level, however,

many issues still occur at the zone-level (Riederer, Marchio et al. 2002, Dong, O'Neill et al. 2014, Goyal, Barooah et al. 2015).

Application examples can be postulated from the previously introduced motivations for zone-level sensing in commercial buildings. Energy monitoring uses sensors to quantify thermo-physical properties of air and water supplied to heating and cooling systems and the electrical energy drawn by lighting, plug in equipment, and electric heating and cooling equipment. Indoor environmental quality often accounts for thermal and visual comfort in buildings and is used to infer overall occupant comfort and can be inferred via the BMS from sensor readings such as indoor air temperature, carbon dioxide concentration, relative humidity, and ceiling illuminance (Spengler, Samet et al. 2001). Passive infrared (beam) occupancy, window contact, and door contact sensors can be integrated with security systems to monitor points of entry and occupant location (Baeg, Park et al. 2007). When considering both fault detection and research applications, all sensors present in the system can be used to provide valuable information to data driven models and expert classification systems (Katipamula and Brambley 2005). Some of the most common zone-level building sensors include (Schachinger and Kastner 2016):

- Indoor air temperature
- Passive infrared occupancy
- Air handling unit (AHU) supply air temperature
- Hydronic system supply water temperature
- Zone air flow rate

In addition to the sensors listed above, illuminance sensors are becoming more common in zone-level lighting systems to provide feedback for daylight harvesting and dimming control as required by modern energy standards (Page 2013, ASHRAE 2016). Additional sensors are often only installed in niche applications and require significant justification.

While sensors, communications, controls, and analytics technologies are rapidly advancing, sensors remain quite costly to install and maintain and the data often requires custom analysis (Wang and Xia 2015, Killian and Kozek 2016, Ruparathna, Hewage et al. 2016, Ferrari and Beccali 2017, Minoli, Sohraby et al. 2017). Therefore, this paper aims to reduce sensor costs by providing a framework for choosing the most valuable sensors for energy use estimation. The methodology and results presented in this paper are intended to be useful for researchers, owners, and operators to go through a selection process for novel ways to use zone-level sensing. Energy use was selected as the target for this study because it is a standard quantitative and actionable metric for owners and operators. Researchers can support owners and operators by applying the modelling approaches presented in this paper to compare sensor value.

The objective of this paper is to develop a methodology to assess which sensors are the most useful for quantifying zone-level energy use and apply this methodology to an institutional building. The following sections describe the motivations from the literature, the approach to generating the hierarchy, the results of the hierarchy with building specific anecdotes, and the conclusions and next steps of this study.

## 4.2 Literature Review

Discussions of the relative rankings of sensors in any form were sparse or tangential in a review of the contemporary research. Some studies did provide context to the current knowledge and goals obtained from past work. Shaikh et al. review optimized control system approaches while discussing sensors, and identified the need for future work to quantify energy savings more effectively (Shaikh, Nor et al. 2014). Dong et al. review smart building sensing systems for indoor environment control with a focus on occupant centric sensing (Dong, Prakash et al. 2019). Dong et al. suggest that cost is a major barrier for occupant sensing. Both review studies infer that sensor cost is an industry-wide challenge for all applications.

Versatile, multi-purpose sensors offer a solution to high sensor costs. One approach to developing multi-purpose sensors is commonly referred to as virtual sensing, where low cost measurements and mathematical models are used to estimate difficult to measure or expensive quantities (Li, Yu et al. 2011). This concept has been explored for many building sensor applications, from individual flow sensors (McDonald, Zmeureanu et al. 2014) to whole building sensor verification (Cotrufo and Zmeureanu 2016). Virtual sensing is one of numerous tags for the same concept. Due to the ambiguity between regular measurement derived values (e.g., temperature from a thermistor) and virtual sensing (e.g., flow from a temperature sensor), this concept will hereafter be referred to as proxy sensing.

Discussion of zone-level energy use proxy sensors are limited in the literature. At the AHU level, Li et al. summarized numerous studies to develop proxy sensing as a fault detection device and to quantify energy use (Li, Yu et al. 2011). Bursill et al. utilized an

inverse modelling approach to quantify the cooling energy use at the zone level for 16 offices with some success in aggregate applications (Bursill, O'Brien et al. 2018). Limitations of current research in energy use at the zone level are highlighted by Amasyali et al., where a keyword search literature review revealed that only 2% of energy estimation papers are applied to loads other than heating or cooling (Amasyali and El-Gohary 2018). Despite the progress in data driven proxy sensing in the literature, which sensors are useful at the zone level for occupant characterization and energy use quantification is unclear.

Many approaches to white- and grey-box proxy sensing were described previously. White-box models are physics based (i.e., significant knowledge of the governing equations and physical properties is required) and black-box models are purely data-driven (i.e., no knowledge of the governing equations and physical properties is required). Grey-box models offer a compromise between white- and black-box models, such that some thermo-physical properties may be maintained that can be used to compare, adjust, and validate the models. Black-box proxy modelling can provide a low cost and low physical information alternative to other approaches, both valuable traits in room- and zone-level environments where thermo-physical information is limited. Artificial neural network (ANN) models are robust, continuous black-box models that are useful for non-linear modelling (Afram and Janabi-Sharifi 2015). Artificial neural network models have been utilized in related applications to model whole buildings and predict energy end usage (Moon and Kim 2010, Ahmad, Hassan et al. 2014, Ahmad, Mourshed et al. 2016, Wei, Zhang et al. 2018). ANN models do not require detailed knowledge of a relationship between model inputs and outputs. This relationship is learned through a training process

(Yang, Rivard et al. 2005). Some success in energy use predictions (heating, cooling, and electrical energy) has been found at the building and zone level using ANNs (Wei, Zhang et al. 2018). Applications at the zone level have been limited in the literature, particularly at the spatial and temporal granularity explored in this paper (Jain, Smith et al. 2014, Dong, Li et al. 2016). Jain et al. utilized machine learning to predict energy use at the unit level in a New York City residential building with data samples at 10 minute intervals and found that hourly samples and floor-level data produced the most accurate models (Jain, Smith et al. 2014). However, in the aforementioned study, data granularity as described by Dong et al. was considered nearly impossible to obtain (Dong, Li et al. 2016).

Selecting an appropriate model to adapt to zone-level energy use is crucial to this study. Moon et al. utilized an eight input, 17 hidden, and single output layer ANN to model a building for predictive control applications based on the principles introduced by Yang et al. (Yang, Rivard et al. 2005, Moon and Kim 2010). The previously justified model form presented by Moon et al. appears to be the most practical adaptation for feasibility modelling at the zone level and was transparent in describing the ANN structure. Artificial neural network models have been able to successfully estimate sub-hourly (15 minute interval) building-level electrical loads, albeit with an ambiguously documented model form (Chae, Horesh et al. 2016). Additional challenges are present when applying the documented approaches at the zone level due to the unpredictability of individual occupants (Chae, Horesh et al. 2016).

Codes and standards have a significant impact on which sensors are common in commercial buildings. New additions to periodical updates of these standards can cause

sensing density (i.e., the number of sensors per unit area of a building) to correlate with building vintage. The 2017 Canadian National Energy Code for Buildings mandates vacancy sensors on automatic control systems for occupancy detection (NRCan 2017). Recent iterations of ASHRAE Standard 90.1 (2013 and onward) require the use of illuminance and occupancy sensing, and multi-level lighting control (ASHRAE 2016). ASHRAE Standard 90.1 also includes provisions for demand control ventilation, which require zone-level carbon dioxide sensor feedback to the BMS. ASHRAE Guideline 14 was written for application to building-level annual energy modelling but provides metrics for energy use model evaluation that can be adapted for other applications (ASHRAE 2002). For aggregate energy use estimation, ASHRAE Guideline 14 recommends a threshold of 10% maximum mean bias error (MBE) for adequate model accuracy (ASHRAE 2002).

Despite the effort placed on individual sensor types and applications previously described, there is a gap in determining which sensors are most valuable for estimating energy use at the zone level. This paper attempts to provide some context to the problem with an exploratory proxy sensor study using ANN inverse models to determine which sensors are most useful for estimating zone-level energy use. The sensor proxy evaluation approach required the development of a novel scoring system to rank sensors relative to each other, which is described with the general approach in the following sections.

### **4.3 Methodology**

The following section describes the sensors present in the building studied, the approach to room-level energy use quantification, an approach to generate proxy models of each

room-level energy end use, the proposed structure of an energy use estimation sensor hierarchy based on the modelling results, and the respective uncertainty analysis of the sensors.

#### **4.3.1 Sensor Introduction**

Data was collected in a new institutional building in Ottawa, Canada from 27 rooms within 11 thermal zones between September 29<sup>th</sup>, 2018 and January 29<sup>th</sup>, 2019. Most of the instrumented rooms (25) are single occupant private offices, while the other two rooms are kitchenettes. Each room has an operable awning window that is 1 m wide by 0.6 m high with a maximum opening angle of 30 degrees. Each office is the same size, shape, and has the same diffuser location and orientation. The kitchenettes have different floor areas (one smaller and one larger than a typical office in this building). Ducts from the variable air volume (VAV) boxes to the diffusers would vary in geometry according to the number of rooms in the zone (e.g., one to four branches). The primary façade for the rooms was oriented 30 degrees North of East and is shown in Figure 1-2. Four of the rooms are located on corners of the building (including the two kitchenettes). Two offices per floor have larger window to wall ratios than the others (0.8 compared to 0.4). In addition, the kitchenettes both have very high window to wall ratios (>0.9) and have Southeast-facing primary facades.

Indoor air temperature, room-level return air temperature, and outdoor air temperature were monitored with BMS integrated 10 kOhm thermistors on the thermostats and at the air intake of one of the floor-level AHUs. The thermistors were calibrated to standard conditioned room air temperatures (between 18°C and 25°C). Rooms were

grouped into zones of one to four rooms that were each supplied cooling and fresh air by one VAV unit. It was assumed the VAV airflow was divided equally by the number of rooms in the zone to estimate room airflow, corresponding with the mechanical drawings and supported by airflow spot measurements with a capture hood when the windows were closed. One zone per floor did not exclusively serve the measured rooms, and the airflow to those rooms was reduced to account for the corridor service. A complete list of the sensors utilized in this study, variable names (for classification learning), and instrumentation error is provided in Table 4-1. Common sensors as indicated in Table 4-1 are adapted from the literature (Schachinger and Kastner 2016). Instrumentation error in Table 4-1 refers to the manufacturer-supplied values associated with random error. The treatment of random error in this study aligns with Cunningham et al. (Cunningham 1981). A diagram of the sensor location and air system ductwork relative to the room geometry is provided in Figure 1-2. The current transducers (one per room for plug loads and lighting) and lighting output feedback sensors were omitted from Figure 1-2 to improve readability. Lighting feedback was provided by three illuminance sensors mounted on the ceiling, dividing the ceiling into quarters with respect to the distance from the exterior wall ( $\frac{1}{4}$ ,  $\frac{1}{2}$ , and  $\frac{3}{4}$  depths). Solar gains were not directly included in this study due to the lack of a BMS-connected weather station at the building site, but ceiling illuminance sensors have been found to be an adequate proxy at the room level in the literature (Gunay, O'Brien et al. 2016, Bursill, O'Brien et al. 2019). All the data utilized in this study was collected and stored in an industrial computer networked to the BMS. A typical location for the in-stream return air temperature sensor access is shown in Figure 4-1.

**Table 4-1: Sensor schedule with simulation abbreviation, sensor type, and instrumentation error.  
Common sensors in commercial buildings are indicated in italics.**

Measured Quantity	Short form	Sensor type	Instrumentation error
Blind position	blind	Controller output feedback	
Hour of day	hour1	Internal clock	
Illuminance near window	lux1	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance in center of room	lux2	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance near wall	lux3	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
<i>Indoor air temperature</i>	IAT	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
<i>Lighting dim level</i>	Ldim	Controller output feedback	
Lighting power	LP	Current transducer	$\pm 2\% \text{ full scale (4.8 W)}$
<i>Lighting state</i>	LS	Controller output feedback	
<i>Occupancy state</i>	occ	Passive infrared	
Outdoor air temperature	OAT	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Plug load	plug	Current transducer	$\pm 0.5\% \text{ of value}$
Radiant heater flow rate	rad_flow	Single jet flow sensor	$\pm 5\% \text{ full scale (30 l/h)}$
Radiant heater inlet temperature	radt1	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Radiant heater outlet temperature	radt2	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
<i>Radiant heater valve position</i>	rad	Controller output feedback	
Reheat coil flow rate	rhc_flow	Single jet flow sensor	$\pm 5\% \text{ full scale (30 l/h)}$
Reheat coil inlet temperature	rhct1	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Reheat coil outlet temperature	rhct2	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Return air temperature	RAT	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
<i>VAV airflow</i>	airflow	Differential pressure transducer	$\pm 1\% \text{ full scale (5 Pa)}$
<i>VAV supply air temperature</i>	sat	4-20 mA Resistance temperature detector	$\pm 0.2^\circ\text{C}$
Window state	win	Wireless dry contact	



**Figure 4-1: Return air temperature sensor access location marked with a red box (view from below). Return air is drawn back to the AHU via the duct at the grated opening in the bottom-left of the figure.**

Supply airflow was driven by a supply fan at the floor-level AHU where the air was initially conditioned to a temperature that is low enough to satisfy the warmest zone in the cooling season, and is low enough to satisfy the lowest zone setpoint temperature in the heating season. There is one AHU per floor. When cooling is required, airflow is modulated to the zones via the VAV zone-level airflow damper to allow the space to achieve a setpoint as defined by the occupant. The setpoint can be adjusted by the occupants of each room via the thermostat and can vary from 20°C to 25°C. Most rooms maintained a setpoint between 21°C and 23.0°C (as a result of user preference), where VAV zone-level supply air temperature setpoints were modulated with the reheat coil to achieve the average of the served room setpoints. If setpoints were outside of the 21°C and 23.0°C range they were reset every Sunday to 21°C in the heating season and 23°C in the cooling season. However, the VAV units are on a schedule to be active from 7:30 AM to 7:00 PM with no airflow

when inactive. Due to the large amount of exhaust air from the lab spaces the outdoor air fraction was greater than 0.9 for most of the data acquisition period to maintain building pressurization.

Data were sampled at 5-minute intervals for analog inputs (e.g., temperature) and upon state change for binary inputs (e.g., window opening). Analog BAS readings were verified within  $\pm 10\%$  via independent spot measurements. All offices are equipped with motorized shades with solar and visible transmittances of 3% and 4% respectively. Motorized blinds were automatically closed after 5:00 PM and reopened at 7:30 AM on weekdays and could be manually overridden by the occupant.

A preliminary mapping of sensors to expected conditions is presented in Figure 4-2, where the available sensors in this study are shown at the top. The bottom half of Figure 4-2 provides a list of expected conditions and indicates the possible physical sensors that could be used to derive estimates for these quantities. The mapping for Figure 4-2 was derived manually (knowledge-based) to remove scenarios that do not make physical sense. One of the objectives of this paper is to provide a systematic approach (shown later) that can be compared to the expected condition mappings. The number of instances that each physical sensor appears in the mapping in Figure 4-2 is shown in Figure 4-3. Indoor air temperature and lighting power are expected to be used for the most derivations.

The relative importance of the zone sensors in Figure 4-2 are quantified in the remainder of the paper. To achieve this, only inferred conditions related to energy use were quantitatively investigated. Other inferred conditions are acknowledged to provide value to building owners and operators but diverged from the focus of this study: direct energy

use. Qualitative analysis is contrary to the goal of this paper: providing a hierarchy of sensors for energy use estimation at the zone level.

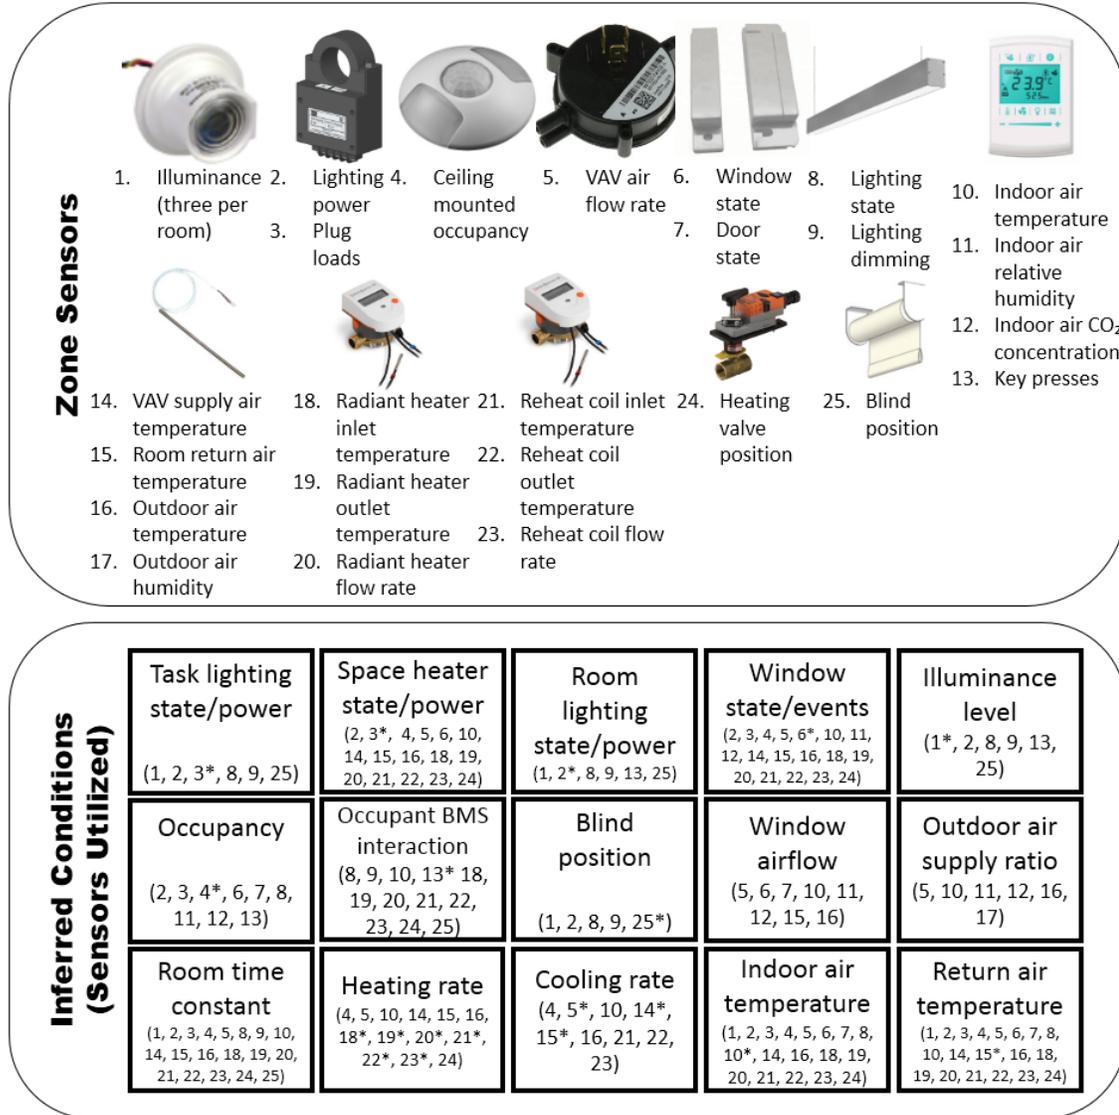
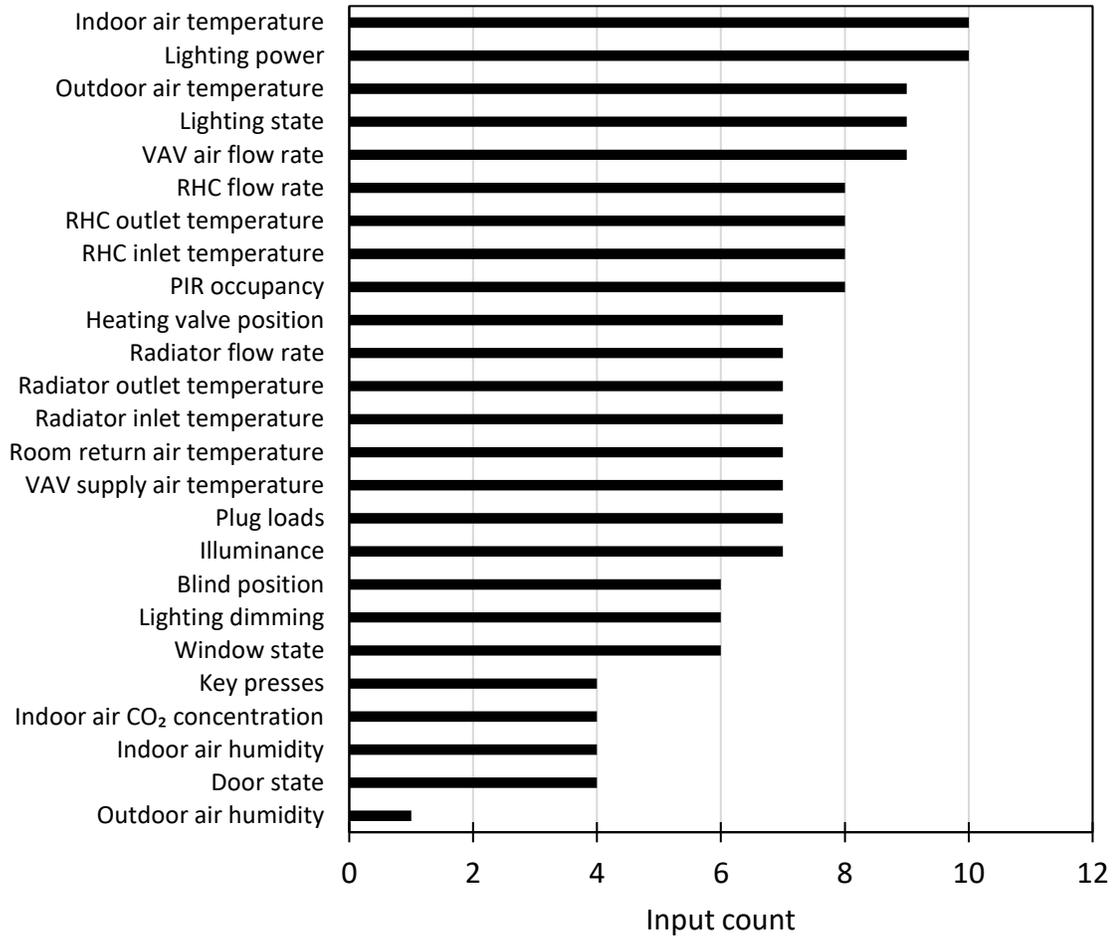


Figure 4-2: Sensors, their inferred conditions, and the sensors that can be utilized to infer such conditions (“\*” indicates sensors intended for direct measurement)



**Figure 4-3: Input count per zone sensor from the inferred conditions of Figure 4-2, where each appearance of a zone sensor in an inferred condition increments the input count of that zone sensor by one (e.g., indoor air temperature appeared in 10 inferred conditions).**

### 4.3.2 Energy Quantification

Steady-state conditions for heat transfer were assumed for the air side VAV and hydronic heating systems for each 5-minute interval. This is supported by the literature, where a zone-level time constant of 15 minutes and an average run time of two hours was found for terminal HVAC equipment (Bursill, O'Brien et al. 2018). For the purpose of

energy use quantification, the steady-state energy use at each sample is considered to be an average over the sample interval as shown in Equation 4-1 where  $E$  is the energy use at each time step,  $Q$  is the average power at each time step, and  $t$  is the index of the time step. Equation 4-1 was applied to both electrical data and thermal (air side and hydronic) data. Electrical demand data was obtained from the BMS; thus Equation 4-1 could be applied directly. A basic steady-state heat transfer equation was utilized to quantify  $Q_h$ , the heating power of the radiant heaters, reheat coils, and AHU coils. This is shown in Equation 4-2 where  $\dot{V}_h$  is the hydronic fluid flow rate,  $\rho_w$  is the density of water at 40°C,  $c_{pw}$  is the specific heat of water at 40°C,  $T_{1w}$  is the water temperature entering the heater/coil, and  $T_{2w}$  is the water temperature leaving the heater/coil (Incropera, Lavine et al. 2007). For room-level energy use estimation reheat heating coil power had to be divided over the number of rooms in the zone (it was assumed that the airflow and reheat coil heat were divided evenly over the rooms served by the VAV units in the zones). Radiant heaters were already measured at the room level, and did not require additional calculations to find their energy use at the room level. Energy added to the air from the floor-level AHU was calculated using Equation 4-3, where  $Q_{SA}$  is the supply air heating energy,  $\dot{V}_r$  is the estimated room-level air flow rate (divided evenly over the number of spaces in the zone served from the total VAV airflow rate), and  $\dot{V}_{AHU}$  is the AHU air flow rate. Energy use for hydronic equipment (radiant heaters, reheat coils, and AHU coils) at the room level was quantified by applying Equation 4-1 to the results of Equations 4-2 and 4-3.

Equations 4-1 to 4-3 allowed for the quantification of the following energy flows at the room level:

- Lighting (Equation 4-1)
- Plug loads (Equation 4-1)
- Hydronic radiant heating (Equation 4-2 and Equation 4-1)
- Hydronic VAV reheating (Equation 4-2 and Equation 4-1)
- Supplied energy from the AHU to condition return and outdoor air (Equation 4-3 and Equation 4-1)

$$E(t) = Q(t)\Delta t \quad (4-1)$$

$$Q_h(t) = \dot{V}_h(t)\rho_w c_{pw}(T_{1w}(t) - T_{2w}(t)) \quad (4-2)$$

$$Q_{SA}(t) = \frac{\dot{V}_r(t)}{\dot{V}_{AHU}(t)} \left( \dot{V}_h(t)\rho_w c_{pw}(T_{1w}(t) - T_{2w}(t)) \right) \quad (4-3)$$

### 4.3.3 Uncertainty Analysis

Instrumentation uncertainty had a significant impact on the reliability of the energy estimation results. The compounding effect of different sensor uncertainty values can generate measurement uncertainty up that approaches 100% of the measured value in some cases (Wang, Duan et al. 2016). The approach to uncertainty analysis and different sources of error in this paper are described as follows.

The thermistors had respective instrumentation uncertainty of 0.2°C. Systematic error due to thermistor positioning at the wall and not the center of the room was not quantified but is acknowledged to induce additional error. The differential pressure transducers measuring the airflow to the zones (via a power relation) had an instrumentation uncertainty of 1% of full scale (2.5 Pa), equating to an average of 30 L/s per room. Individual room maximum airflow rates ranged from 61 L/s to 140 L/s. The assumption of equal division of zone airflow to each room is acknowledged to induce additional error with respect to airflow. The aggregate instrumentation uncertainty (i.e., the uncertainty of the compound measurement due to the combined instrumentation uncertainty of all input quantities) of air side sensible cooling and heating energy estimates was found to be 31.6% of the average value using a separation of variables approach described by Equation 4-4, where  $s$  is the aggregate uncertainty,  $Q$  is the energy quantity being assessed for uncertainty,  $x$  is the input variable,  $i$  is the index of the input variable, and  $n$  is the number of input variables required to derive the energy quantity (Smith 2013). The aggregate precision error (i.e., error due to the smallest level of precision of the measuring equipment) of the air side sensible cooling and heating energy estimates was found to be 7.0% of the average value of air side sensible cooling and heating energy. Hydronic side aggregate instrumentation uncertainty of radiant heating energy estimates was found to be 13.7% of the average value. Precision error in the hydronic case was negligible (less than 0.1%). Other individual instrumentation uncertainties are provided in Table 4-1. Analog to digital conversion values of 4096 were standard for all inputs on the BMS. The largest source of measurement uncertainty was from the air side pressure

transducer, where uncertainties of 1% full scale were observed. Pressure transducer accuracy is often limiting in in-situ research and improving the economics of higher accuracy embedded zone-level air measurements could make future work more conclusive. Despite the high measurement uncertainties calculated from manufacturer data, spot measurements from using a capture hood on room-level diffusers found an average airflow measurement uncertainty of 4.0% corresponding to an aggregate heating energy estimate uncertainty of 9.0% which was assumed to be acceptable for this study.

$$\Delta s_Q^2 = \sum_{i=1}^n \left( \frac{\partial Q}{\partial x} \right)^2 \Delta x_i^2 \quad (4-4)$$

#### 4.3.4 Sensor Proxy Approach

Sensor proxies were developed using feed forward artificial neural network (ANN) models and combinations of inputs from the test facility with the MATLAB ANN Toolbox to find a near best case model for each input and end use combination (Hagan, Demuth et al. 1996). Cross-correlation was not found to be sufficient for the analysis because many input-output relationships were highly nonlinear. More detailed models would be more useful to practitioners, but at this point the sample space of input and output combinations is large and unexplored. The data collected over four months was split into training and validation sets, where the training set was also sub-divided for feed forward ANN training and validation.

Firstly, each input was assessed individually for its ability to predict each given energy end use output based on MBE over the validation period. Each input and output

combination was considered a “case”, where 85 case combinations utilizing single inputs and single outputs (SISO) were proposed for the five end uses previously introduced (Lighting, Plug loads, Hydronic radiant heating, Hydronic VAV reheating, and Supplied energy from the AHU to condition return and outdoor air). The energy use estimation results of the 27 rooms were averaged for each case when scored.

Next, when the best model for end use SISO cases achieved MBE greater than 5% (e.g., OAT as an input for AHU cooling power estimation) forward selection was utilized to add inputs to the models systematically (based on incrementing SISO case MBE) to reduce the MBE and generate multiple input single output (MISO) models. Forward selection for ANN inputs involves adding inputs systematically to improve the model accuracy (May, Dandy et al. 2011). In this study, inputs were added to the MISO models based on SISO MBE. Models became increasingly complex in MISO cases by adding inputs with forward selection, and the incremental improvement in model accuracy from an additional input was used to add to the score of individual inputs. The 5% MBE threshold was selected from ASHRAE Guideline 14, where 10% MBE corresponds to a minimally acceptable model and 5% MBE is twice as accurate, indicating incremental accuracy may not be worth the additional data acquisition and modelling efforts for the application in this study (ASHRAE 2002). Finally, the resulting SISO and MISO cases of energy use estimation accuracy (incremental accuracy in the MISO cases) were summarized and ranked by end use and input for comparison. The end use rankings (i.e., how accurate the illuminance-based model was at estimating lighting loads) are assumed to be more widely applicable and less building specific than the input rankings (i.e., how

much energy illuminance sensors can accurately estimate).

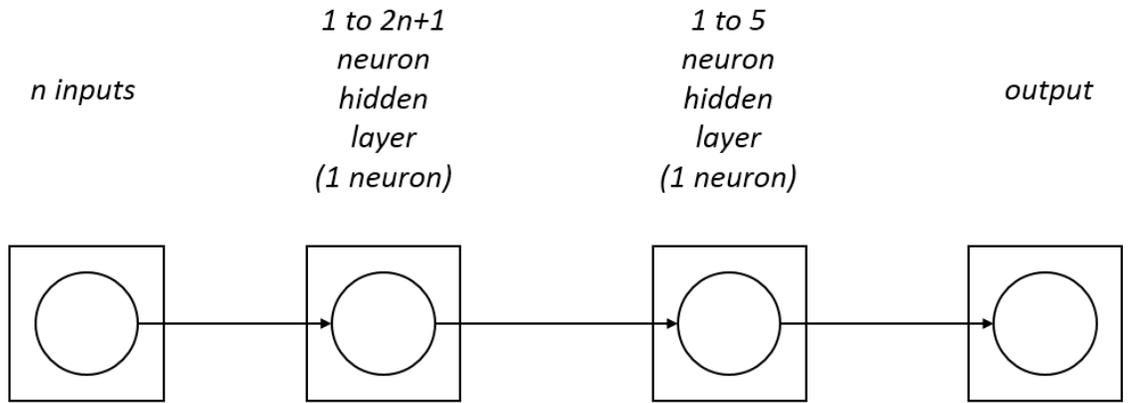
Proxy models were generated based on the mapping from Figure 4-2. The requirement of ASHRAE Guideline 14 for time-step CV(RMSE) to be less than 30% was not considered in this study because this modelling exercise was focused on estimating energy use over a relatively long period. Additionally, the model and Guideline 14 time-steps (5 minutes vs. 60 minutes) were not as comparable as the estimation horizons (2 months vs. 12 months). The MBE utilized to select models in this study was calculated as a percentage as shown in Equation 4-5, where  $M$  is the measured data,  $S$  is the simulated data,  $i$  is the instance, and  $N$  is the number of samples.

$$MBE = \frac{\sum_{i=1}^{N_i} (M_i - S_i)}{\sum_{i=1}^{N_i} M_i} \quad (4-5)$$

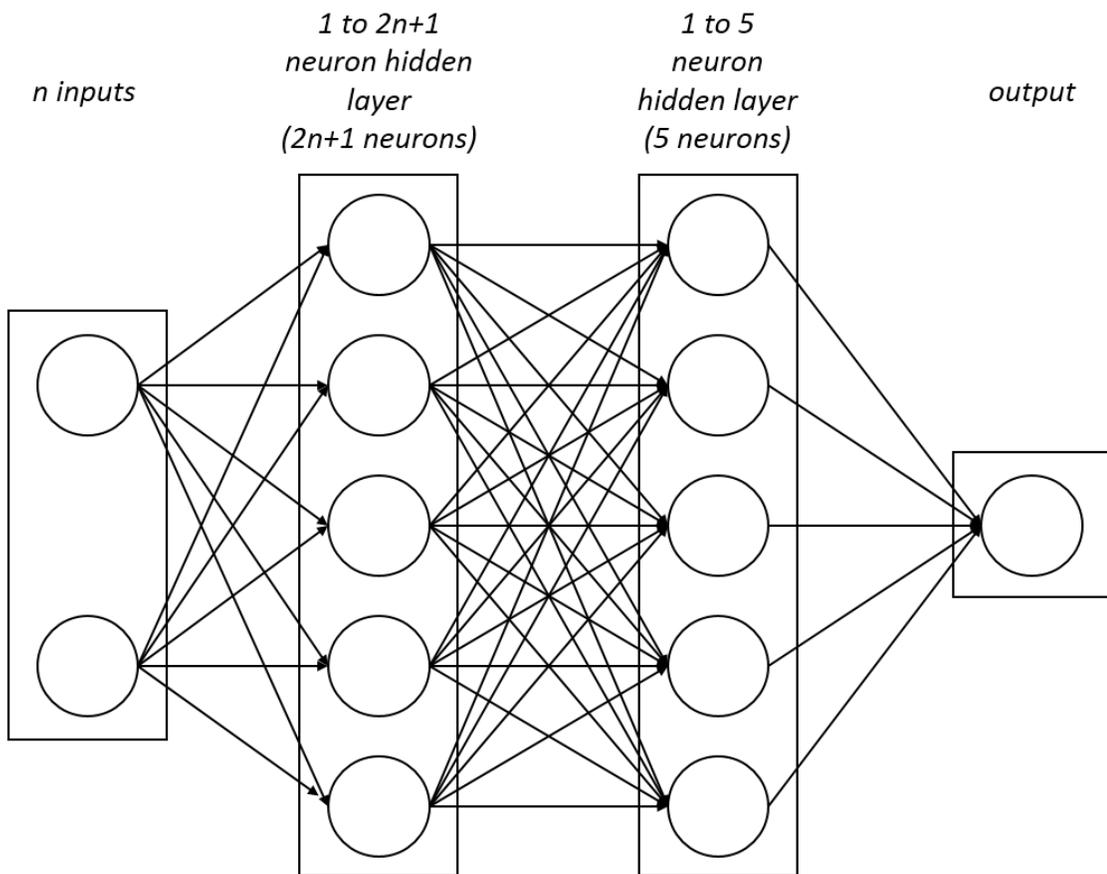
One input, hidden, and output layer is defined in the literature as being sufficient for most building modelling applications (Hecht-Nielsen 1992). The input layer should have a number of neurons equal to the number of inputs to the model,  $n$ . Under this premise, if there are  $n$  inputs and one output, it is suggested that the hidden layer consist of  $2n + 1$  hidden neurons (Hecht-Nielsen 1992). To provide additional model flexibility in this study the first hidden layer was permitted to have between one and  $2n + 1$  hidden neurons, and a second hidden layer was added that could have between one and five neurons. Model flexibility was added due to early modelling iterations where model instability was found due to inappropriate model form and increased the required simulation time exponentially. Regardless of the suggestions for model form in the

literature, the presence of redundant neurons should not pose a significant problem and a good learning algorithm tends to ignore the excessive parameters in the system (Yang, Rivard et al. 2005). Logistic sigmoid functions were utilized as the activation functions for the neurons. Gradient descent is the preferred training method in building modelling applications for large data sets, but Levenberg–Marquardt is more robust for smaller data sets or if the computation time is not large (Yang, Rivard et al. 2005). Due to the data size and computational power available, the Levenberg–Marquardt training method was chosen.

A script was developed for this study using forward selection with the MATLAB ANN toolbox. Model input and output data were provided for the SISO and MISO models from a master database that contained eight months of five-minute time-step data (four months each of heating and cooling season data) for all of the sensors used in this study. The seasonal data was divided in half (half for training and half for validation/estimation) for each season during modelling. In the script, models were generated for all  $10n + 5$  ANN structures for each case (the product of the up to  $2n + 1$  neuron first layer and the up to five neuron second layer), and the model with the lowest MBE over the validation data set was selected for each SISO and MISO case. MISO cases had much longer simulation times than the SISO cases because the first ANN layer increased at a rate of  $2n + 1$  for every input,  $n$ . Illustrations of two sample ANN structures from this study is shown in Figure 4-4 and Figure 4-5.



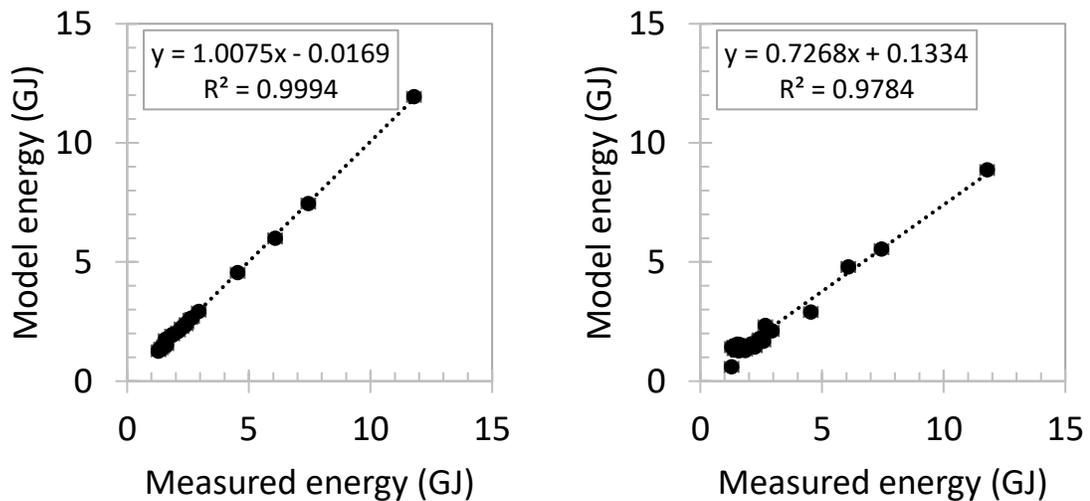
**Figure 4-4: ANN structure of proxy sensor selection script for a SISO case with one input on the 1x1 hidden neuron iteration (the minimum complexity iteration for the SISO case)**



**Figure 4-5: ANN structure of proxy sensor selection script for a MISO case with two inputs on the 5x5 hidden neuron iteration (the maximum complexity iteration for the two input single output case)**

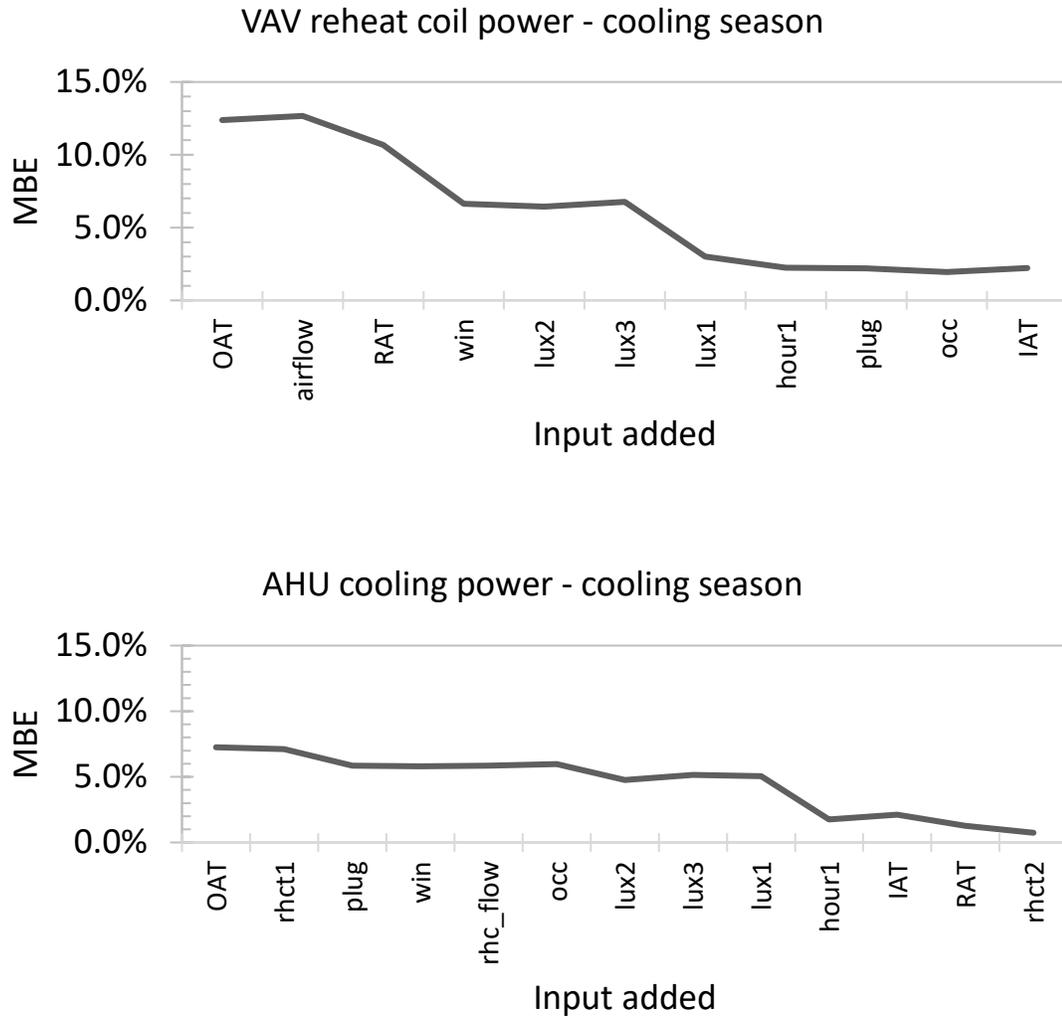
MBE was chosen over mean absolute error (MAE) of all samples in this case because the goal of the modelling exercise was to estimate mid-term to long-term energy usage for building owners (rather than time step power which would be more appropriately assessed with MAE). However, this is one approach as a feasibility study and may not be the optimal or most accurate way to model sensors via proxy for other applications (e.g., fault detection, indoor environment quality, real-time control decisions).

Individual input cases varied significantly in average MBE. Two such cases are shown in Figure 4-6, where the aggregate energy predictions for the radiant heating energy of all 27 rooms during the validation period are plotted for the singular inputs of OAT and IAT in the heating season (0.1% and 29.9% MBE respectively). Error bars are applied to the measured energy using the values from the uncertainty analysis. Most of the error can be demonstrated by the difference in slopes of the trendline.



**Figure 4-6: Aggregate energy prediction over the two-month validation period for the validation data using individual sensor proxy models for office-level radiant heater energy using only OAT (left) and IAT (right). Each dot represents the aggregate measured and modelled energy use for one room over the validation period.**

Forward selection modelling was executed in a systematic manner, where the MBEs of the single input models were used to rank the order in which the inputs were added (May, Dandy et al. 2011). It was assumed that useful information from additional inputs would improve the model, and useless information would not significantly degrade the model due to the sample volume (13,500 samples each for training and validation data sets). Forward selection modelling was executed for 24 MISO cases in the cooling season and 38 MISO cases in the heating season using the available sensor data as shown in Figure 4-7 and Figure 4-8. The MISO models start with only the left-most input and become additively complex by adding the input to the right (i.e., the OAT, airflow, RAT, and window model in Figure 4-7 has a MBE that corresponds with the window variable on the x-axis). In Figure 4-8 the MBE increases twice during forward selection, on each occasion when a temperature is added from the radiant heater hydronic loop (first the inlet, then the outlet temperature). This likely indicative of a time lag in temperature between the AHU heating coil water and the radiant heater water. Only current time-step sensor values were used in this study as a result of preliminary analysis suggesting time lag did not have a significant impact for energy use estimates using this proxy approach. Adding previous time-step data as inputs to the ANN models is something that is suggested for future work to mitigate this effect when combining same time-step inputs that exhibit thermal lag. The tabulated results of the individual models and forward selection models are presented in the Results and Discussion section.



**Figure 4-7: Forward model MBE progression of ANNs for VAV reheat coil power (top) and AHU cooling power (bottom) in the cooling season. The input labels on the x-axis are provided in Table 4-1.**

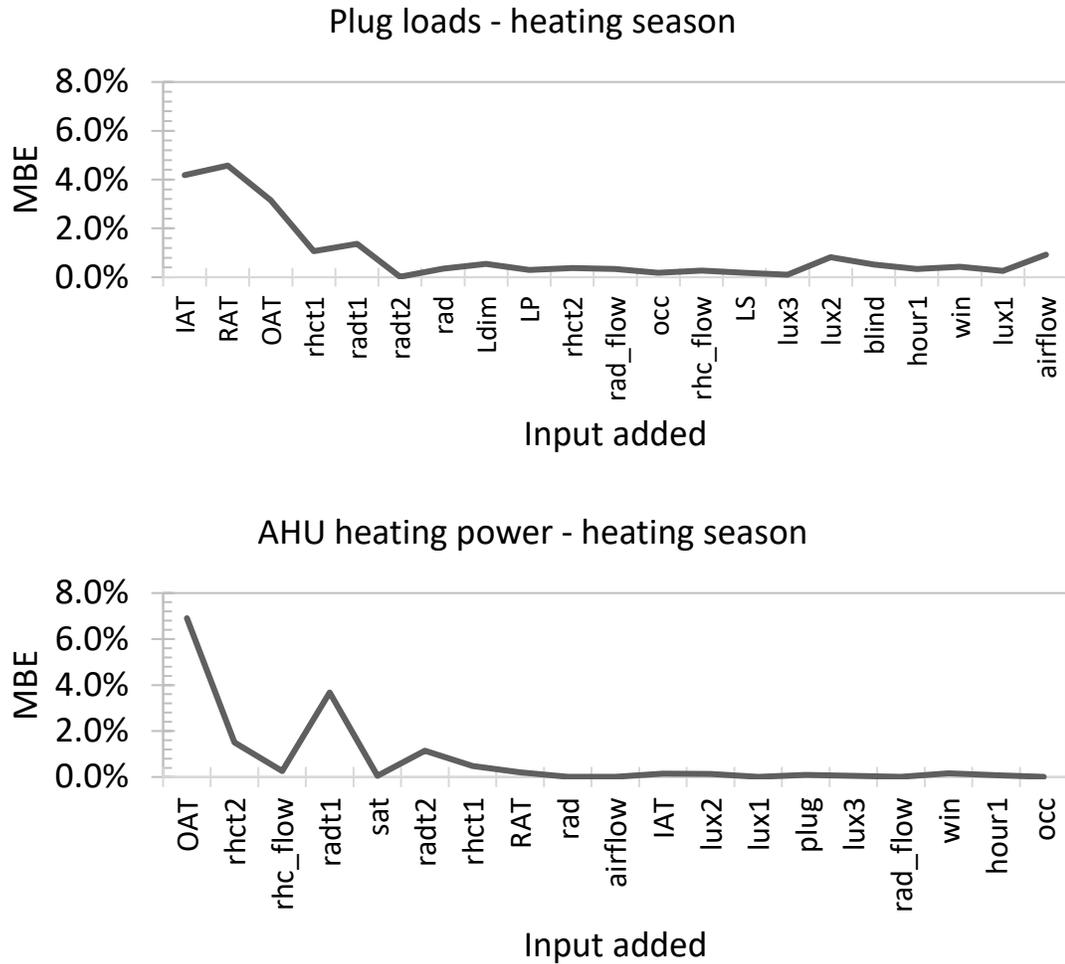


Figure 4-8: Forward model MBE progression of ANNs for plug loads (top) and AHU heating power (bottom) in the heating season. The input labels on the x-axis are provided in Table 4-1.

### 4.3.5 Ranking Sensors

The energy estimated by the SISO and MISO proxy models was used to rank the sensors for each energy end use and cumulatively for all end uses. SISO sensor models for end uses were given the energy accurately predicted as per Equation 4-6, where  $IND$  is the individual input score and  $E(t)$  is the measured energy use found for either electrical

(Equation 4-1) or thermal energy (Equation 4-2 or Equation 4-3 and Equation 4-1). Forward selection MISO energy use prediction scores were based on the incremental benefit of adding the input, where the incremental benefit of the first input added was split between the first two inputs (and thereafter each singular input added would get the entire incremental score). The incremental score of forward selected inputs was calculated as shown in Equation 4-7, where  $FWD$  is the forward modelling input score, and  $n1$  and  $n2$  represent the MBE before and after the new input has been added to the model respectively. The  $IND$  and  $FWD$  energy use prediction scores are added to produce the total score for each sensor. By design, an individual input model will have a much higher score than an incremental input to a forward selection model. Additionally, the value of each sensor is also representative of the quantity of energy it can estimate in the building studied (while meeting acceptable model criteria). Some sensors appeared in multiple end use models, and thus their energy use prediction scores are additive and higher to acknowledge that they can be used to estimate multiple energy end uses.

$$IND = (1 - MBE) \sum_{i=1}^t E(t) \quad (4-6)$$

$$FWD = (MBE_{n2} - MBE_{n1}) \sum_{i=1}^t E(t) \quad (4-7)$$

#### 4.4 Results and Discussion

The results of the energy use quantification, sensor proxy modelling, and sensor hierarchy are discussed in the following sections. Additionally, notable anomalies such as space

heater usage and the difference between indoor air temperature measured at the thermostat and return air temperature in the duct are also discussed anecdotally.

#### **4.4.1 Energy End Uses**

Room-level energy use at the office level for the entire four months of data was quantified using the BAS integrated sensors and the equations introduced in the Methodology section. A breakdown of the occupancy fraction (ratio of the occupied period to total measurement time), design lighting power density (from the electrical drawings), floor area, energy end usage, and total metered room energy usage is provided in Table 4-2. Asterisks (“\*”) in Table 4-2 indicate missing values that were replaced by the average where sensor errors occurred and double asterisks (“\*\*”) indicate rooms with high window to wall ratios. Air side VAV zoning is also provided to explain some of the duplicate values from VAV energy division. One of the goals of this study was to determine the largest energy end uses for the monitored offices and which sensors the most useful for estimating energy use. Notably, the end use breakdown in this study will not be applicable to all buildings. As such, this data is provided to the reader to assess the applicability of the results of this study to individual buildings and end uses.

The lighting power density was based on electrical drawing fixtures and did not include ballast factors. Occupancy ratios were significantly lower than the average 40-hour work week (which corresponds to 0.238), which could be due to a combination of including the winter break and the nature of professor duties. Heating energy in three different modes (AHU preheating, VAV reheating, and room-level radiators) comprised most of the heating

season energy usage. Electrical end uses were an order of magnitude lower than the sum of hydronic (radiant panel, reheat coil, and AHU coil) heating end uses. The energy end usage and occupancy fractions of Table 4-2 are represented graphically in Figure 4-9. Rooms 26 and 27 were both kitchenette/meeting rooms that had highly glazed ( $>0.9$  window to wall ratio) south facing walls. Additionally, rooms 21 and 25 were corner offices that had window to wall ratios of roughly 0.8. Room 21 would have had higher energy use, but the absence of radiant heater data was replaced with the average (which likely is lower than if data was present).

**Table 4-2: Energy use breakdown for measured rooms for the cooling and heating seasons. “\*” indicates missing data replaced by the average value, “\*\*” indicates rooms with high window to wall ratios**

Room	Zone	Lighting power density (W/m <sup>2</sup> )	Area (m <sup>2</sup> )	Cooling season									Heating season								
				Occupancy fraction	Radiant heater energy (GJ)	Reheat coil energy (GJ)	Supply air energy (GJ)	Plug loads energy (GJ)	Lighting energy (GJ)	Total (GJ)	Area normalized total (GJ/m <sup>2</sup> )	Energy use Rank	Occupancy fraction	Radiant heater energy (GJ)	Reheat coil energy (GJ)	Supply air energy (GJ)	Plug loads energy (GJ)	Lighting energy (GJ)	Total (GJ)	Area normalized total (GJ/m <sup>2</sup> )	Energy use Rank
1	10	3.4	14.7	0.026	0.00	0.02	7.17	0.15	0.02	7.36	0.50	19	0.019	2.26	1.74	2.83	0.48	0.04	7.35	0.50	1
2	10	3.4	14.7	0.025	0.00	0.02	7.17	0.22	0.02	7.43	0.50	21	0.007	3.07	1.74	2.83	0.07	0.03	7.75	0.53	2
3	3	3.4	14.7	0.037	0.00	0.06	4.23	0.18	0.02*	4.50	0.31	7	0.044	3.14	1.56	2.61	0.45	0.05*	7.77	0.53	3
4	2	3.4	14.7	0.001	0.00	0.00	5.85	0.16	0.02	6.03	0.41	15	0.020	2.17	2.18	3.16	0.21	0.04	7.80	0.53	4
5	3	3.4	14.7	0.043	0.00	0.06	4.23	0.03	0.01	4.34	0.29	3	0.064	3.65	1.56	2.61	0.10	0.04	7.97	0.54	5
6	10	3.4	14.7	0.072	0.00	0.02	7.17	0.20	0.02	7.40	0.50	20	0.023	2.74	1.74	2.83	0.69	0.04	8.05	0.55	6
7	7	3.4	14.7	0.008	0.00	0.01	4.80	0.01	0.02	4.84	0.33	8	0.031	2.37	2.58	3.04	0.02	0.07	8.06	0.55	7
8	4	3.4	14.7	0.024	0.00	0.01	5.66	0.01	0.02	5.70	0.39	10	0.020	4.18	1.19	2.80	0.01	0.05	8.24	0.56	8
9	7	3.4	14.7	0.021	0.00	0.01	4.80	0.05	0.02	4.88	0.33	9	0.053	2.45	2.58	3.04	0.24	0.05	8.36	0.57	9
10	4	3.4	14.7	0.030	0.00	0.01	5.66	0.08	0.01	5.76	0.39	12	0.040	4.24	1.19	2.80	0.16	0.05	8.44	0.57	10
11	3	3.4	14.7	0.018	0.00	0.06	4.23	0.02	0.01	4.33	0.29	2	0.041	4.18	1.56	2.61	0.06	0.05	8.47	0.58	11
12	9	3.4	14.7	0.055	0.00	0.01	5.68	0.50	0.03	6.22	0.42	17	0.037	2.28	1.85	3.69	0.69	0.05	8.57	0.58	12
13	8	3.4	14.7	0.004	0.00	0.00	3.91	0.11	0.02	4.04	0.27	1	0.006	2.68	1.37	4.37	0.17	0.02	8.62	0.59	13
14	9	3.4	14.7	0.083	0.00	0.01	5.68	0.49	0.02	6.19	0.42	16	0.046	2.16	1.85	3.69	1.16	0.05	8.92	0.61	14
15	3	3.4	14.7	0.020	0.00	0.06	4.23	0.09	0.02	4.41	0.30	4	0.042*	4.54	1.56	2.61	0.18	0.05	8.94	0.61	15
16	4	3.4	14.7	0.006	0.00	0.01	5.66	0.03	0.02	5.72	0.39	11	0.013	4.97	1.19	2.80	0.01	0.03	9.00	0.61	16
17	2	3.4	14.7	0.019	0.00	0.00	5.85	0.03	0.02	5.90	0.40	13	0.039	3.55	2.18	3.16	0.21	0.04	9.12	0.62	17
18	2	3.4	14.7	0.060	0.00	0.00	5.85	0.08	0.03	5.96	0.41	14	0.020	3.65	2.18	3.16	0.09	0.04	9.14	0.62	18
19	9	3.4	14.7	0.081	0.00	0.01	5.68	0.58	0.02	6.28	0.43	18	0.052	2.47	1.85	3.69	1.25	0.05	9.32	0.63	19
20	8	3.4	14.7	0.024	0.00	0.00	3.91	0.50	0.02	4.44	0.30	5	0.029	2.57	1.37	4.37	1.37	0.05	9.73	0.66	20
21**	6	3.4	14.7	0.052	0.00	0.01	11.71	0.17	0.02	11.91	0.81	25	0.099	4.80*	2.11	3.86	0.40	0.05	11.21	0.76	21
22	8	3.4	14.7	0.043	0.00	0.00	3.91	0.52	0.02	4.45	0.30	6	0.097	4.80*	1.37	4.37	1.02	0.16	11.72	0.80	22
23**	6	3.4	14.7	0.040	0.00	0.01	11.71	0.13	0.02	11.87	0.81	24	0.065	7.66	2.11	3.86	0.32	0.04	13.99	0.95	23
24**	1	3.4	14.7	0.017	0.00	0.05*	9.31	0.13	0.02	9.51	0.65	23	0.031	10.87	2.18*	2.48	0.32	0.05	15.90	1.08	24
25**	1	3.4	14.7	0.021	0.00	0.05*	9.31	0.02	0.00	9.39	0.64	22	0.086	12.88	2.18*	2.48	0.13	0.11	17.79	1.21	25
26**	11	3.9	12.7	0.060	0.00	0.00	13.28	0.08	0.01	13.37	1.05	26	0.030	4.80*	8.71	15.85	0.18	0.01	29.56	2.32	26
27**	5	8.3	29.3	0.083	0.00	0.87	30.20	0.05	0.02	31.14	1.06	27	0.071	20.44	5.20	28.69	0.20	0.11	54.64	1.87	27
Average uncertainty					0.00	0.00	1.94	0.01	0.04	1.99	0.14			0.57	0.25	1.01	0.02	0.04	1.89	0.13	
Average offices		3.40	14.72	0.03	0.00	0.02	6.14	0.18	0.02	6.36	0.43		0.04	4.17	1.80	3.19	0.39	0.05	9.61	0.65	
Standard deviation offices		0.00	0.00	0.02	0.00	0.02	2.24	0.18	0.01	2.22	0.15		0.03	2.65	0.41	0.61	0.41	0.03	2.63	0.18	
Average all		3.60	15.19	0.04	0.00	0.05	7.29	0.17	0.02	7.53	0.48		0.04	4.80	2.18	4.60	0.38	0.05	12.02	0.76	
Standard deviation all		0.95	2.84	0.02	0.00	0.17	5.24	0.18	0.01	5.35	0.22		0.03	4.03	1.51	5.43	0.40	0.03	9.68	0.43	
Total			410.00		0.00	1.39	196.89	4.62	0.49	203.39				129.56	58.91	124.33	10.21	1.42	324.44		

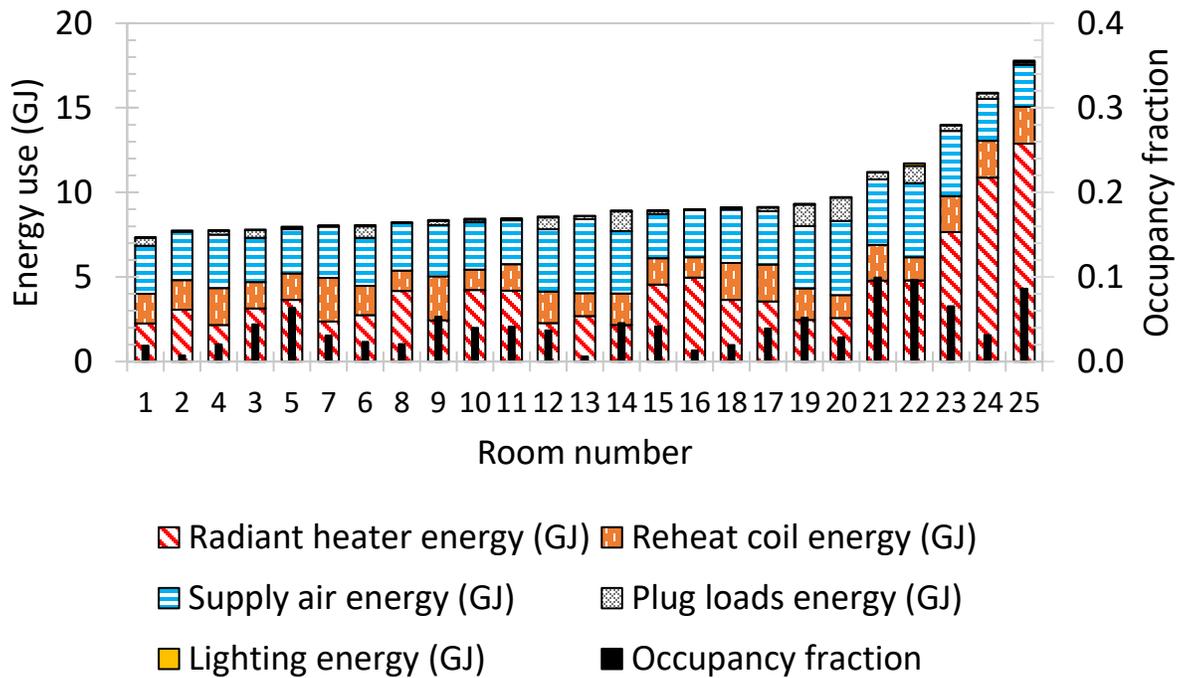
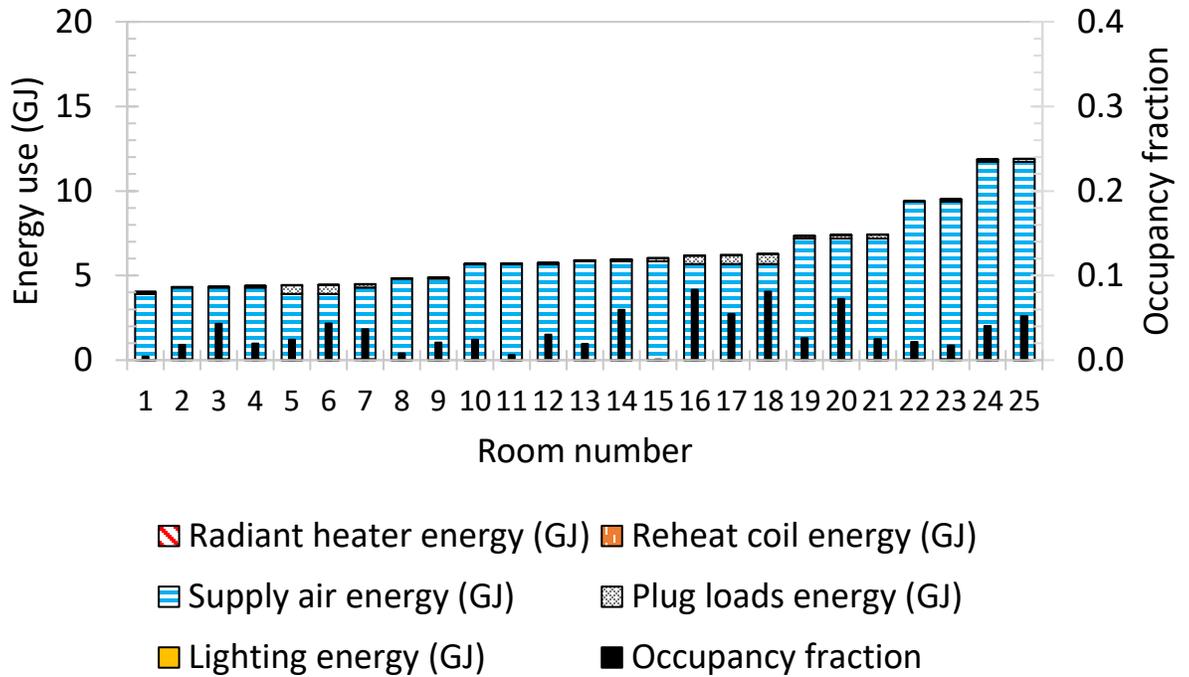


Figure 4-9: Energy use breakdown and occupancy fraction for monitored offices for the cooling season (top) and heating season (bottom). Lighting energy use is at the top and not visible for most offices.

Lighting energy usage may have been low due to the effective usage of daylighting, dimming ability, operable blinds, manual on/automatic off control, and low power density due to LED technology (Gilani and O'Brien 2018). Lighting operating in a typical office from the study for 8 hours per day would use 0.176 GJ over the data collection period, confirming low lighting energy use. This also confirms that even in the design case lighting uses significantly less energy than other end uses in the building studied. The average plug load of the measured offices equates to 2.5 W/m<sup>2</sup> over all samples (occupied and unoccupied), less than the range of values in the literature (5.4 – 32.2 W/m<sup>2</sup>) (ASHRAE 2013, Fuertes and Schiavon 2014, Gunay, O'Brien et al. 2016). Plug loads were also low in comparison to heating and cooling (but were on average four times higher than lighting loads). The low plug loads could be caused by the low occupant density (25 of the 27 room are single occupant offices).

The mean and median room-level energy breakdowns are provided in Figure 4-10 to demonstrate typical energy breakdowns within the rooms studied. In the 21 rooms of similar geometry, a range of 4.0 to 7.4 GJ from April 29<sup>th</sup> to August 29<sup>th</sup> and 7.4 to 11.7 GJ of total energy usage was measured from September 29<sup>th</sup> to January 29<sup>th</sup>, demonstrating significant variation based on occupant preferences and space utilization. Due to the dominance of heating and cooling energy usage (on average >95% of total metered energy for the rooms studied), temperature setpoints would likely have the greatest immediate and controllable operational impact on total office energy usage. Of the HVAC heating energy usage, only the radiant heating portion (41% of heating energy use) is fully controlled by the individual office occupants. The reheat coil energy (19% of heating energy use) is

influenced by individual office setpoints but is controlled to the office within the zone with the highest heating or cooling demand. This suggests that zoning has a significant role in energy usage variation of similar offices.

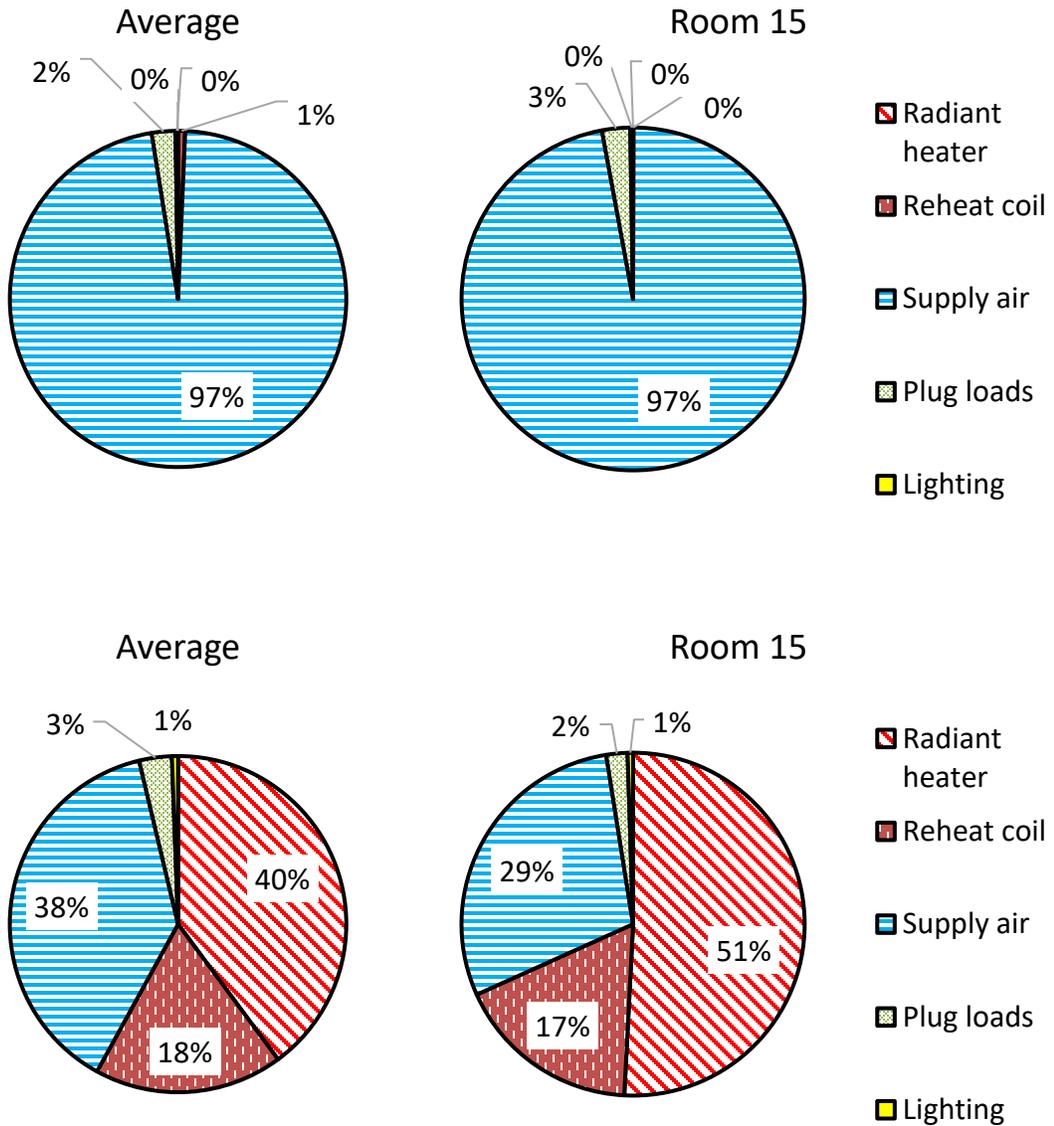


Figure 4-10: Energy usage per end use by percent for the average of the rooms (left) and the median total energy room (right) for the cooling season (top) and heating season (bottom)

#### 4.4.2 Sensor Proxy Results

The results of the approach presented in the Methodology section for the estimated energy end uses are presented in Table 4-3. The score (in GJ) is the amount of energy that is enabled to be estimated by the presence of the sensor over the two-month validation period. Only Equation 4-6 was used to calculate the score for lighting power and plug loads in the cooling season and lighting power, radiant heater, and reheat coil energy in the heating season. The sum of Equation 4-6 and Equation 4-7 were used to calculate the energy use prediction scores for reheat coil energy and supply air energy in the cooling season and plug loads and supply air energy in the heating season. Energy use prediction scores are averaged for all 27 rooms for each row of Table 4-3, and scores of zero (or very close to zero) were removed. Table 4-3 is intended to be used as a reference for each individual end use, particularly for the smaller end uses such as lighting or plug loads which may be a larger portion of the energy use in other buildings.

Notable counter-intuitive sensing scenarios include: IAT as a proxy for plug loads, illuminance near window and window state as proxies for reheat coil energy and reheat coil outlet temperature as a proxy for AHU supply air heating energy. IAT was likely a good proxy for plug loads in this case due to the presence of electric plug-in space heaters in some of the offices. Offices with higher plug loads were likely using electric personal heaters, which is investigated further in the following section. Illuminance near the window is closely related to solar radiation because it is not as affected as illuminance sensors deeper in the room by indoor shading and light diffusion conditions. Resultantly, the reheat coil energy usage is inversely related to the solar thermal gains of the rooms studied.

Additionally, window state will have a significant effect on the heating requirement above what is supplied by the AHU in the heating season. The reheat coil outlet temperature can provide information about both the supply water temperature and the amount of heat used in the zone-level coil to the ANN, which explains the usefulness for estimating AHU heating loads. The previously described notable cases are examples of how the approach taken in this paper revealed useful sensor proxy relationships that otherwise may not have been identified and quantified.

The scores in Table 4-3 can also be used to compare individual sensors to each other for energy use estimation accuracy. The MISO score additions from forward selection had a limited effect on the scores, and the SISO scores were dominant. It can generally be assumed that combining sensors with high scores in Table 4-3 will improve the model estimation accuracy (as is shown in most cases in Figure 8 and Figure 9). For example, rooms in the measured building with only outdoor air temperature and light dimming state sensors would be assumed to have better energy use estimation accuracy than similar rooms using lighting state, lighting power, and reheat coil flow rate sensors in the cooling season using the approach to comparing sensors presented in this paper. This approach assumes there would still be a requirement for temporary sensors to calibrate proxy sensor models to use the permanent sensors. Specific multiple input combinations for energy use estimation are considered beyond the scope of this study due to the lack of an existing approach for comparison of proxy sensors and were not explicitly explored. It is acknowledged that this should be an area of future work, and that a full factorial exploration of sensor combinations could be beneficial (but computationally intensive). The same

scoring methodology proposed in this paper could then be applied to compare the specific multiple input combinations for estimating various quantities.

**Table 4-3: Energy use prediction score for each input with respect to each energy end use studied for the cooling season (top) and heating season (bottom)**

Input	Target	Total score (GJ)
Blind position	Lighting power	0.035
Lighting state	Lighting power	0.034
Illuminance near window	Lighting power	0.034
Lighting dim level	Lighting power	0.034
Illuminance near wall	Lighting power	0.034
Occupancy state	Lighting power	0.034
Illuminance in center of room	Lighting power	0.034
Hour of day	Lighting power	0.034
Plug load	Lighting power	0.034
Window state	Lighting power	0.033
Lighting dim level	Plug load	0.157
Lighting power	Plug load	0.156
Reheat coil outlet temperature	Plug load	0.156
Illuminance in center of room	Plug load	0.156
Illuminance near wall	Plug load	0.155
Reheat coil flow rate	Plug load	0.154
Hour of day	Plug load	0.154
Outdoor air temperature	Plug load	0.154
Illuminance near window	Plug load	0.154
Occupancy state	Plug load	0.154
VAV airflow	Plug load	0.153
Blind position	Plug load	0.153
Lighting state	Plug load	0.152
Window state	Plug load	0.152
Indoor air temperature	Plug load	0.144
Reheat coil flow rate	Reheat coil energy	0.034
Reheat coil outlet temperature	Reheat coil energy	0.034
Outdoor air temperature	Reheat coil energy	0.034
Window state	Reheat coil energy	0.001
Illuminance near window	Reheat coil energy	0.001
Outdoor air temperature	Supply air energy	8.069
Hour of day	Supply air energy	0.145
Plug load	Supply air energy	0.055
Illuminance in center of room	Supply air energy	0.053
Return air temperature	Supply air energy	0.037
Reheat coil outlet temperature	Supply air energy	0.023
Illuminance near window	Supply air energy	0.004
Reheat coil inlet temperature	Supply air energy	0.003
Window state	Supply air energy	0.002

Input	Target	Total score (GJ)
Illuminance near window	Lighting power	0.040
Illuminance near wall	Lighting power	0.040
Lighting dim level	Lighting power	0.040
Illuminance in center of room	Lighting power	0.040
Lighting state	Lighting power	0.039
Plug load	Lighting power	0.039
Occupancy state	Lighting power	0.038
Window state	Lighting power	0.037
Hour of day	Lighting power	0.037
Blind position	Lighting power	0.037
Indoor air temperature	Plug load	0.157
Return air temperature	Plug load	0.152
Outdoor air temperature	Plug load	0.002
Radiant heater outlet temperature	Plug load	0.002
Outdoor air temperature	Radiant heater energy	2.580
Radiant heater inlet temperature	Radiant heater energy	2.515
Radiant heater flow rate	Radiant heater energy	2.490
Radiant heater outlet temperature	Radiant heater energy	2.440
Radiant heater valve position	Radiant heater energy	2.416
Illuminance near window	Reheat coil energy	0.994
Occupancy state	Reheat coil energy	0.991
Window state	Reheat coil energy	0.990
Hour of day	Reheat coil energy	0.990
Plug load	Reheat coil energy	0.985
Radiant heater valve position	Reheat coil energy	0.984
VAV supply air temperature	Reheat coil energy	0.984
Reheat coil outlet temperature	Reheat coil energy	0.974
Illuminance in center of room	Reheat coil energy	0.971
Illuminance near wall	Reheat coil energy	0.960
VAV airflow	Reheat coil energy	0.949
Radiant heater flow rate	Reheat coil energy	0.928
Outdoor air temperature	Reheat coil energy	0.926
Indoor air temperature	Reheat coil energy	0.905
Outdoor air temperature	Supply air energy	3.410
Reheat coil outlet temperature	Supply air energy	0.187
Reheat coil flow rate	Supply air energy	0.043

### **4.4.3 Hierarchy of Commercial Building Sensors**

The energy use prediction scores in GJ from Table 4-3 were ranked as the sum of all energy use prediction scores for each sensor and rearranged in Table 4-4 grouping the cooling and heating season to generate a relative hierarchy with respect to energy estimation potential for the measured building. Within this ranking, five different tiers were developed based on significant gaps in energy use prediction scores. The most uncommon but incrementally valuable (with respect to energy end use estimation) sensors in modern buildings reside in tiers two and three, where hydronic radiant heater information can provide reasonable proxies for the estimated energy use of multiple end uses in the measured building. Additionally, window state, plug load, and multi-stage illuminance offer similar value with respect to energy use estimation to IAT sensors. The most valuable sensor from an energy use estimation perspective in the measured building was the OAT sensor. This study suggests that OAT sensors can be used for many energy estimation applications when coupled with spot measurements and that incremental accuracy of OAT sensors could have the greatest benefit of all the sensors assessed. However, the sensor hierarchy in Table 4-4 is dependent on the energy use breakdown of the building studied (which had a high outdoor air ratio) and could vary significantly for other buildings. Table 4-4 should be used as an example of applying the hierarchy approach to each end use and demonstrating the best value for specific building types (i.e., the building measured in this study). The end use-specific rankings of Table 4-3 are intended to be more widely applicable to other buildings.

**Table 4-4: Sensor hierarchy using aggregated input scores from all end uses for both the heating and cooling season for the measured building. The sensors have been divided into tiers based on distinct gaps in total score.**

Input	Total score (GJ)	Tier
Outdoor air temperature	15.18	1
Radiant heater flow rate	3.42	2
Radiant heater valve position	3.40	2
Radiant heater inlet temperature	2.52	3
Radiant heater outlet temperature	2.44	3
Reheat coil outlet temperature	1.37	4
Hour of day	1.36	4
Illuminance in center of room	1.25	4
Illuminance near window	1.23	4
Occupancy state	1.22	4
Window state	1.22	4
Indoor air temperature	1.21	4
Illuminance near wall	1.19	4
Plug load	1.11	4
VAV airflow	1.10	4
Reheat coil flow rate	0.23	5
Lighting dim level	0.23	5
Lighting state	0.23	5
Blind position	0.22	5
Return air temperature	0.19	5
Lighting power	0.16	5

#### 4.4.4 Limitations

Numerous limitations exist in this study with respect to general applicability, including but not limited to:

- Measurements are required over a multi-week period for each season to calibrate proxy sensor models. These measurements include all the sensors to quantify the end use that is modelled (e.g., hydronic flow, inlet temperature, and outlet temperature for radiant heating). Some spot measurements (e.g., hydronic temperature and flow) can be cumbersome and require expensive equipment.
- Building thermal and occupant properties from this study may not be indicative of typical buildings, or even other buildings of the same type with similar HVAC systems. The results of this study may also only be applicable to a humid continental climate as per the Köppen-Geiger climate classification (where the measured building is located) (Kottek, Grieser et al. 2006).
- Other model forms and structures may yield better results. Feed forward ANNs were chosen because of their robust and flexible continuous non-linear modelling characteristics. Other ANN structures and types may also generate different results with a full factorial approach.
- Other model selection targets would yield different results (e.g., fault detection, indoor environment quality, real-time control decisions). This study focused on mid-term to long-term energy use, and it is acknowledged that the models selected were not the best models for time step level accuracy as required by the exempli gratia.

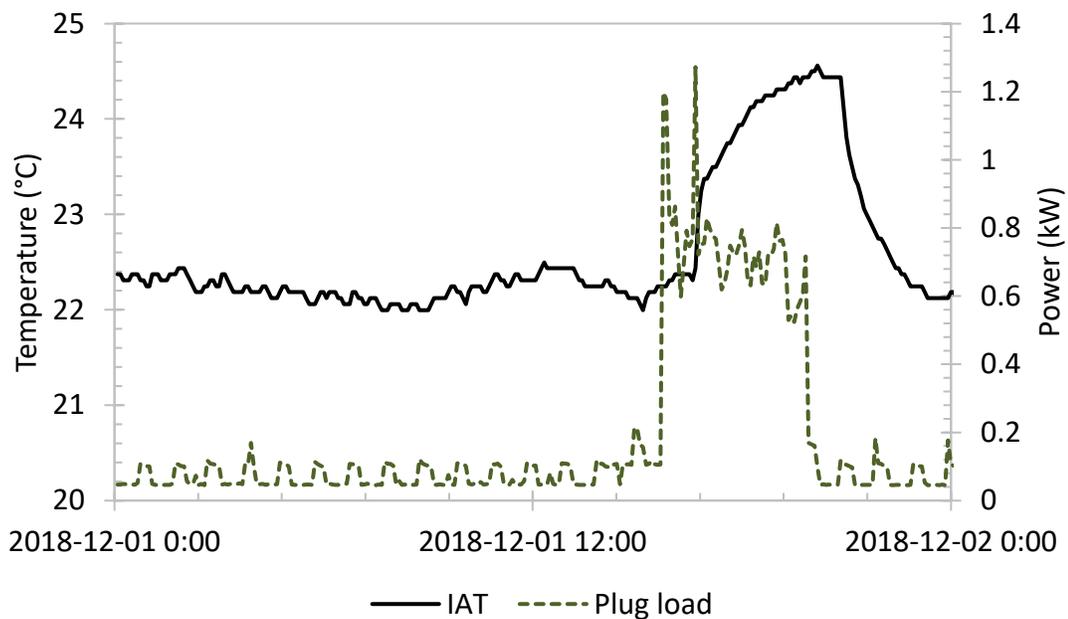
- Past sensor data was not used on the current time-step in the proxy sensor models. It is indicated in Figure 4-8 that thermal lag may have an impact on the MISO proxy sensor accuracy.
- Only part of the heating season was evaluated. The cooling season will not have radiant heating end uses but will have cooling end uses. As such, a different sensor hierarchy would occur, potentially even generating different hierarchies for plug loads and lighting in the cooling season. Additionally, AHU coil energy use analysis in the cooling season could introduce complexity due to latent coil loads.
- Mechanical systems (e.g., constant air volume, fan coil, entirely radiant, natural ventilation) and control methods (e.g., manual, deadband, proportional-integral) could affect the order of the hierarchies.
- The outdoor air ratio of the fresh air supply system could have a significant effect on the hierarchy at the building level (and potentially the heating and cooling specific end uses).
- The higher cost of electrical energy was not distinguished from thermal energy. Additionally, factors beyond the building level (e.g., emission factors, conversion factors, combustion efficiency) were not considered in the valuation of energy.
- The sensor hierarchy is end use magnitude specific (i.e., the sensors relevant to the higher end uses received higher energy use prediction scores). In this case lighting energy use was so low that sensors only viable for lighting estimation did not receive a high score.

#### 4.4.5 Anecdotal Observations

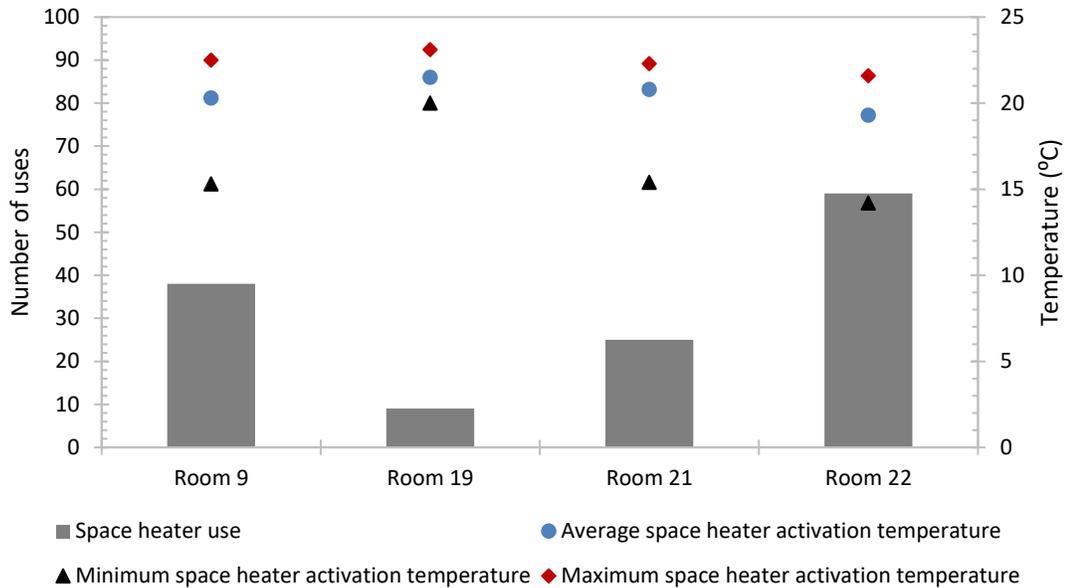
The systematic approach described in the Methodology section provides the framework for scoring and ranking sensors for energy use estimation but does not provide context to some anomalies observed during the measurement period. Two anomalies that would not have been easily noticeable without the BMS embedded sensors were the use of personal heaters and the differences between the room-level indoor and return air temperatures.

During this study electric personal heaters were observed to be present in several offices. Attempts were made to identify the frequency and duration of space heater usage from the plug load data. Space heaters were considered on if the plug load for the office was higher than 500W, a value inferred by manual inspection of the data. Normal electrical plug usage was typically below 300W outside of these observed anomalies, and spikes in indoor air temperature were also observed within an hour of these events. A limitation of this portion of the study is the inconclusiveness of the plug load being a space heater due to the similar sensible heat gains from any high-power plug-in equipment (e.g., kettles, refrigerators) and the lack of ground truth. Suspected space heater usage in an office is compared to IAT in Figure 4-11 for reference. A new space heater usage was conservatively logged when there was eight hours or more between periods of plug loads greater than 500W to estimate the number of workdays the space heater was used. Four rooms exhibited the characteristics of regular space heater usage over the 123 days of measurement. These rooms (9, 19, 21, and 22) are also observed to use less heating energy in Figure 4-9. The indoor air temperature at the moment that space heaters were turned on and the number of days of their usage is shown in Figure 4-12. The total energy usage

exclusively due to the previously defined space heater conditions was estimated to be up to 1.72 GJ over the four-month measurement period. Only two rooms had space heater usage that was estimated to be greater than 0.35 GJ (room 16 and 22). Additionally, occupants in these rooms were recently relocated from another building on campus and previously had space heaters to cope with poor thermal conditions in an older building suggesting usage may have been more related to habits than extreme discomfort.



**Figure 4-11: Suspected space heater usage in an office due to high plug load and the corresponding effect on IAT**

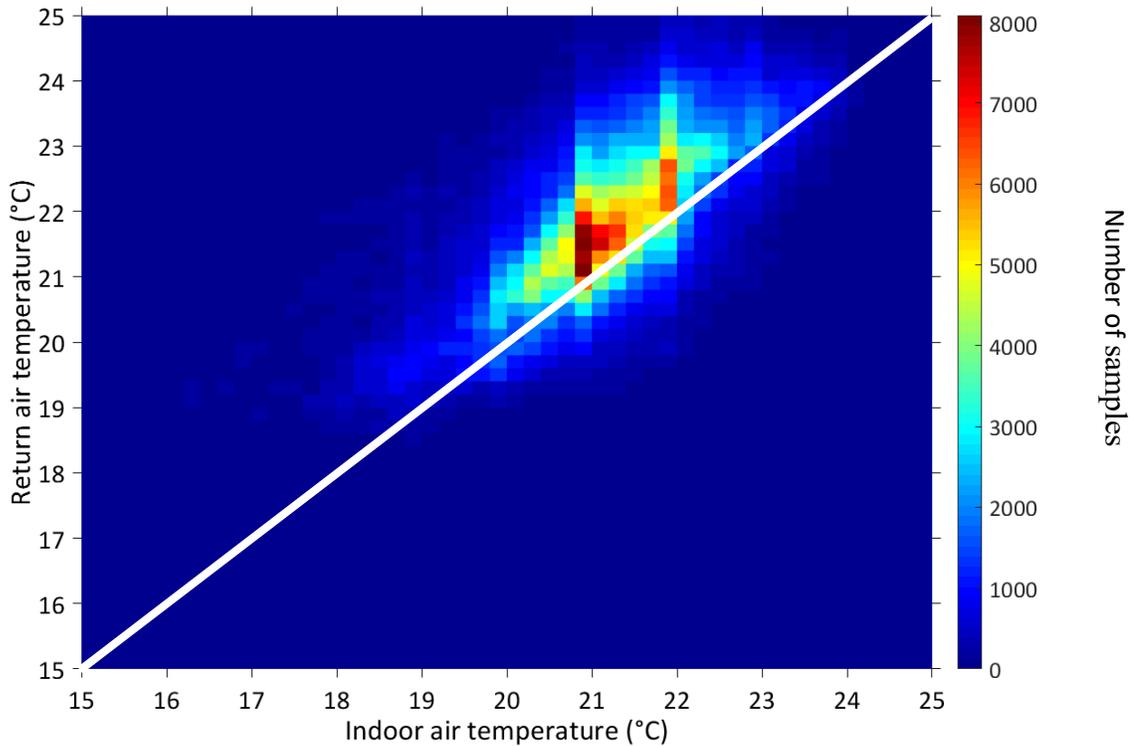


**Figure 4-12: Suspected space heater temperatures and usage data**

Many data-driven room-level models rely on indoor air temperature as an input. Another common assumption is that indoor air temperature can be used as a proxy for room temperature and return air temperature (Gunay, O'Brien et al. 2016, Bursill, O'Brien et al. 2018). However, this assumption may not be accurate in practice. Room-level return air sensors are not common in commercial buildings while AHU-level return air sensors are required for economizer control to adhere to ASHRAE Standard 90.1 (ASHRAE 2016). In this study the indoor and return air temperature (sensor location is shown in Figure 4-1) readings were compared for all rooms over the measurement period. A heatmap plot of the return air temperature and indoor air temperature over the entire measurement period is provided in Figure 4-13 where the number of temperatures within 0.2°C by 0.2°C bins is represented by the color gradient. A line where the slope is equal to one is also overlaid on

the graph to demonstrate the bias between IAT and RAT for the rooms. The r-squared value for Figure 4-13 was 0.52. Time lags from zero to 120 minutes were evaluated for fit, and the r-squared value improved to 0.57 if a lag of 90 minutes was introduced between the return air temperature and indoor air temperature. This confirms that lag due to air change is present, but also that it is not enough to improve the correlation between IAT and RAT significantly.

Airflow changes did not appear to have a significant effect on the temperature difference. Sensor position is likely a significant factor because the IAT sensor is on the wall thermostat where it may be influenced by longwave radiation and solar gains, while the RAT sensor is in the return air stream and shielded from these effects. Air stratification and supply to return short circuiting could also have added to the bias in Figure 4-13. The results of this sensor comparison demonstrate the importance of using a direct return air temperature sensor for sensitive modelling applications and the possible error of assuming the IAT and RAT are equal. An average difference between IAT and RAT of 1.06°C was measured in this study, which could be significant when calculating air side energy usage for low temperature differential systems.



**Figure 4-13: Heatmap comparison of indoor air temperature and return air temperature for all rooms.**

#### 4.5 Summary

Determining the optimal sensors for commercial buildings zone-level energy use estimation is an open challenge in the building control industry. This study attempted to provide some insight into this problem by analyzing measurements from a group of heavily instrumented rooms in an institutional building in Ottawa, Canada. It was found that many of the sensors added for research purposes were not more useful for accurately estimating aggregated zone-level energy use than common sensors (e.g., OAT, lighting state, occupancy state, and hour of day). The most valuable sensor from an energy use estimation perspective was the OAT sensor due to its ability to be a proxy for estimating multiple

energy end uses accurately over the validation period of two months. The following tier of useful sensors was from the flow and temperature sensors on the hydronic radiant heating system. The lowest scoring sensors on the hierarchy were only useful for lighting energy use estimation (or in the case of reheat coil flow rate, incremental energy estimation in a forward selection model). Sensors not present in the hierarchy were not found to be suitable proxies for energy use estimation and should be valued lower than the bottom tier. Despite the promising results found in this study, model calibration is required for each proxy sensor.

Some anecdotal findings identified additional sensor benefits that were not encompassed by the sensor score and ranking. Addressing the usage of personal space heaters in these offices could both save electrical energy and improve occupant comfort. Room-level return air temperature was found to vary unpredictably by an average and standard deviation of 1.06°C and 1.03°C respectively when compared to indoor air temperature measured at the thermostat. Hydronic heating energy (room-level radiant heating, VAV reheat, and AHU coil heating) was found to comprise an average of over 95% of the energy usage of the 27 rooms in this study, confirming that heating loads are dominant in the heating season despite the trending increase in plug loads in commercial buildings.

The results of this study can be assumed to be generalizable to new construction (e.g., low lighting loads and high HVAC loads) and buildings subject to similar conditions. However, this study also has many limitations due to the data-driven approach used (all of which are presented in detail in the Results and Discussion sections). The most pertinent

limitation is the requirement of model validation measurements from costly and cumbersome temporary equipment.

## Chapter 5: Lab Implementation of Classification Learning-based sMPC

This chapter is based on the supplementary material of:

Jayson Bursill, William O'Brien, Ian Beausoleil-Morrison. 2019. Experimental application of classification learning to generate simplified model predictive controls for a shared office heating system. *Science and Technology for the Built Environment* 25 (5), 615-628.

The application of model-based predictive control (MPC) to commercial buildings has been a topic of increasing interest in the literature. Despite significant energy saving potential, commercial adoption has been limited. The literature suggests that the limited proliferation of MPC in buildings is due to a mismatch between the mathematical complexity of MPC and the expertise of practitioners. Accordingly, simplified MPC (sMPC) approaches may improve uptake. This paper applies an approach to MPC where rules are extracted from a detailed MPC algorithm to produce an algebraic model of fewer than 50 lines of code. The goal is to produce sMPC that will be more accepted by practitioners. Testing in a shared office was executed over a five week period of the heating season from March 11<sup>th</sup> to April 14<sup>th</sup>. The detailed MPC algorithm presented in this paper reduced heating energy consumption by 23% compared to conventional reactive control when tested under similar weather conditions. The sMPC algorithm achieved 95% of the heating energy savings of the detailed MPC algorithm with significantly less complexity. Based on this favorable result for a single case, further work should be conducted with larger sample sizes to determine if rule extraction can produce effective sMPC.

## 5.1 Introduction

Reducing building energy consumption is one of the central approaches toward addressing the long-term goals of sustainability. Further, reducing energy use in building heating and cooling systems via improved control presents an opportunity with low capital costs and tangible operating cost savings. Significant research efforts have been made to explore advanced algorithms implementing various forms of model predictive control (MPC). Many of these studies however, are limited to a simulation environment, thus minimizing the ability to adjust to disturbances and stochastic occupant conditions (Braun and Chaturvedi 2002, Henze, Felsmann et al. 2004, Oldewurtel, Parisio et al. 2012, Coffey 2013, Afram and Janabi-Sharifi 2015). Building upon this work, at least 11 experimental and in-situ studies have been executed with promising results (Roth, Llana et al. 2006, Privara, Široký et al. 2011, Široký, Oldewurtel et al. 2011, Afram and Janabi-Sharifi 2014, Shaikh, Nor et al. 2014, Touretzky and Baldea 2014, Goyal, Barooah et al. 2015, Gunay, O'Brien et al. 2016, Gunay, O'Brien et al. 2016, Hilliard, Swan et al. 2016, Thieblemont, Haghghat et al. 2017). Despite this effort, there has been minimal commercial adoption of predictive controls. The reasons for slow commercial uptake of MPC in buildings are difficult to diagnose and address, but have been observed industry-wide (Killian and Kozek 2016). A similar delay of proliferation was observed in the process industry (Li, Yu et al. 2011), where sole ownership of equipment from design to active lifetime renders incremental MPC related capital spending more attractive when compared to the high turnover building construction industry.

Commercial buildings represent a unique challenge for control implementation, where the requirements and sequences differ significantly from project to project. This necessitates simple and transparent solutions for building control that can be applied quickly by technicians. The proposed approach involves first establishing a detailed MPC algorithm (Wang 2009) to control an office using weather and occupant predictions to determine optimal current and future heating energy input. After two weeks of operation, classification learning in the form of simplified model-based predictive control (sMPC) was used to develop simple thresholds for heating operation from the detailed MPC data. The resulting control algorithms are compared with two weeks of run time to the baseline control implemented by the control contractor during construction. The objective of this paper is to provide a workflow for assessing the viability of rule extraction from detailed MPC in a simple form for an experimental office. It is assumed that this approach can be scaled for multiple offices on future implementations. The following sections of this paper describe the process of MPC and sMPC development, energy savings quantification, and the respective results observed in an active office.

## 5.2 Literature Review

The literature review in this study initially introduces the potential benefits of MPC identified in previous studies. Pertinent studies and their contributions to the literature (relevant to this study) are then presented. From the body of previous work, gaps that present opportunities for novelty in predictive controls are identified.

Control and operational improvements have been estimated to provide potential site energy savings of up to 50% (Tzempelikos and Shen 2013, Dong, O'Neill et al. 2014). With new construction estimated to account for only 1-2% of the annual building stock (Popescu, Bienert et al. 2012), the importance of retrofit based control and operation improvements is apparent. MPC presents a solution that can be entirely software-based. Further, MPC has experienced an increase in journal and conference publications over the last decade in building literature. The foremost barrier to implementation of MPC in commercial projects is the multidisciplinary complexity of the embedded algorithms (Killian and Kozek 2016). An online MPC algorithm may draw from the following bodies of knowledge: occupant comfort, weather prediction, thermal modelling, optimization, and statistics.

Maintaining or improving occupant comfort is important for any control modification made to an existing building. The primary purpose of the heating and cooling system is to provide comfort to the occupants, and neglecting this can lead to unpredictable adaptive occupant behaviors (Gunay, O'Brien et al. 2013). The quintessential metrics for occupant comfort were developed by Fanger using surveys, experiments, and extensive modelling approximations (Fanger 1970). The Fanger approach is fundamental in ASHRAE Standard 55, and inspired modern work in the field (ASHRAE 2017).

MPC has been defined as an algorithm that optimizes the setpoint trajectory by limiting equipment demand (Lee and Braun 2008). Input prediction methods are required, and weather forecasting is a popular approach to obtaining environmental predictions. Periodic weather predictions can facilitate MPC by predicting input values over several hours (Gwerder, Gyalistras et al. 2013). Florita & Henze suggest that neural network models are acceptable for weather forecasting for control applications, and that more complicated models are not necessary (Florita and Henze 2009). For Canadian locations, Natural Resources Canada has developed CanMETEO, a program that utilizes a numerical approach that divides the atmosphere into numerous control volumes and solves periodically based on Environment Canada weather observations contained within grib files (Candanedo 2018). The availability and connectivity of modern weather services has provided numerous feasible solutions for real-time weather forecasting.

Fundamentally, greater model complexity can improve the energy savings of a MPC algorithm with better prediction accuracy. However, higher complexity will increase the required computation resources, mathematical complexity, and time. In previous studies a prediction horizon of two to three hours was found to be appropriate for buildings without abundant thermal mass (Gunay, Bursill et al. 2014). Under this assumption simplified “grey-box” state-space models can provide appropriate thermal modelling accuracy. In model-based predictive control applications, lumped capacitance and state-space models have been applied successfully for decades (Penman 1990, Gouda, Danaher et al. 2000, Braun and Chaturvedi 2002, Chen and Athienitis 2003, Bacher and Madsen 2011, Saberi Derakhtenjani, Candanedo et al. 2015). System identification is required for grey-box

MPC, and can be based upon the MATLAB System Identification Toolbox approach (Ljung 1995). Best practices of model selection and fitting have been summarized in the following points by Rouchier (Rouchier 2018):

- The identifiability of the model structure can be checked before solving the inverse problem to avoid parameter interactions and redundant parameters.
- The impact of approximation errors cannot be overlooked. All models are biased to some extent, which adds up to the effect of measurement uncertainty, and may result in disastrous estimation errors if the proper validation steps are not carried out.
- Validation steps should include a form of residual analysis to diagnose unaccounted phenomena, and an estimation of the confidence regions of parameter estimates.
- Choosing the appropriate model structure is a bias-variance trade-off: a detailed model offers a better fit with data than a simple one, but excessive complexity will offer poor prediction accuracy through over-fit.

When the appropriate steps of model prediction have been followed, a cost function can be minimized to determine control decisions. One of the simplest and most robust forms in the literature is presented by Wang (2009). Many applied MPC algorithms have been based on state-space inverse modelling of controlled spaces. Uncertainty in weather prediction has been addressed by an approach involving the addition of stochastic elements to MPC, yielding some improvements in setpoint control by modifying the state-space model (Oldewurtel, Parisio et al. 2012). Different state-space zone air temperature models have also been explored to determine suitability for private offices (Gunay, O'Brien et al.

2016). A recent exception to the state-space modelling approach involved weather forecasts and detailed modelling using EnergyPlus building simulation software to predict the building response over the prediction horizon as an application to complicated heating and cooling systems (Hilliard, Swan et al. 2016). In many global regions, time of use billing is present and can necessitate more complex optimization functions to take further advantage of short-term thermal storage in building thermal mass. The application of modified MPC approaches using time of use weighting factors has observed additional baseline energy savings compared to flat rate approaches in the literature (Henze, Brandemuehl et al. 2008, Ma, Qin et al. 2012, Cole, Powell et al. 2014, Li and Wen 2014).

Despite the clear and tangible energy savings that can be obtained from MPC in buildings, commercial adoption of advanced control algorithms at the zone level is virtually non-existent (Killian and Kozek 2016). Commercial adoption is an ongoing challenge, and the delays observed in early MPC adoption for the manufacturing industry may be exacerbated in the turnover-oriented construction industry. Simplified approaches to MPC based on applied statistics have presented promising solutions to expedite commercial usage of predictive algorithms. May-Ostendorp (2011) utilized classification learning to determine combinations of different input thresholds for control outcomes categories in the form of “rule extraction”, borrowing the concept from applications in other fields. The importance of large quantities of data to validate classification approaches is highlighted by Coffey (2013), where simulation was used to develop optimal control rules for a modelled building stock in the form of a lookup table.

The primary gap in the literature that this paper addresses is the lack of simple, applied predictive algorithms tested in-situ. Additional gaps include: qualitative and quantitative insights from live implementation, documentation of challenges and unexpected disturbances, and commentary on scalability. The following sections will describe the forecasting, modelling, optimization, and classification approaches, and the accompanying results obtained from a single office implementation of a simplified form of rule extraction.

### **5.3 Methodology**

Extracting rules from detailed MPC involves first developing an MPC algorithm that can improve building operation, then applying classification learning to determine thresholds of the inputs and outputs of the algorithm. The goal of simplified rule extraction is to find the minimum number of inputs, thresholds, and parameters required to maintain most of the MPC energy savings while maintaining the same level of comfort for the occupant. This approach provides a practical commercial implementation by removing the optimization calculation from building automation controllers where matrix operations are limited. Matrix multiplication and inversion cannot be executed in database programs of most building controllers without manual addition of code, rendering commercial usage of database matrix operations limited.

Simplified MPC based on rule extraction requires knowledge of the conditioned space characteristics in the form of a model. One month of preliminary sensor data was collected from the building automation system serving the office to train and validate inverse models for predicting future office air temperatures in the MPC algorithm.

Following detailed MPC testing, classification learning was utilized to extract rules for sMPC control and corresponding testing. Simplified MPC still required a six parameter model to predict office air temperatures for the next time step, but all other prediction information was not required. Hydronic input to the radiant heating panels and the supplied air from the AHU were compared between the different weeks to quantify energy savings. The following sections of the methodology describe in detail the previously mentioned processes.

### **5.3.1 Experimental Office Characteristics and BAS infrastructure**

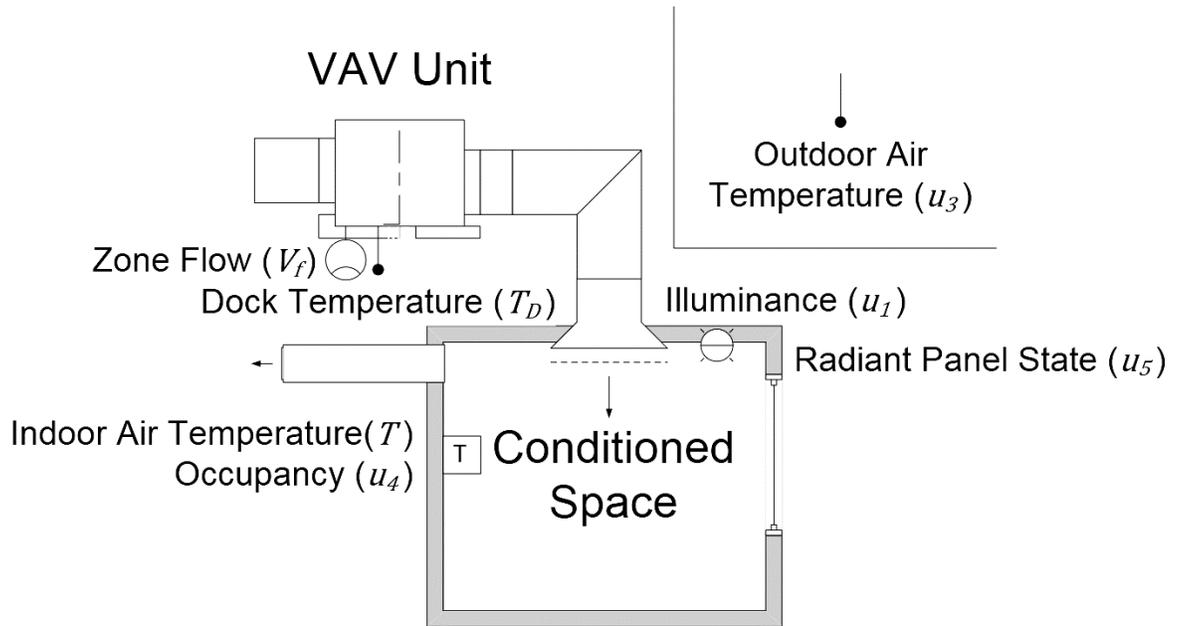
The office studied was a multiple occupant corner office on the 7<sup>th</sup> floor of an 8-year-old institutional building in Ottawa, Canada with an external wall facing 30 degrees south of West. The external wall has an area of 15 m<sup>2</sup>, a window to wall ratio of 0.5, and contains RSI 3.0 rigid insulation. The office has a floor area of 15 m<sup>2</sup>, a 25 cm thick concrete slab floor covered with carpet, and a concrete ceiling slab with an acoustic tile dropped ceiling. Adjacent to the office is a corridor and another office, both of which are conditioned. The office features automatic roller blinds with 10% solar transmittance, 30% solar absorptance, and 10% visible transmittance that were manually controlled in this study. Lighting was provided to the office by four 32W T8 recessed fluorescent tubes that were manually controlled by the occupants. Table 5-1 provides a summary of the sensors used for the MPC reduced order inverse models of the space, and Figure 5-1 shows the sensor locations within the space. Data was collected at 10-minute sample intervals over five weeks from March 6<sup>th</sup> to April 15<sup>th</sup>, 2018 in the order shown in Table 5-2. Occupancy data was recorded such that if an occupancy sensor was triggered after the last sample and prior

to the upcoming sample, the occupant state would read as “1”.

The general control schemes that will be evaluated in this paper can be categorized as baseline reactive control, model-based predictive control (MPC), and rule extraction based sMPC. The reactive control is an unmodified form of the stock controls provided by the control contractor for the building during construction. This form of control includes several physical inputs (Table 5-1) and variables to provide information to the algorithm, and two previously existing proportional-integral (PI) control loops on heating and cooling outputs that were configured by contractors when the building automation system was initially programmed. In this case, the heating and cooling outputs are the radiant heating panel and the VAV box flow rate respectively. A user-defined setpoint is maintained by controlling the VAV damper and radiant heating panel valve using the previously mentioned PI loops.

**Table 5-1: Input sensors for Reduced Order Models**

Input	Unit	Sensor
Indoor air temperature ( $T$ )	°C	10 kOhm thermistor
Illuminance ( $u_1$ )	Lux	Photodiode
VAV proxy ( $u_2$ )	L-°C/s	Differential pressure transducer and 4-20 mA resistance temperature detector
Outdoor air temperature ( $u_3$ )	°C	10 kOhm thermistor
Occupancy ( $u_4$ )	None	Passive infrared
Radiant panel valve position ( $u_5$ )	%	10 VDC signal



**Figure 5-1: Location of building automation system sensors used for reduced order modelling within the measured office**

### 5.3.2 Experimental Approach

Experimental evaluation of the energy saving potential (relative to reactive control) of the different algorithms was tested by alternating algorithms within the space on a weekly basis as shown in Table 5-2. Two weeks of heating season data were obtained for reactive control

and MPC, and one week of data was obtained for sMPC, where the two weeks of reactive control and MPC were not contiguous. A single space was used sequentially (rather than multiple spaces in parallel with a test and control group) which produced different weather and occupant conditions that varied as shown in Table 5-2.

**Table 5-2: Weekly algorithm schedule, average OAT, and average occupant fraction**

Week	Start	End	Algorithm	Average outdoor air temperature (°C)	Fraction of time with at least one occupant
1	March 11 <sup>th</sup>	March 17 <sup>th</sup>	Reactive	-6.1	0.05
2	March 18 <sup>th</sup>	March 24 <sup>th</sup>	MPC	-2.0	0.06
3	March 25 <sup>th</sup>	March 31 <sup>st</sup>	Reactive	1.8	0.07
4	April 1 <sup>st</sup>	April 7 <sup>th</sup>	MPC	-2.2	0.08
5	April 8 <sup>th</sup>	April 14 <sup>th</sup>	sMPC	-1.2	0.04

The MPC was executed using a control agent to determine the result of the cost function in Equation 5-1 following the approach outlined in (Wang 2009). Equation 5-1 is the general form of an MPC cost function that can contain individual values or matrices where  $J$  is the arbitrary cost constant,  $SP$  is the temperature setpoint provided by the occupant,  $T$  is the air temperature as observed by the BAS or predicted using modelling,  $\Delta U$  is the energy input from the heating and cooling system, and  $R$  is a matrix of penalty weights for different time steps within the prediction horizon. The application of Equation 5-1 using matrices will be described later in this chapter. In this study the radiant panel valve position was the only input controlled by the MPC. All other inputs were for the thermal model. The proposed cost function is formulated to minimize the difference between the air temperature and setpoint and minimize energy consumption in the conditioned space simultaneously. The matrix  $R$ , set as the identity matrix, is intended to

provide a compromise between maintaining the setpoint and minimizing energy use. The thermal model for the space was as shown in Equation 5-2a, where  $t$  is the time step index,  $u$  is the input array, and  $A$  and  $B$  are the linear time invariant constants for each respective space. States and inputs were constrained by the limits of the sensor readings (and respective equipment capacities). The parameter values of  $A$  and  $B$  were unconstrained. The model must be retrained when input ranges change significantly (e.g., seasonal transition, furniture addition or removal, envelope upgrade) to maintain prediction accuracy and account for differing environmental and occupant conditions. The input array for the model selected in this study consisted of the five sensors in Table 5-1 adapted from recommendations in previous literature where  $u_1$  is the office ceiling illuminance in lux,  $u_2$  is the VAV proxy,  $u_3$  is the outdoor air temperature,  $u_4$  is the occupancy state, and  $u_5$  is the radiant panel valve position. This expansion modified Equation 5-2a to Equation 5-2b. The illuminance sensor is used as a proxy for the solar gains of the space because surface irradiation correlates with ceiling illuminance, and the office studied had illuminance sensors available. The use of illuminance sensors as a solar proxy was also found convenient and practical by Gunay (2014). The VAV proxy is the product of the VAV unit discharge air temperature and the flow rate. Those inputs were chosen because it was observed that illuminance correlated adequately with solar irradiation on the window while accounting for blind position. Further, the VAV proxy improved orthogonality (reduced correlation of model inputs) with respect to indoor air temperature when compared to a typical first principles-based air side heat balance.

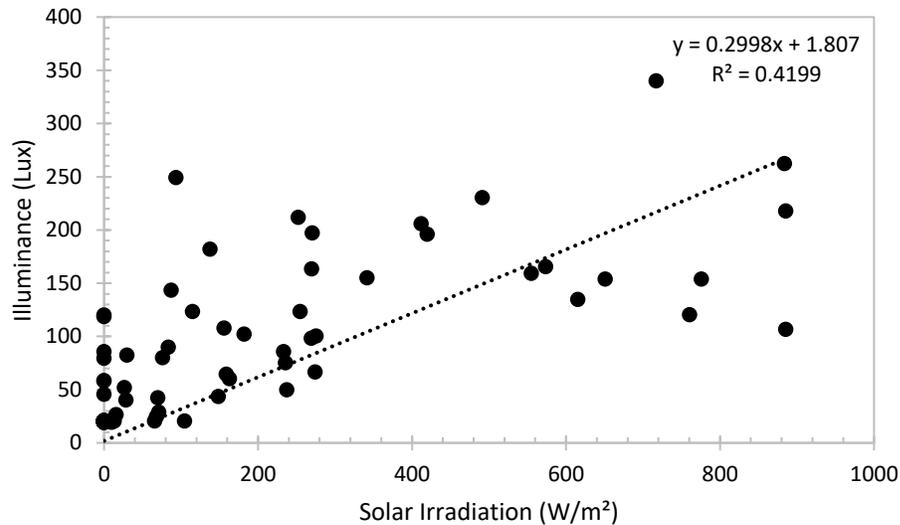
$$J = (SP - T)^T(SP - T) + \Delta U^T R \Delta U \quad (5-1)$$

$$T(t + 1) = AT(t) + Bu(t) \quad (5-2a)$$

$$T(t + 1) = AT(t) + B_1u_1(t) + B_2u_2(t) + B_3u_3(t) + B_4u_4(t) + B_4u_4(t) + B_5u_5(t) \quad (5-2b)$$

Following the selection and training of an appropriate model (Wang 2009, Gunay, Bursill et al. 2014), a method of prediction was developed. The algorithm used local weather station predictions provided at six-hour intervals from the CanMETEO program developed by NRCan for illuminance (converted from solar radiation, used as a solar proxy) and outdoor air temperature prediction. The predicted incident solar radiation on the wall from CanMETEO was adjusted to an equivalent illuminance (lux) value using a linear fit to five days of hourly data with a linear  $R^2$  of 0.42 when comparing recorded predictions to observed real values as shown in Figure 5-2. Both the site observations and CanMETEO predictions use different sensors in different orientations to measure different quantities, thus significant error is expected. Degradation of the line fit can be attributed to lighting usage at the room level and the differing physical quantities and local solar conditions being measured between the airport weather station and building location (roughly 8km apart). Differences in observed weather and weather predictions is a significant source of error for indoor air temperature predictions within the MPC. MAE of 100.1 lux and 1.8 °C were found for converted illuminance and outdoor air temperature

predictions respectively when compared to corresponding Environment Canada observations. The sensitivity of the reduced order thermal model to the illuminance and outdoor air temperature MAE (as a product of the linear parameter and MAE) was  $0.06^{\circ}\text{C}$  and  $0.02^{\circ}\text{C}$  respectively. Despite the low  $R^2$  value between the solar irradiation and illuminance predictions, this fit is considered adequate for the application in this study due to the cost of installing surface irradiation sensors in most commercial buildings. Of the commonly available sensors in conditioned spaces (and available in this study), illuminance sensors are the most direct proxy for estimating room-level solar gains and can infer the effects of glazing geometry and blind position. Additionally, illuminance sensor can be used to simultaneously infer the heat contribution of solar gains and lighting. The building used in this study did not have particularly large solar gains (which may have improved the effectiveness of this approach). While useful in this study for prediction future indoor air temperature, this linearization approach to solar gains may not be as effective when solar gains are more significant.

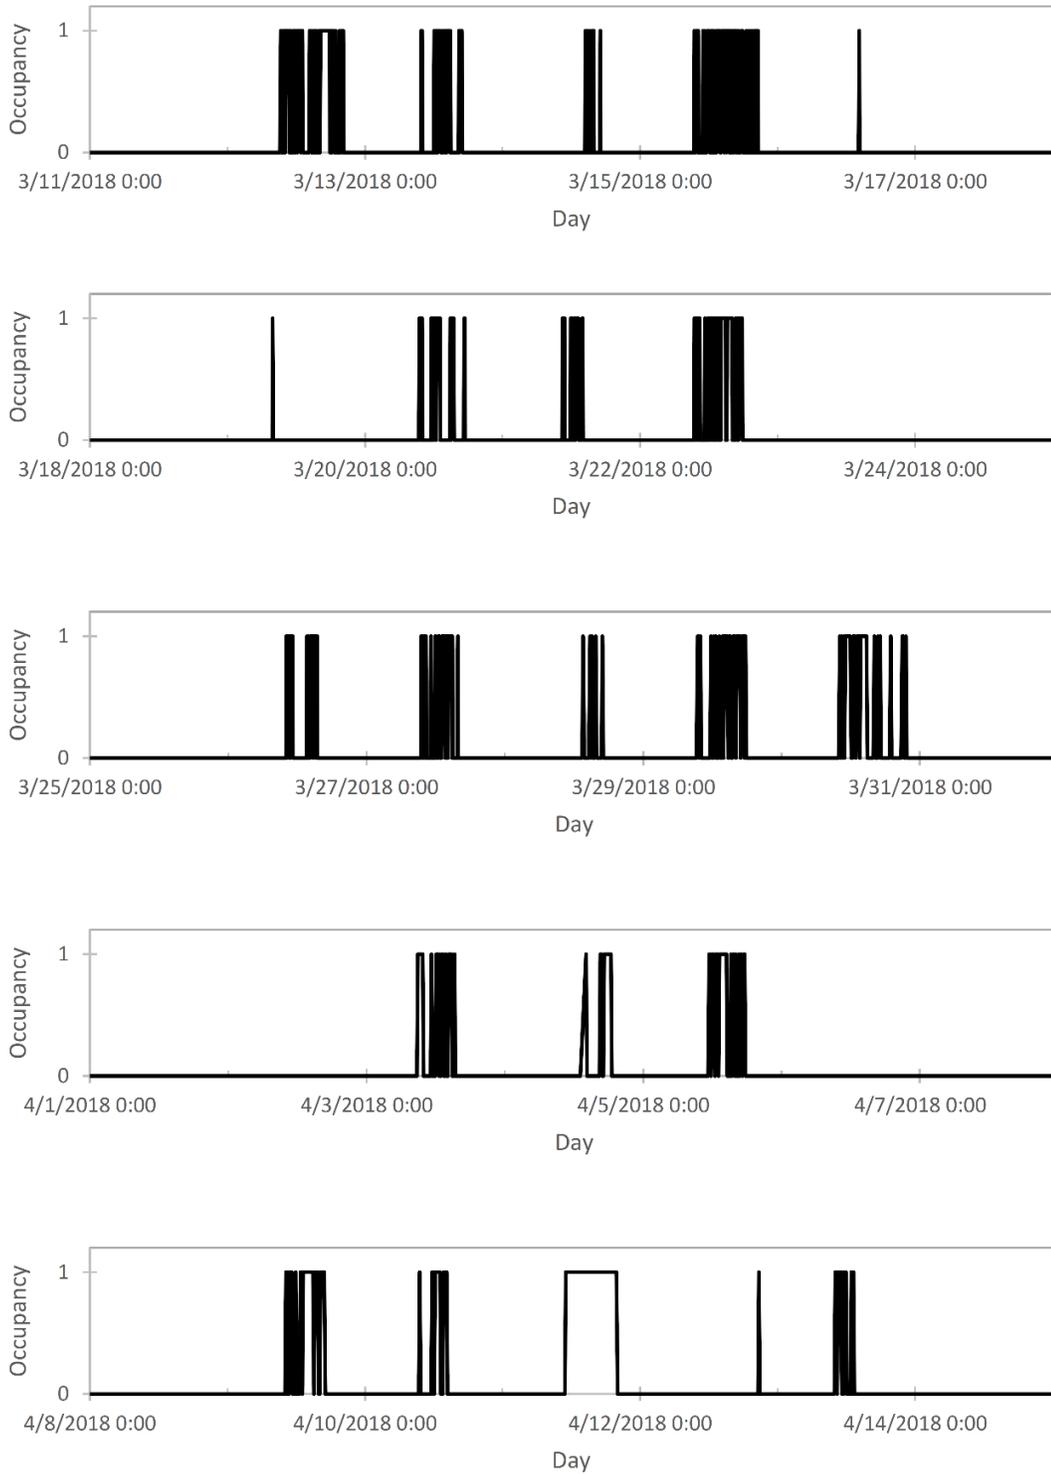


**Figure 5-2: Observed ceiling illuminance linear fit to CanMETEO solar irradiation predictions at the office wall orientation.**

The timing of first occupant arrival and last occupant departure events were conservatively predicted based on the previous two weeks of occupancy data. These events were assumed normally distributed. The occupied schedule was taken as the mean first arrival minus 1.5 standard deviations to the mean last departure plus 1.5 standard deviations. This approach accounts for roughly 90% of all arrivals and departures. The time between the estimated arrival and departure times on weekdays was assumed to be occupied in the prediction algorithm, and the weekends were assumed to be unoccupied for temperature prediction. The observed event-based occupancy over the five-week testing period is provided in Figure 5-3, where 0 and 1 indicate the room is unoccupied and occupied respectively. The purpose of Figure 5-3 is to provide context of typical actual arrival and departure times over the testing period. At the time scale in Figure 5-3, brief absences (e.g., five minutes) are indistinguishable from longer absences (e.g., a half hour).

Oscillation of occupancy state is due to the five-minute time-out period of the occupancy sensor in the office.

Single node state-space models were selected for this application for simplicity and acceptable accuracy for short prediction horizons (Gunay, Bursill et al. 2014). Additional nodes could be included with the same inputs and disturbances, however diminishing returns in prediction accuracy were observed for increasing complexity (particularly for shorter prediction horizons). The cost function minimizes energy consumption. It was assumed that the occupant could control the room setpoint to a desirable temperature, and that ASHRAE 55 comfort criteria and a low number of thermostat interactions were indicative of occupant comfort. A desirable MPC algorithm will reduce energy consumption compared to reactive control while maintaining, and hopefully improving, occupant comfort. As such, the occupant comfort metrics of indoor air temperature range and thermostat interaction were used to compare the baseline reactive control with the MPC and sMPC algorithms. Details of the MPC modelling and control, and the rule extraction classification learning for sMPC are provided in the following sections.



**Figure 5-3: Occupancy profiles for the testing period where a 0 indicated the room was unoccupied and a 1 indicated the room was occupied. The minimum occupied time for the sensor was 15 minutes.**

### 5.3.2.1 Modeling and Control

The MPC in this paper was developed in a similar form as described by Wang (2009), with adaptations for a single node state-space room thermal model. For this study, Equation 5-3 is formulated given Equation 5-1 and Equation 5-2, where  $\varphi$  is a square 6x6 matrix containing the state parameters as shown in Equation 5-4. The state parameters ( $A$ ) are the same as in Equation 5-2. Most of the resulting matrix operations in Equation 5-3 can be reduced to the 6x6 Hessian,  $H$ . The simplified form of Equation 5-3 can then be used to find six time steps of predicted energy changes in the space with the 6x1 matrix,  $\Delta U$ , using 6x1 input matrices of current and predicted setpoint and indoor air temperature,  $SP$  and  $T$  respectively. The matrix in Equation 5-3 is representative of a six time step prediction horizon with 30 minute time steps giving a three hour prediction horizon. Only the first time step of Equation 5-3 uses real data, and the remaining time step data for  $T$  are predicted by using Equation 5-2 incrementally. The setpoint,  $SP$ , is assumed to remain constant over the prediction horizon unless there is a change in the real setpoint (in which case the new setpoint will be assumed to remain constant over the prediction horizon). Time steps of 30 minutes were chosen to maintain thermal model stability and dynamic response to changing input conditions. The algorithm was executed every 10 minutes and was limited by network speed. Half of the prediction horizon (three time steps) was utilized for the control horizon. The sign of the net energy balance over the control horizon and setpoint deviation were compared, and if the net energy balance over the control horizon was negative (indicating energy loss) and the temperature was below the room setpoint, the algorithm would return a 1 indicating the need of heating over the control horizon.

Otherwise, the algorithm would return a 0, allowing the space to thermally free float. Pseudocode of this process is provided in Figure 5-4, where the returned value allows the space to operate using the PI loops programmed by the controls installer (returns 1) or suppresses the radiant panel valve position to zero (returns 0). The algorithm output was appended to the existing code in the reactive control sequence to act as an override to the existing controls. The algorithm was executed at 10 minute intervals as a compromise between encompassing changing office conditions and minimizing building automation system network traffic. This allowed for three state and input updates between each model time step and prevented controllers from having overloaded network buffers.

$$\Delta U = (\varphi^T \varphi + R)^{-1} \varphi^T (SP - T) = H(SP - T) \quad (5-3)$$

$$\varphi = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ A & 1 & 0 & 0 & 0 & 0 \\ A^2 & A & 1 & 0 & 0 & 0 \\ A^3 & A^2 & A & 1 & 0 & 0 \\ A^4 & A^3 & A^2 & A & 1 & 0 \\ A^5 & A^4 & A^3 & A^2 & A & 1 \end{bmatrix} \quad (5-4)$$

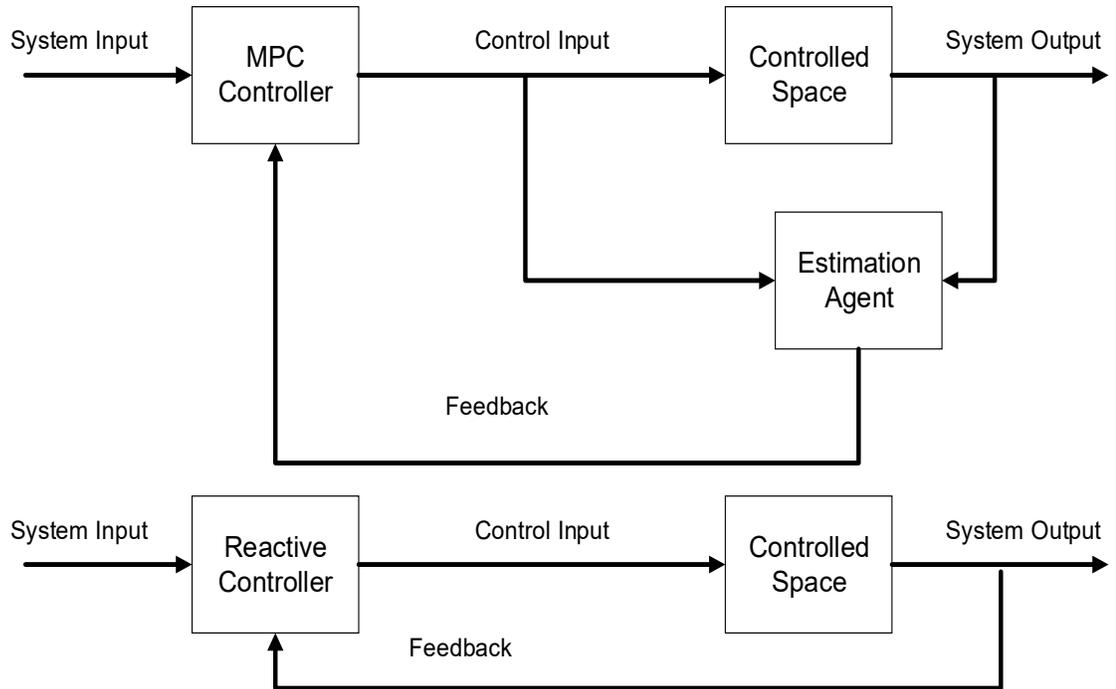
```

IF (Indoor air temperature < setpoint) AND (sum(ΔU(t + 1): ΔU(t + 3)) < 0)
    MPC output = 1;
ELSE
    MPC output = 0;
END

```

**Figure 5-4: MPC override classification pseudocode for heating season. A 0 indicates the heating equipment should be overridden and a 1 indicates it should continue using reactive control.**

The MPC algorithm has the added complexity of including an estimation agent that required significant passing of information between the inputs and outputs of the controlled space and the MPC controller itself. This differs from the complexity of the reactive controller and is illustrated with block diagrams in Figure 5-5.



**Figure 5-5: MPC (top) and reactive control (bottom) block diagrams**

The control agent was intended to be modular, and to interface with any BACnet building automation system (BAS). In the studied implementation, drivers were used to transfer data between the MATLAB simulation environment and the automation system using data stored in comma delimited files. The same principles could be applied to systems operating on other protocols, but the information gateway operation would differ. The

output of the MPC algorithm was selected to be binary for simplicity of implementation.

### 5.3.2.2 Classification Learning for Rule Extraction

Control decisions from the MPC algorithm and the respective inputs (and input predictions) were recorded and utilized as the basis for extracting rules to construct a Boolean decision tree. The necessary inputs and thresholds were found using the MATLAB classification toolbox classification tree functions (Duin, Juszczak et al. 2000). The binary output of the MPC algorithm allowed the use of a classification tree rather than a regression tree to minimize the presence of arbitrary thresholds and improve readability for operators and technicians. The exact method was utilized to fit the classification tree, exploring all possible combinations of inputs. The minimum leaf size was maintained at 1% of the data set to minimize the loss of prediction accuracy of multiple splits, while maintaining a reasonable number of conditions. Prediction accuracy with respect to the MPC algorithm was assessed using confusion matrices, and the importance of minimizing false positives (e.g., activating HVAC equipment unnecessarily) was highlighted.

The decision tree developed for the office studied is shown in Figure 5-7, where  $T(t + 1)$  is the predicted air temperature ( $^{\circ}\text{C}$ ) for the next time step using the linear parameter model,  $u_2(t)$  is the illuminance (lux), and  $u_3(t)$  is the outdoor air temperature ( $^{\circ}\text{C}$ ). These thresholds were determined by the classification algorithm using the exact method, where all possible classifications are explored and the version with the lowest prediction error is selected (Maji, Berg et al. 2008). Other similar forms of decisions that included temperature predictions from the edge of the control horizon (three time steps

ahead) were generated but diminishing returns on classification accuracy were found for the required connectivity of weather predictions. Practical commercial building implementation of sMPC would require a dedicated embedded database object with threshold variables and real-time inputs as shown in the pseudocode of Figure 5-6. Each threshold corresponds to a value from Figure 5-7, where the inputs are predicted air temperature one time step ahead (IATp1), ceiling illuminance (Lux), outdoor air temperature (OAT), and the output (out) is binary. This decision tree utilizes no weather prediction data and derives prediction information from only the linear model parameters and current time step inputs and states. Slight improvements in classification were observed by utilizing indoor air temperature predictions from the edge of the control horizon using weather predictions, but the incremental benefit of ~1% classification accuracy was outweighed by the simplicity of the rules selected.

Room air temperature setpoint is a notable absence in the resulting decision tree. This occurred because the setpoint was 21°C and during the entire testing period (and did not change), and there was no setback in the initial reactive control algorithm. Setback would necessitate a different set of rules for the unoccupied period but could yield a similar structure. The setpoint may not be required in the decision tree if the conditioned space is a single occupant office because a single occupied setpoint would satisfy preferences. However, this would require seasonal retraining of the parameters and thresholds, and an assessment of input validity and model form.

The extracted rules in the decision tree generated in this study are based on zone air temperature predictions, solar input, and outdoor air temperature. This follows a logical

sequence, where a low predicted temperature will activate the heating system, a high predicted temperature will deactivate the heating system, and a predicted temperature within a specific band will activate sets of conditions based on solar input and outdoor air temperature. If a solar input is observed within this band (indicating daytime), the heating will not activate. Meanwhile, if solar gains are low (indicating nighttime) and outdoor air temperature is also low, the heating system will activate. The intuitive nature of the previous statements highlights the ease of understanding of the proposed rule extraction approach in this study, where unique paths can be intuitively analyzed.

```

if IATp1 < 20.866 then
  if Lux < 30.139 then
    out = 1
  else if IATp1 < 20.711 then
    out = 1
  else
    out = 0
  End If
End If
else if IATp1 < 20.953 then
  if Lux < 17.276 then
    if OAT < -4.075 then
      out = 0
    else
      out = 1
    End If
  else
    out = 0
  End If
else
  out = 0
End If
End If

```

**Figure 5-6: sMPC classification pseudocode implemented in conditioned space for heating season**

Qualitatively examining the decision tree in Figure 5-7 yields several intuitive inferences from the building conditions. The three predictors ( $T(t + 1)$ ,  $u_2(t)$ , and  $u_3(t)$ ) essentially align with thermal mass, time of day, and the balance temperature of the building respectively. A product of the state-space model, the IAT prediction defines four distinct temperature regimes, where the two constrained regimes have supplementary conditions dependent on other inputs. Predicted temperatures at the setpoint and predicted temperatures more than  $0.29^{\circ}\text{C}$  below the setpoint always result in suppressed and enabled heating respectively. The illuminance threshold is very low, suggesting it is an indication of day or night rather than a precise level of solar radiation by proxy. The outdoor air temperature aligns well with the point at which, even with solar gains, the office would still drop in temperature. This outdoor air temperature threshold may also be correlated with thermal mass and other terms of the state-space model, however there is not enough data present from this single office analysis to validate this postulation.

Different rules would need to be extracted for different spaces based on thermal characteristics and user preferences. However, similar tree structures and model forms could be obtained and adjusted using a local optimization approach for similar space types. The decision tree for the space studied is presented in Figure 5-7, where a 0 endpoint indicates that the reactive control should be suppressed, and a 1 indicates that the reactive control should remain in normal operation. The training and validation confusion matrices are presented in Figure 5-8 and Figure 5-9 indicating errors of 6.8% and 6.4% respectively. It is clear in the graphical version of the decision tree that when the available inputs are reduced, the classification tree develops ranges for certain decisions within the hierarchy.

These conditions can clearly be verified by adjusting inputs, thus improving transparency to the operator or technician. Over-fit to the MPC conditions was avoided by utilizing one week as a training set and another week as a cross-validation data set, where similar accuracy to the training set was observed.

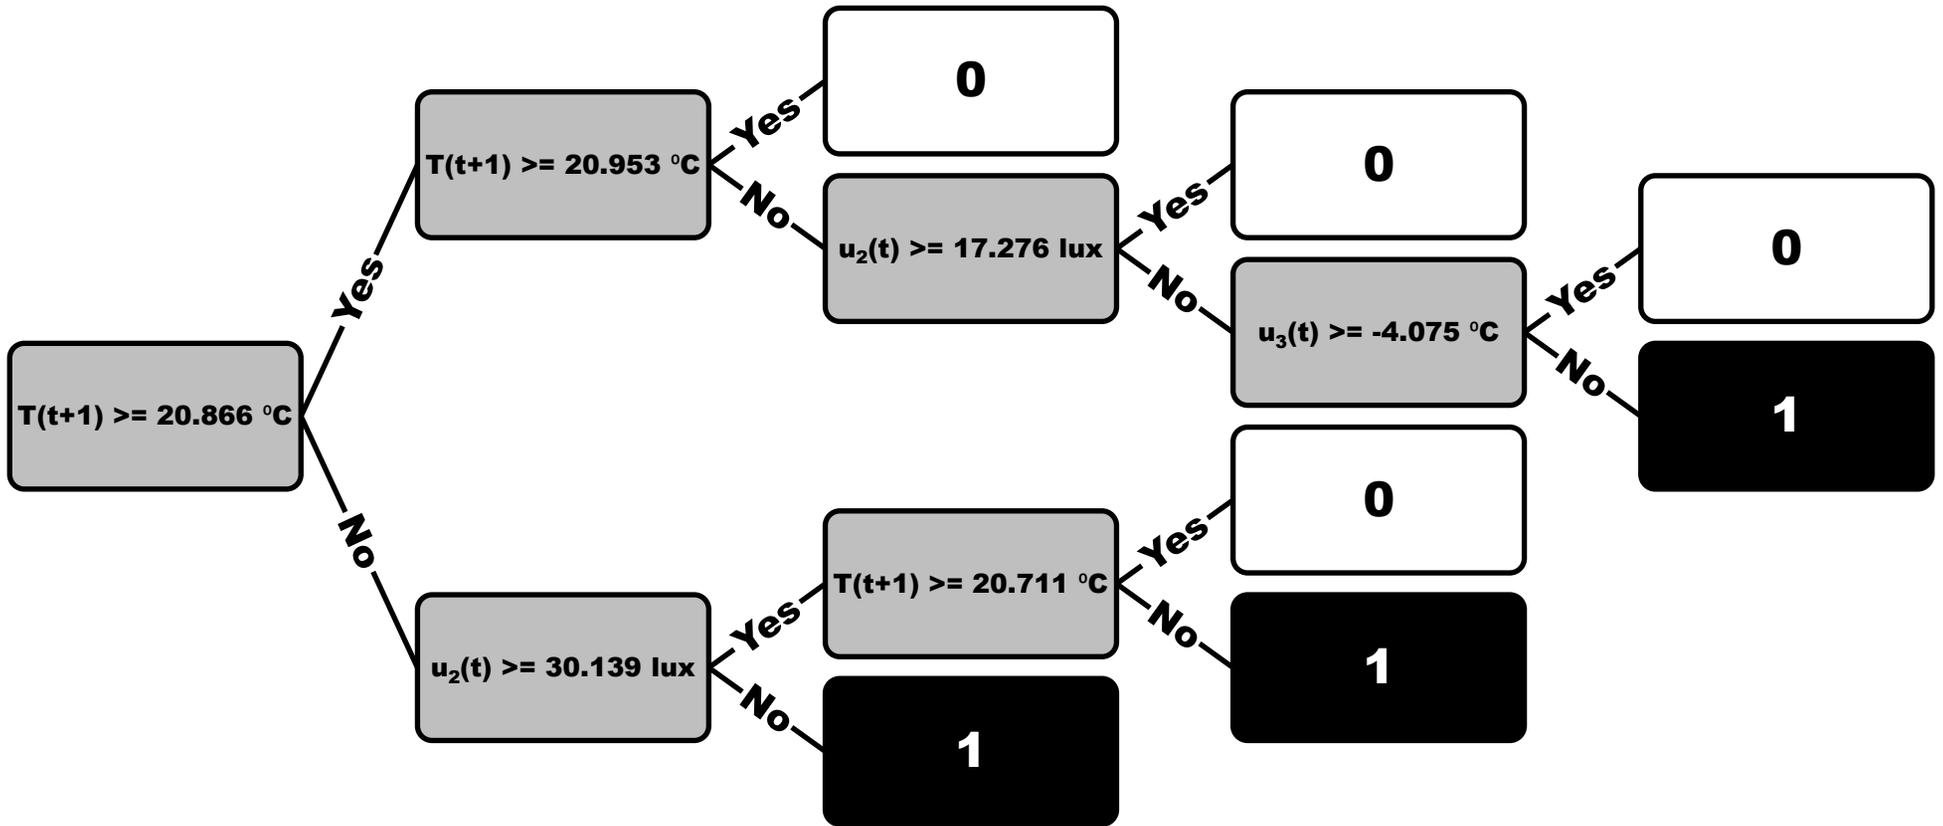


Figure 5-7: Sample classification tree for simulated predictive control output. A 0 indicates the heating equipment should be overridden and a 1 indicates it should continue using reactive control.

Output Class	0	172 45.1%	6 1.6%
	1	20 5.2%	183 48.0%
	Errors 6.8%	0	1
	Target Class		

**Figure 5-8: Confusion matrix for simulated predictive control output training. The target class and output class correspond with the MPC and sMPC outputs, respectively. Classification error is indicated on the diagonal of the matrix.**

Output Class	0	228 63.6%	14 3.9%
	1	9 2.5%	109 30.3%
	Errors 6.4%	0	1
	Target Class		

**Figure 5-9: Confusion matrix for simulated output validation. The target class and output class correspond with the MPC and sMPC outputs, respectively. Classification error is indicated on the diagonal of the matrix.**

### 5.3.2.3 Energy and Comfort Quantification

This work attempts to quantify energy savings using flow meters and thermistors measuring hydronic heating system properties, while allowing for the adaptation of simple virtual sensors to estimate flows from valve position for spaces that do not have direct sensing of flows and temperatures. Some success has been found using empirical models to estimate flow using valve position as a proxy (Song, Joo et al. 2012). The radiant panel inlet temperature, outlet temperature, and hot water flow rate were used to estimate the heating energy cost at each time step. Temperatures were obtained by pipe surface mounted thermistors (beneath the pipe insulation), and the flow rate was measured using an ultrasonic flow meter. The average observed values, tabulated thermo-physical properties (Incropera, Lavine et al. 2007), and manufacturer reported bias errors are provided in Table 5-3. These quantities were used to estimate energy consumption at individual time step resolution with Equation 5-5. Equation 5-5 assumes the system is at steady state during each time step, neglecting transient effects. The transient period for the radiant panels was observed to be roughly 15 minutes, which even during periods of frequent heating activity oscillation, accounted for less than 9.8% of daily total energy consumption. Quantifying useful energy delivered to occupied space (and discomfort due to asymmetry) is not within the scope of this study, and has been covered extensively in the literature and the ASHRAE Handbook of Fundamentals (ASHRAE 2009). Utility rates and time of use billing are not accounted for in this study due to significant regional variability.

**Table 5-3: Input average value over the testing period and manufacturer-provided instrument errors**

	Average Value	Instrument Error
$V$ (m <sup>3</sup> /s)	$8.3 \times 10^{-6}$	$8.3 \times 10^{-8}$
$\rho$ (kg/m <sup>3</sup> )	1,000	0.0025
$c_p$ (J/kg-K)	4,200	5.0
$\Delta T$ (°C)	10	0.4

$$\Delta U = \rho V C_p \Delta T \quad (5-5)$$

The sums of the heating energy for each time step were compared to estimate the energy savings over the five different weeks of testing. Plug and lighting loads were excluded from this analysis because they were insignificant in comparison to the thermal loads. The variability of casual gains was also excluded because occupancy was assumed to be the proxy, and the passive infrared occupancy sensors do not allow for quantification of occupants, only a binary state.

Instrumentation uncertainty was quantified as the root square of the bias error and precision error. The bias error and instrumentation uncertainty were calculated using Equation 5-6 and Equation 5-7. The bias error was found using the separation of variables approach of Equation 5-6, where Equation 5-5 was differentiated for each variable. Typically, bias error is much larger than precision error (Wang, Duan et al. 2016). This was the case for the radiant panel instrumentation, where the bias error was 14.4 W and the precision error was 1.2 W. Using Equation 5-7, the instrumentational uncertainty was found to be 14.5 W. Over the entire observation period, the uncertainty ranged from 3.6% to 9.1%, with an average of 4.5%.

$$\Delta s_b^2 = \sum_{i=1}^n \left( \frac{\partial U}{\partial x} \right)^2 \Delta x_i^2 \quad (5-6)$$

$$\Delta s = \sqrt{\Delta s_b^2 + \Delta s_p^2} \quad (5-7)$$

## 5.4 Results

This section presents the in-situ testing results of the MPC and sMPC approaches described in the methodology section. The weather prediction accuracy, initial reactive control case, online MPC case, online sMPC case, and energy quantification results are presented, and observations and anomalies are discussed.

### 5.4.1 Weather Forecast Accuracy

Forecast accuracy was verified during initial reactive control monitoring. An example day is provided from the second MPC testing week in Figure 5-10. Error between predicted and observed weather conditions was previously quantified in the methodology section, and all forecasted inputs were determined to have uncertainty of  $\leq 0.06^\circ\text{C}$  for all inputs of the state-space model for each time step. In Figure 5-10 the absence of bias in the weather predictions is observed. Additionally, the conservative overestimation of occupancy is demonstrated. It was assumed that the discomfort caused to the occupant by underestimating their occupancy would be a larger detriment to thermal prediction and energy consumption than the potential for oversupply of heat during absence. This assumption would also be of benefit with setback implementation.

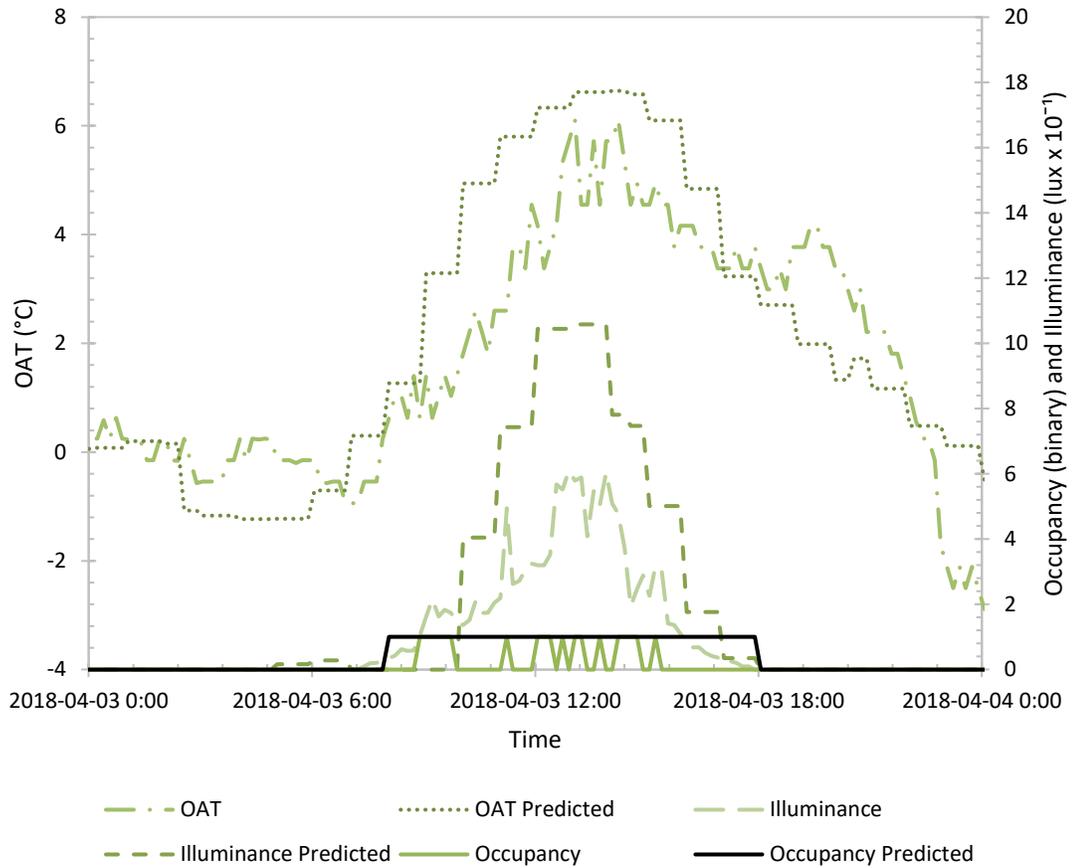


Figure 5-10: Input prediction accuracy on April 3<sup>rd</sup>, 2018 for OAT, illuminance, and occupancy.

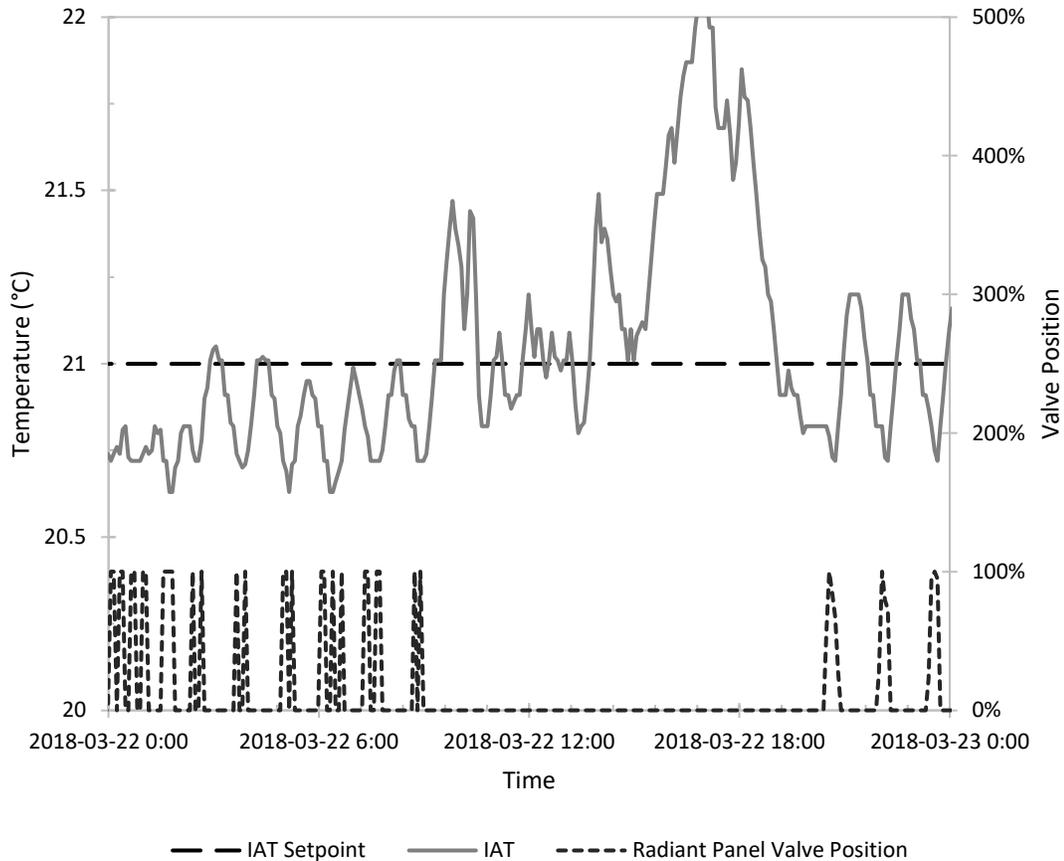
#### 5.4.2 Reactive Control Baseline

A sample day of reactive control with indoor air temperature, indoor air temperature setpoint, and radiant panel valve position for the reactive control case is provided in Figure 5-11. Figure 5-11 demonstrates that a setpoint deviation of nearly 0.5°C occurs within three hours of heating activation, just before 9:00 on March 22<sup>nd</sup>. This overheating could have been predicted with the prediction horizon of the MPC algorithm and prevented both energy use and overheating simultaneously. In similar situations, allowing the building to

“free float” with the knowledge that the energy balance is neutral or positive over the control horizon can be advantageous.

Cooling was also provided by the air side system to modulate the temperature during overheating conditions and to provide outdoor air. Note that the cooling is a result of the supply air (which is a combination of the return air and supplied air) being cooler than the room air temperature rather than active cooling with a cooling coil. Despite the availability of cool air to reduce the room air temperature, parameters of the reactive PI control loop would often reset and caused unpredictable system response (e.g., periods of very high or low loop bias). This is visible around 18:00 in Figure 5-11, where the indoor air temperature climbs despite being well past the setpoint and PI loop deadband of 0.5°C. The PI settings were not modified in this study and were set by a controls contractor during initial building controls programming. The proportional band was 1.7°C and the integral rate was 1% per minute.

Further weekly averaged input and energy information for the baseline case is provided in Table 5-4.

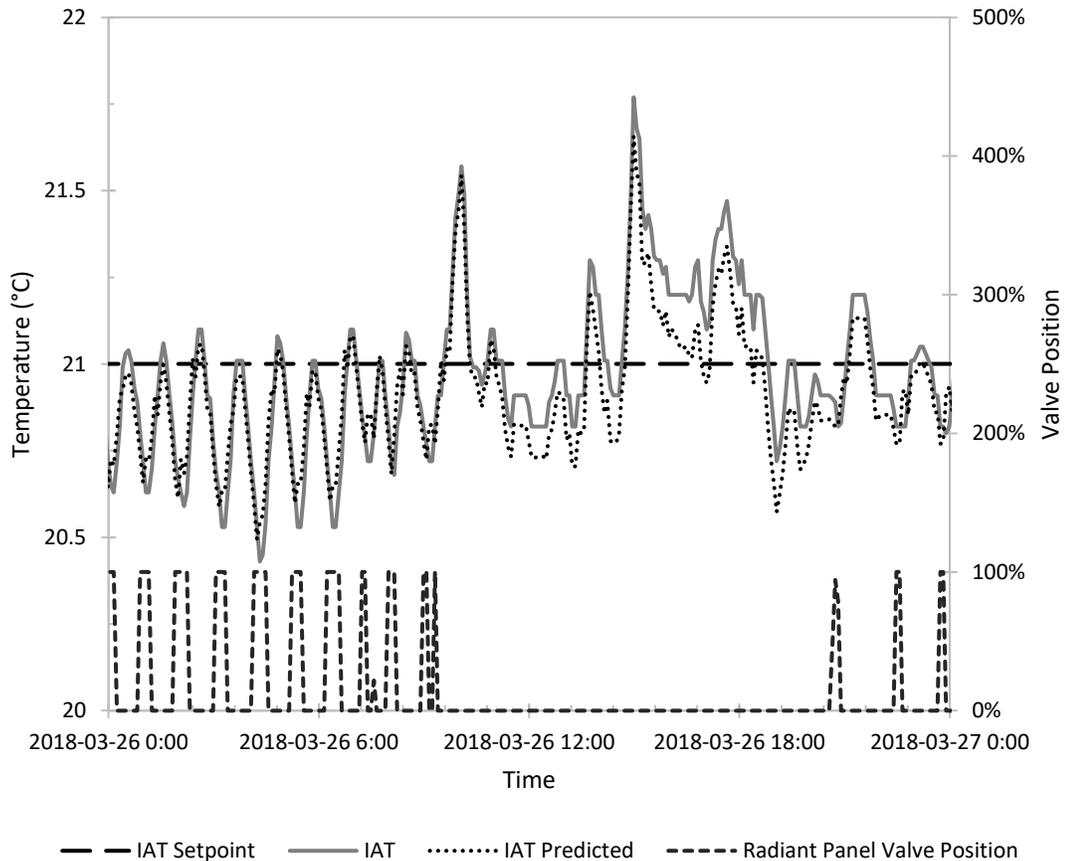


**Figure 5-11: Typical temperature predictions and control outputs for reactive control baseline on March 22<sup>nd</sup>, 2018, where the OAT was on average 1.8°C**

### 5.4.3 Online MPC

The MPC algorithm was implemented after two weeks of office sensor data was gathered to train the inverse state-space model. In the MPC case, indoor air temperature overshoot was both less frequent and less severe. Indoor air temperature overshoot was observed to be impacted by solar gains that could not be estimated or predicted by the reactive control algorithm. This can be seen to be mitigated by the MPC algorithm in Figure 5-12, where the setpoint deviation remains significantly lower and overshoot is less common.

Additionally, predicted thermal loads are considered and heating equipment is stopped early to utilize the thermal mass. Large disturbances can still cause sub-optimal heating operation and are usually due to solar or occupancy fluctuations. The occupant and solar disturbance effects can be attributed to the non-linearity of the respective inputs when utilized in a linear model. Solar energy is often modelled with higher order relationships using temperature, and occupancy in this case is binary. Both are cases that highlight the limitations of using a simplified linear time-invariant model but demonstrate that these limitations do not inhibit potential energy savings.



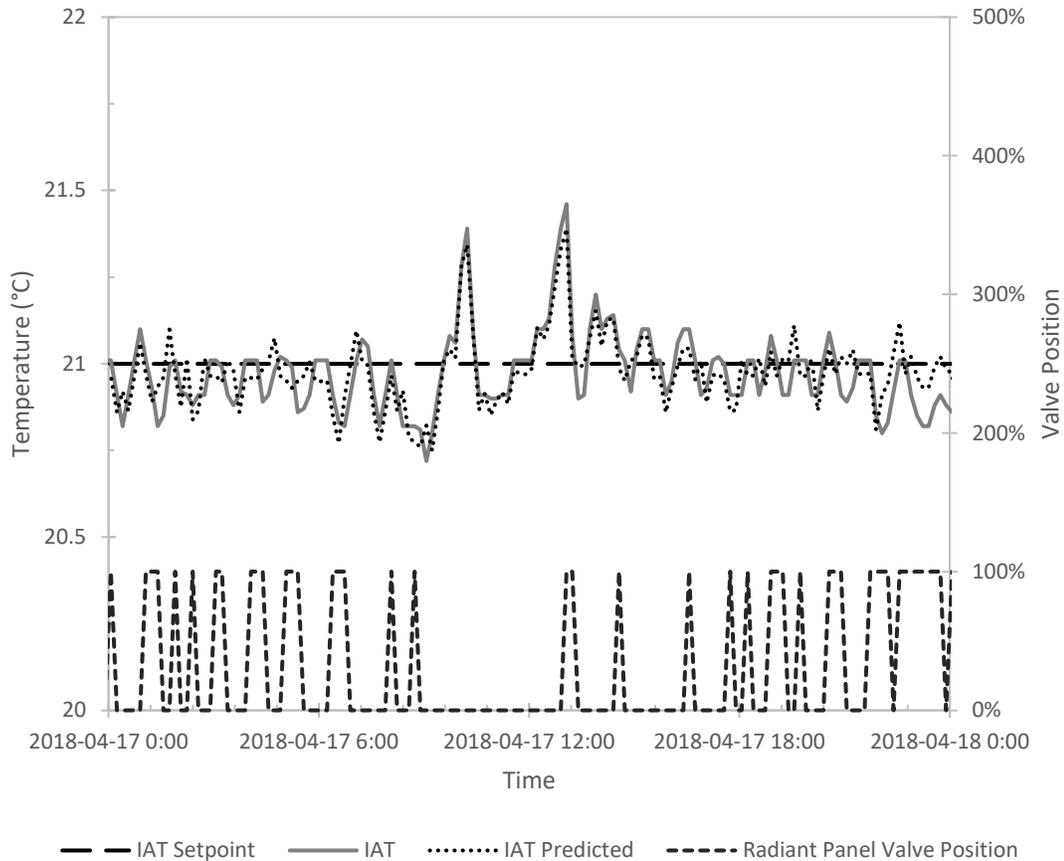
**Figure 5-12: Typical temperature predictions and control outputs for MPC case on March 26<sup>th</sup>, 2018, where the OAT was on average -2.0°C**

#### 5.4.4 Online sMPC

The prior MPC testing provided one week of training and one week of validation data for the rule extraction approach to develop thresholds for different inputs with respect to sMPC control decisions. The resulting decision tree utilizes one state, five inputs, six linear time-invariant model parameters, and six thresholds. More data will be required in future work to determine if a 12 parameter (six linear time-invariant parameter and six threshold) decision tree is appropriate for more archetypal spaces. However, if a similar model form is appropriate, this approach could be coupled with reinforcement learning to circumvent the need for computationally and mathematically cumbersome forms of MPC. Future work should include a statistically significant sample size of office spaces to determine the range and appropriate discretization granularity for the optimization problem to identify the feasibility of optimizing the sample space within a typical building automation controller. The single office space studied in this paper demonstrates a simplified approach but does not provide enough information to make statements about generic offices.

Operation of the heating system under the sMPC algorithm exhibited much closer control to the setpoint as seen in Figure 5-13. This was a result of precise static thresholds of predicted air temperature that circumvented any hysteresis that would occur in either the reactive control loop or MPC control horizon. An unintended side-effect of using classification learning is the occasional erratic heating system operation (or lack thereof). This will be inherently present in all statistical modelling methods that utilize deterministic models for stochastic processes but can be minimized with enough data and by avoiding overfit through appropriate model validation.

Another observation of the sMPC case was that changes in input conditions outside of the training range have a large effect. The passive infrared occupant sensors in the office studied provide only a binary indicator of occupancy and cannot determine the quantity of occupants. Anecdotally, high occupancy was observed to cause overheating in the sMPC case because the occupancy input is only a small component in the rules (one of six inputs to the predicted temperature). While occupancy triggered overheating was observed primarily in the sMPC case, it would be an issue in the MPC case as well and is a limitation of using binary occupancy sensors when modelling a multi-occupant space. Future work should attempt to utilize a different combination of sensors to quantify the number of occupants in the space for better thermal response.



**Figure 5-13: Typical temperature predictions and control outputs for sMPC case on April 17<sup>th</sup>, 2018, where the OAT was on average -1.2°C**

#### 5.4.5 Seasonal Energy Savings and Thermal Comfort

Heating and cooling energy results are provided in Figure 5-14 with weekly average input values to provide context. Heating energy use by the radiant heating panels (the only source of heat in the room) was estimated using ultrasonic flow meter, inlet temperature, and outlet temperature data with Equation 4-2, and the cooling energy use was estimated using the AHU SAT, VAV airflow rate, and indoor air temperature as per Equation 6-5. The use of a singular test space for this study was a limitation regarding the environmental input

effects on different weeks of evaluation and will be addressed in future work. As a result, this study is not intended to be a truly quantitative assessment of the energy consumption of the different cases. The MPC and sMPC cases both saw reduced heating energy consumption with respect to the base case. Average input conditions for each case week have been reported to provide context to how some weeks may have been affected by weather and/or occupant conditions. Notably, the reactive control week with low energy consumption also saw high solar gains (albeit also with a low OAT) likely from open blinds. For much of the rest of the testing period the blinds remained closed due to occupant preference.

When comparing MPC week 1 to MPC week 2, the similarity of average input conditions does not reveal the entire context of the energy results. Less cooling was used in MPC week 2 than MPC week 1 despite a higher heating energy, indicating the heating energy was not wasted. This difference between MPC week 1 and 2 could be due to a difference in the transfer of heating from the space to the corridor. It was verified that the door was open for 6.2% more time during the second week of MPC testing. Further, it could also be attributed to a larger quantity of occupants.

Cooling was not accounted for in the optimization process because heating is the focus of this iteration of the algorithm, and overheating was assumed to be avoidable due to the MPC handling of thermal mass. Large solar gains were observed to cause undesired overheating in the office (as indicated by the required cooling energy), and an extended control horizon could better adapt to these conditions. The 1.5 hour control horizon was chosen to be occupant-centric to facilitate passing information between occupant

prediction sub-routines and the building thermal prediction algorithm. Further, a longer control horizon would increase the forecast dependence of the algorithm, including the accompanying uncertainty. Table 5-4 provides a summary of the setpoint deviation, setpoint bias, model inputs, cooling energy, heating consumption, and energy savings for all five weeks of testing. The high setpoint deviation of Reactive Control Week 1 provided an energy usage advantage at the potential expense of occupant comfort. As a result, it was used in Reactive Control Total to compare with MPC total. The individual weeks of MPC and sMPC were compared to Reactive Control Week 2 as the equivalent run. The limitations of using a single space to test in series (rather than in parallel) are apparent, but the variation of weather and occupant conditions were limited throughout the testing period. Further, the averaging of the two reactive control and MPC weeks improves averaging of the input values. The sMPC approach was observed to achieve 94% of the heating energy savings of full MPC which aligns with the findings of previous related work in the literature (May-Ostendorp, Henze et al. 2011). It is assumed that local energy rates can be applied (in a blended format for time of use regions) to estimate regional applicability.

An average indoor air temperature of 21.12°C, with standard deviation of 0.55°C and a minimum of 19.38°C was observed over all testing conditions. Further, no setpoint overrides were requested, and anecdotally the occupants preferred colder air temperatures. Mean radiant temperature (and operative temperature) were not obtained, but under the assumption of equal operative temperature to air temperature for the radiant heating system, the indoor air temperatures fall within or near ASHRAE 55 criteria.

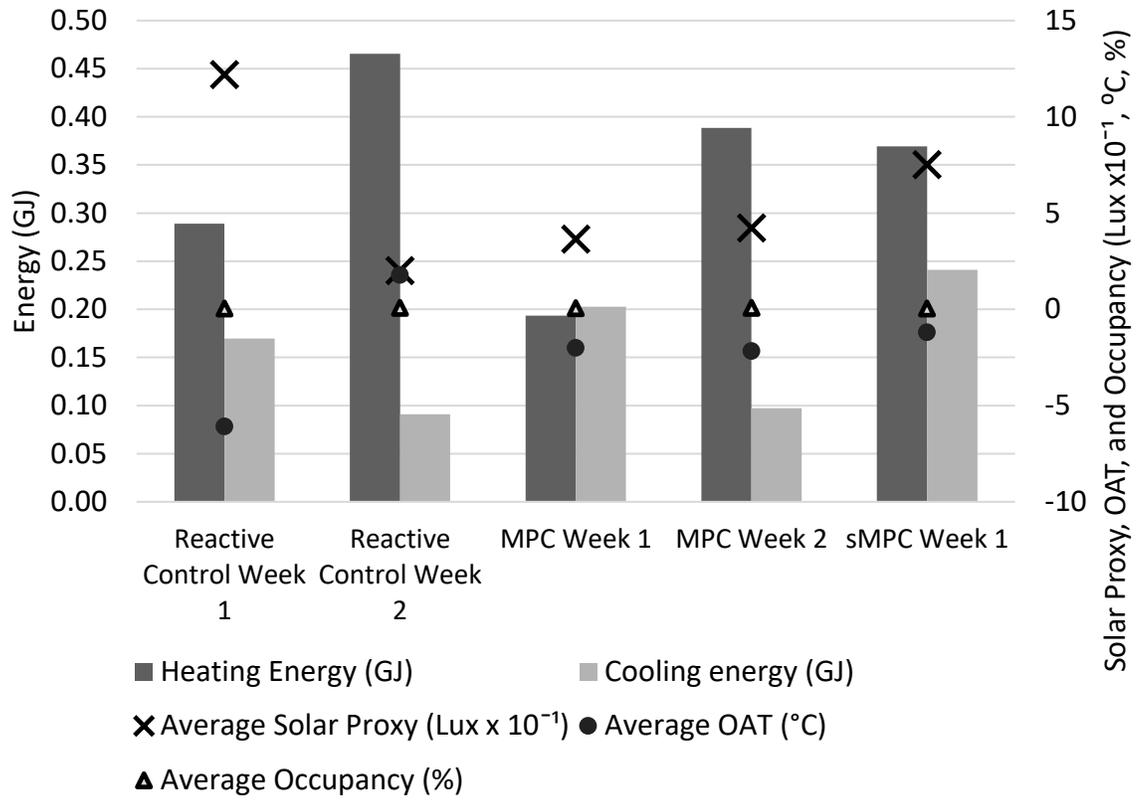


Figure 5-14: Weekly heating and cooling energy use for the testing period. Average solar proxy, OAT, and occupancy for each week is also indicated.

**Table 5-4: Weekly energy consumption, average solar proxy, average OAT, and average occupancy values for each algorithm**

Mode	Average Setpoint Deviation (°C)	Average Setpoint Bias (°C)	Average IAT (°C)	Average Solar Proxy (Lux)	Average VAV Proxy (kW)	Average OAT (°C)	Occ rate (%)	Cooling Energy (GJ)	Heating Energy (GJ)	Heating energy savings compared to Equivalent run (GJ)	Savings compared to Equivalent run (%)
Reactive Control Total	0.57	0.46	21.5	86.6	2.5	-3.4	6%	0.19	0.75		
Reactive Control Week 1	0.77	0.73	21.7	122.0	2.7	-6.1	5%	0.17	0.29		
Reactive Control Week 2	0.21	-0.04	21.0	20.2	2.1	1.8	7%	0.09	0.47		
MPC Total	0.23	-0.11	20.9	38.5	2.5	-2.1	7%	0.22	0.58	0.17	22.6%
MPC Week 1	0.22	-0.08	20.9	36.5	2.7	-2.0	6%	0.20	0.19	0.28	59.6%
MPC Week 2	0.26	-0.16	20.8	42.2	2.3	-2.2	8%	0.10	0.39	0.08	16.5%
RE Week 1	0.11	-0.02	21.0	75.1	2.6	-1.2	4%	0.24	0.37	0.10	21.3%
Average	0.34	0.11	21.1	60.2	2.5	-2.2	6%	0.17	0.36		
Range	0.65	0.89	0.89	101.8	0.6	7.9	4%	0.15	0.27		

## 5.5 Summary

In this chapter two model predictive control approaches were assessed using a single institutional multi-occupant office. A single node linear time-invariant state-space thermal model with five sensor inputs was developed to characterize the temperature response of the space for the MPC cost function. The first approach generated a binary override for the BAS based on the sum of the MPC predicted output over the control horizon and the setpoint deviation. The second approach used classification learning to extract a decision tree to predict the MPC binary override from the sensor information and weather predictions for the office.

It was found that the MPC implemented in this study achieved heating energy savings of 22.6% with respect to the reactive control installed during construction. Further, the sMPC approach was observed to achieve 94% of the heating energy savings of the MPC approach (21.3% with respect to reactive control) using no forecast information or matrix inverse operations.

Thermal comfort was assessed by occupant feedback via button presses to adjust setpoints and air temperature. No overrides were made over the entire testing period, and an average indoor air temperature of 21.12°C, with standard deviation of 0.55°C and a minimum of 19.38°C was observed over all testing conditions.

The rule extraction approach provided a simplified predictive control scheme that could save heating energy while maintaining (or improving) occupant comfort. Setpoint deviation was lower with the sMPC algorithm than any of the other algorithms tested due to static temperature control thresholds, likely an artifact of the small deadband of the base

control. The simplicity of the sMPC algorithm offers a degree of transparency unachievable with most data-driven predictive control algorithms. The developed decision tree contained six thresholds and required six linear time-invariant parameters from the original state-space model. Various paths and thresholds could be identified and justified using the simplified decision tree, and intuitive actions for the zone level BAS were identified. The sMPC algorithm provides a feasible retrofit solution for existing buildings where older control hardware, mathematical operator limitations in programming, and congested networks would make detailed MPC application difficult or impossible without extensive hardware upgrades.

Despite the observed energy savings, the use of a single office without parallel testing of competing algorithms generates a need for further verification. Moreover, the building-scale impact on energy (e.g., for boilers, chillers, and fans) should be assessed in future research. The classification tree utilized for this office could not be verified as a generic solution, and the logical sequence could vary from case to case. However, the proposed workflow and approach proposed in this paper could expedite future work in larger office sample sizes that can facilitate testing and tuning in parallel with a control group. The model form in this study worked well for one office, but the assessment of different model forms over a larger set of offices would be of benefit to improving the temperature prediction accuracy. Future work will include the study of a larger quantity of offices in a parallel testing format. Under these conditions, most of the results uncertainty will be correlated to occupancy rather than weather conditions. A longer period of study with algorithms swapping between the control and test groups would also provide more

consistent results. Preliminary results are promising, and when applied to a larger dataset, could transform the commercial MPC problem into a reinforcement learning local optimization problem with rapid commercialization potential.

## Chapter 6: Field Implementation of Rule Extraction Supervisory Control

This chapter is based on the supplementary material of:

Jayson Bursill, William O'Brien, Ian Beausoleil-Morrison. (In review). Multi-zone field study of rule extraction control to simplify implementation of predictive control to reduce building energy use. *Energy and Buildings*, Accepted for publication.

Model-based predictive control (MPC) has been promoted as a software-based energy saving approach for buildings since the 2000s. However, there has not been significant industry adoption of MPC using current approaches. This paper proposes and tests an approach to MPC using rule extraction (RE) that can be easily implemented in building controllers to override sub-optimal control. The resulting decision trees could be implemented by building control programmers to save energy when ambient conditions are predicted to satisfy the thermal requirements of the spaces. A detailed MPC algorithm using inverse models was implemented in 27 rooms of an institutional building to provide data for a classification learning approach. Cooling and heating season decision trees were generated based on the inputs and outputs of the detailed MPC algorithm. An ensemble and sample randomization were used to generalize the trees across rooms and prevent overfitting to individual rooms. Cooling season MPC and RE energy savings were 42% and 27% respectively and heating season MPC and RE energy savings were 18% and 33% respectively. Both algorithms also reduced the difference between the temperature setpoint

and measured indoor air temperature compared to the reactive controls.

## **6.1 Introduction**

Model-based predictive control (MPC) has demonstrated the potential to save up to 50% of heating and cooling energy in commercial buildings by preventing overheating and overcooling through the use of thermal mass (Tzempelikos and Shen 2013, Dong, O'Neill et al. 2014). The goal of this paper is to provide a simplified and experimentally evaluated approach to MPC utilizing rule extraction (RE) from a detailed MPC algorithm to generate a Boolean logic-based simplified algorithm. The simplified algorithm can be mapped to thresholds of physical inputs and implemented by building automation system (BAS) contractors.

Numerous studies have been conducted to evaluate the energy saving potential of MPC. Many of these were simulation-based studies rather than physical experiments, and therefore these did not examine the impact of disturbances and uncontrolled conditions (Braun and Chaturvedi 2002, Henze, Felsmann et al. 2004, Oldewurtel, Parisio et al. 2012, Coffey 2013, Afram and Janabi-Sharifi 2015). At least 11 experimental and in-situ studies have been conducted with promising results reflecting the previous claim of up to 50% energy savings (Roth, Llana et al. 2006, Privara, Široký et al. 2011, Široký, Oldewurtel et al. 2011, Afram and Janabi-Sharifi 2014, Shaikh, Nor et al. 2014, Touretzky and Baldea 2014, Goyal, Barooah et al. 2015, Gunay, O'Brien et al. 2016, Gunay, O'Brien et al. 2016, Hilliard, Swan et al. 2016, Thieblemont, Haghghat et al. 2017). Despite this effort, there has still been minimal commercial adoption of predictive controls (Killian and Kozek 2016).

Applied, simplified approaches are sparse within the experimental subset of the literature on predictive controls. Gunay et al. adapted a detailed MPC algorithm for recursive parameter estimation within a BACnet BAS controller environment (Gunay 2016). Noted limitations of most commercial BAS controllers in Gunay's work are the inability to perform matrix operations and low physical memory for program code storage and intermediate variable storage. It is also stated that the hardware limitations of BAS controllers significantly limit the algorithm complexity that can be programmed directly into controller databases. Remote reading and writing of controller commands to a computer that can perform detailed MPC calculations can also pose operational and safety risks due to the need for a constant network connection.

Classification learning has been used to identify combinations of physical input thresholds for control outcomes in the form of RE from detailed MPC (May-Ostendorp, Henze et al. 2011). May-Ostendorp's approach was applied to a mixed-mode building simulation, and later extended to a test in a controlled and instrumented climate controlled chamber (May-Ostendorp 2012). Bursill et al. adapted this approach for experimental application to the existing BAS of a single multi-occupant office, but cited the need for a larger set of offices to verify the effectiveness of simplified MPC at reducing energy use (Bursill, O'Brien et al. 2018). The importance of large quantities of data to validate classification approaches is further highlighted by Coffey et al. (2013), where simulation was used to develop optimal control rules for a modelled building stock in the form of a lookup table.

Ambiguity of the practical energy savings of both MPC and RE control in real buildings is a concern when advocating for predictive controls. Few mechanically ventilated facilities demonstrate the control and measurement capability to implement real-time MPC and validate the energy use of multiple thermal zones. Herrick Labs at Purdue University has extensive instrumentation to monitor occupants and the air supply system, with a focus on occupant and indoor environment quality studies (Deng and Chen 2019). The National Renewable Energy Laboratory's Research Support Facility is another commercial building with in-situ research applications, largely in occupant comfort and net-zero energy (Lobato, Pless et al. 2011). The Canadian National Research Council's Canadian Centre for Housing Technology research facility contains several highly instrumented residential buildings that can measure supplied energy from heating and cooling systems (Zirnhelt and Richman 2015). FLEXLAB at Lawrence Berkeley Laboratory can evaluate the energy efficiency of major building systems as an integrated system under real-world conditions (McNeil, Kohler et al. 2017). The Health Sciences Building at Carleton University contains 25 offices with the ability to measure supplied terminal hydronic heating, lighting, and plug loads at the room level (Bursill, O'Brien et al. 2019).

A gap in the literature is addressed in this paper by combining advanced control algorithms (MPC), simplified predictive approaches (RE), in-situ application and testing, and room-level measurement into a comprehensive study. Prior to this study there has not been a documented application of RE control to multiple rooms within an actively occupied building. The novelty of this work is in both the field implementation approach and the

scale of the field implementation itself. In contrast to previous studies which typically looked at entire buildings or single rooms (Miriél, Serres et al. 2002, Rahimi and Sabernaemi 2010, Bălan, Cooper et al. 2011, Candanedo and Athienitis 2011, Široký, Oldewurtel et al. 2011, Bengea, Kelman et al. 2012, Karava, Athienitis et al. 2012, Ma, Kelman et al. 2012, Dermardiros, Chen et al. 2015, Li, Wen et al. 2016, Bursill, O'Brien et al. 2019), this study used a much larger sample of 27 offices, thus increasing confidence of the results.

First, this paper describes the test facility and instrumentation uncertainty. Second, it develops a methodology for zone-level MPC and adapt it to zone-level RE control. Third, it quantifies the energy use savings of each algorithm when applied at the zone-level using in-situ approaches in an occupied building. Finally, it provides an evaluation of the relative effectiveness of RE compared to detailed MPC and the original reactive control of the building.

## **6.2 Methodology**

This chapter builds upon the findings and recommendations of the lab implementation in Chapter 5. Resultantly, some modifications were made for field implementation. The most notable modifications include: removal of the embedded indoor air temperature prediction model, additional RE testing time (compared to sMPC), parallel control (untested reactive control) and test (MPC or RE control) groups, inclusion of both the heating and cooling season, modification of the override algorithm for zone control, improved automation of reading and writing between the detailed algorithm and the BAS.

The following sections of this methodology describe the experimental facility, measurement uncertainty, detailed MPC algorithm, workflow to generate simplified RE decision trees to emulate the detailed MPC, and the formulation of energy use to compare the performance of the respective algorithms. A high-level overview of the procedure for implementing the MPC and RE control for field implementation is provided in Figure 6-1.

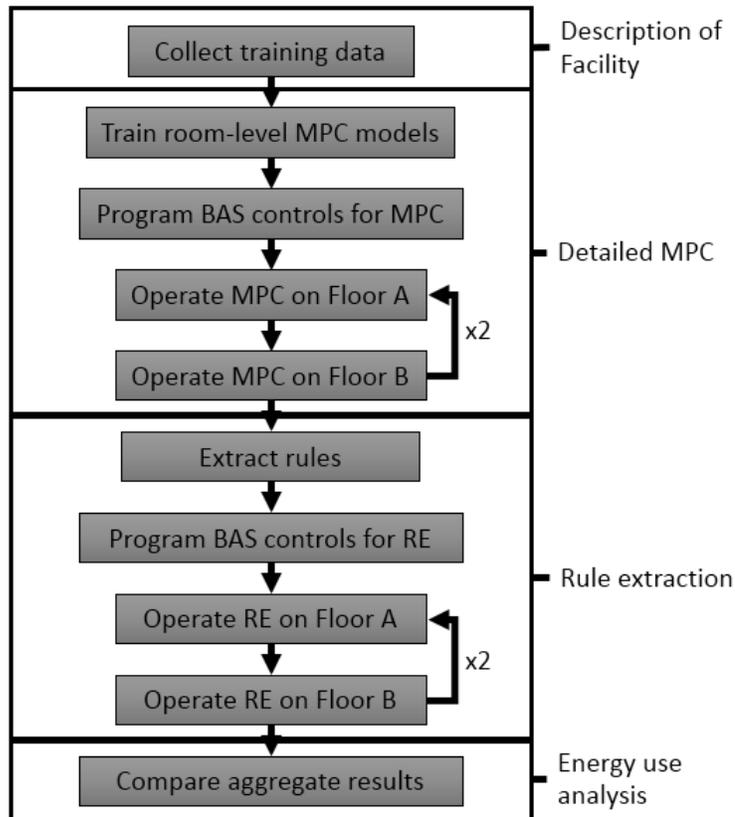


Figure 6-1: High-level MPC and RE design and implementation workflow diagram

### 6.2.1 Description of Facility

Control algorithms were compared using 27 rooms contained within 11 thermal zones distributed between two similar floors in an institutional building in Ottawa, Canada.

Twenty-five of the measured and controlled rooms are single occupant private offices and two are kitchenettes that are used as communal spaces.

Each office is the same size and shape, and has the same sensors and variable air volume (VAV) diffuser location and orientation. VAV units are shared between 2-4 offices depending on the location and floor. The kitchenettes have different floor areas than the offices and each other, and each have dedicated VAV units. The VAV terminal units are equipped with hot water reheat coils and the rooms have floor mounted hot water radiant heating panels near the windows as shown in Figure 6-2. The corner rooms (two kitchenettes and two offices, one of each per floor) have larger window to wall ratios than the other rooms (0.8 compared to 0.4).

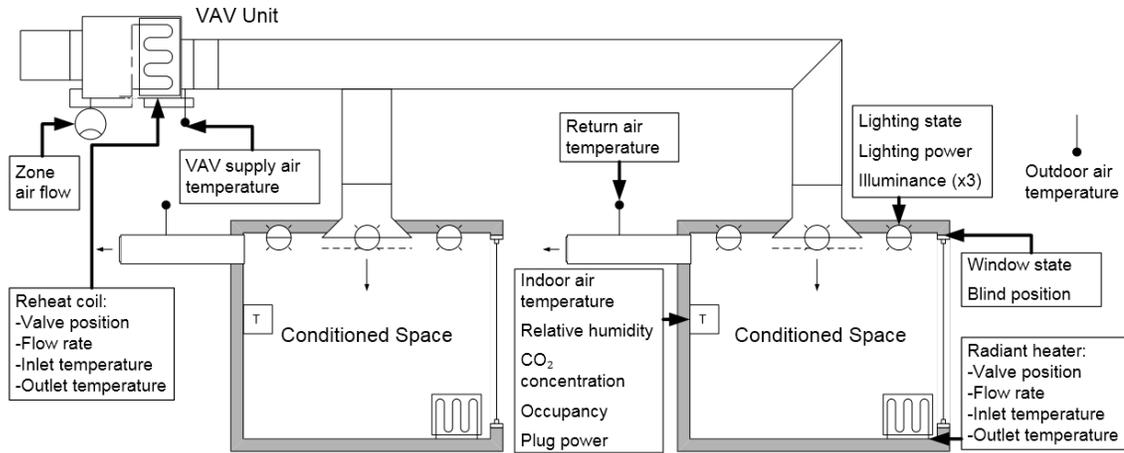
The reactive control was implemented during construction by the controls contractor using the sequence of operation documents. When heating or cooling are required, the reactive control modulates the VAV damper to adjust the airflow to each zone to meet the air requirement of the room with the greatest thermal demand in the zone. In the heating season the radiant heater is modulated to a maximum temperature drop (inlet to outlet) of 5°C. The room air setpoint can be adjusted by the occupants of each room via the thermostat and can vary from 18°C to 24°C. All room setpoints were reset weekly at midnight on Monday morning by the original controls if they deviated by more than 2°C from the seasonal setpoints (21°C in the heating season and 23°C in the cooling season). The air side control has multiple rooms per zone, while the hydronic radiant heating has only one office per unit. The combination of grouped room air side zone control and individual room setpoint adjustability can yield situations where meeting the setpoint of all

rooms is unlikely.

The VAV box also provides the outdoor air supply requirement and is active on a schedule from 7:00AM to 7:00PM supplying an average of 1,700 L/s to the entire office group when active (and no airflow to the offices when inactive). The floor-level AHU operates 24 hours per day every day of the week because it also supplies lab spaces. Due to the large amount of exhaust air from the lab spaces the outdoor air fraction was greater than 0.9 for most of the data acquisition period to maintain building pressurization for fume hoods. A complete list of the sensors used in this study and the corresponding bias error is provided in Table 6-1. A diagram of a typical room with sensor and heating, ventilating, and air conditioning (HVAC) equipment locations is provided in Figure 6-2.

**Table 6-1: Sensor schedule with sensor type and manufacturer-provided instrumentation error**

Measured Quantity	Sensor	Instrumentation error
Blind position	Controller output feedback	
Hour of day	Internal clock	
Illuminance near window	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance in center of room	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance near wall	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Illuminance in center of room	Photodiode	$\pm(6 \text{ lux} + 5\% \text{ of value})$
Indoor air carbon dioxide concentration	Dual beam nondispersive infrared	$\pm(50 \text{ ppm} + 2\% \text{ of value})$
Indoor air relative humidity	Integrated relative humidity and temperature	$\pm 3\%$
Indoor air temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Lighting power	Current transducer	$\pm 2\%$ full scale (4.8 W)
Lighting state	Controller output feedback	
Occupancy state	Passive infrared	
Outdoor air temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Plug load	Current transducer	$\pm 0.5\%$ of value
Radiant heater flow rate	Single jet flow sensor	$\pm 5\%$ full scale (30 l/h)
Radiant heater inlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Radiant heater outlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Radiant heater valve position	Controller output feedback	
Reheat coil flow rate	Single jet flow sensor	$\pm 5\%$ full scale (30 l/h)
Reheat coil inlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Reheat coil outlet temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
Reheat coil valve position	Controller output feedback	
Return air temperature	10 kOhm thermistor	$\pm 0.2^\circ\text{C}$
VAV airflow	Differential pressure transducer	$\pm 1\%$ full scale (5 Pa)
VAV supply air temperature	4-20 mA Resistance temperature detector	$\pm 0.2^\circ\text{C}$
Window state	Wireless dry contact	



**Figure 6-2: Diagram of a two-room zone with zone-level HVAC system and BAS-connected sensors.**

To represent a realistic case for the proposed algorithms the baseline control was not modified in the reactive offices. Existing control was only modified for MPC and RE testing by the insertion of a binary variable before the final output of the existing reactive control programming for overrides in the MPC and RE offices. While the baseline reactive control predated and was likely not compliant to ASHRAE Guideline 36 (High-Performance Sequences of Operation for HVAC Systems) (ASHRAE 2018), researchers were involved when the controls were defined and design effort was made to reduce energy use (e.g., automated blinds that open and close at night to conserve thermal energy, seasonal setpoint deadbands, nighttime setbacks).

Ethics approval was obtained for the logging and usage of room-level data in the measured building. Continuous data (e.g., temperature) was sampled at 5-minute intervals and binary data (e.g., window opening) was sampled when events occurred. Binary values from passive infrared occupancy sensors were used to infer occupancy and accompanying

loads (e.g., lighting and plug in equipment) as suggested in the literature (Gunay, O'Brien et al. 2016).

### 6.2.2 Uncertainty Analysis

A separation of variables approach was used to quantify aggregate instrumentation uncertainty of energy use estimation in this study (Smith 2013). The thermistors had bias uncertainty of 0.2°C. Bias due to thermistor positioning at the wall and not the center of the room was not quantified but is acknowledged to induce additional error. The differential pressure transducers had bias errors of 1% of full scale (2.5 Pa), corresponding to an average of 30 L/s per room. Room-level maximum airflow rates ranged from 61 L/s to 140 L/s. Air side sensible cooling and heating energy estimates were found to have errors of 32% of the average values using Equation 6-1, where  $s$  is the aggregate uncertainty,  $Q$  is the energy quantity being assessed for uncertainty,  $x$  is the input variable,  $i$  is the index of the input variable, and  $n$  is the number of input variables required to derive the energy quantity. Precision error,  $\Delta s_p$ , and bias error,  $\Delta s_b$ , are added vectorially in Equation 6-2 to estimate the total measurement uncertainty. The aggregate precision error of the air side sensible cooling and heating energy estimates was found to be 7.0% of the average value of air side sensible cooling and heating energy. Hydronic aggregate bias error of radiant heating energy estimates was found to be 13.7% of the average value. Precision error in the hydronic case was negligible (less than 0.1%). Analog to digital conversion values of 4096 were standard for all inputs on the BAS. The largest source of measurement uncertainty was from the air side pressure transducer, where bias errors of 1% full scale were reported by the manufacturer. Despite the high bias errors from manufacturer data,

spot measurements found an average airflow error of 4.0% of the measured values corresponding to an aggregate heating energy estimate error of 9.0%. The previously introduced errors are summarized in Table 6-2.

$$\Delta s_Q^2 = \sum_{i=1}^n \left( \frac{\partial Q}{\partial x} \right)^2 \Delta x_i^2 \quad (6-1)$$

$$\Delta s = \sqrt{\Delta s_b^2 + \Delta s_p^2} \quad (6-2)$$

**Table 6-2: Sensor error summary table**

Type of error	Error magnitude
Differential pressure transducer bias error	1.0% of full scale
Sensible cooling and heating energy aggregate error	32.0% of average value
Air side sensible cooling and heating energy precision error	7.0% of average value
Radiant heating energy estimate bias error	13.7% of average value
Radiant heating energy estimate precision error	<0.1% of average value
Average measured airflow error	4.0% of average value
Average heating energy estimate error using measured airflow error	9.0% of average value

### 6.2.3 Detailed Model Predictive Control Development

Extracting rules from detailed MPC involved first developing an MPC algorithm that could improve building operation by reducing energy use while maintaining or improving occupant comfort. The predictive control objectives were accomplished by minimizing energy use and setpoint deviation using a cost function described in the Model Development section. The following sections describe the test plan, model development, and control strategy for the MPC algorithm.

### 6.2.3.1 Test Plan

One month of preliminary sensor data was collected from the BAS serving the offices to train and validate inverse models for predicting future office indoor air temperatures in the MPC algorithm. Following detailed MPC testing, classification learning was used to extract rules for RE control and testing.

Floor-level algorithms were alternated on a weekly basis as shown in Table 6-3. Under this testing plan effects from outdoor environment conditions and occupant behavioral effects on floor-level energy use are minimized. The fraction of time with at least one occupant is the portion of time when the ceiling motion detector senses movement, where 0.24 is the equivalent of a 40-hour work week. An average of 0.042 was found over the entire testing period for all rooms, corresponding to seven hours per week. These offices were primarily occupied by professors with significant lecture and meeting obligations, explaining the low occupancy ratios. Additionally, many of the rooms had weeks where they were only entered briefly by the cleaning staff, bringing down the average from more regularly occupied rooms.

**Table 6-3: Floor-level weekly algorithm schedule, average OAT, and average occupant fraction**

Season	Week	Start Date	End Date	Algorithm - Floor A	Algorithm - Floor B	Average outdoor air temperature (°C)	Fraction of time with at least one occupant - Floor A	Fraction of time with at least one occupant - Floor B
Cooling	1	23-07-2018	27-07-2018	MPC	Control	22.2	0.031	0.034
Cooling	2	30-07-2018	03-08-2018	Control	MPC	23.6	0.030	0.036
Cooling	3	06-08-2018	10-08-2018	MPC	Control	23.2	0.020	0.023
Cooling	4	13-08-2018	17-08-2018	Control	MPC	21.9	0.033	0.041
Cooling	5	27-08-2018	31-08-2018	RE	Control	21.6	0.018	0.034
Cooling	6	03-09-2018	07-09-2018	Control	RE	20.4	0.051	0.049
Cooling	7	10-09-2018	14-09-2018	RE	Control	18.2	0.046	0.057
Cooling	8	17-09-2018	21-09-2018	Control	RE	17.4	0.042	0.051
Heating	1	21-01-2019	25-01-2019	MPC	Control	-8.5	0.022	0.049
Heating	2	28-01-2019	01-02-2019	Control	MPC	-14.6	0.052	0.064
Heating	3	04-02-2019	08-02-2019	MPC	Control	-6.9	0.054	0.049
Heating	4	11-02-2019	15-02-2019	Control	MPC	-7.6	0.031	0.038
Heating	5	25-02-2019	01-03-2019	RE	Control	-9.1	0.046	0.059
Heating	6	04-03-2019	08-03-2019	Control	RE	-7.8	0.040	0.061
Heating	7	11-03-2019	15-03-2019	RE	Control	0.0	0.035	0.056
Heating	8	18-03-2019	22-03-2019	Control	RE	-1.1	0.030	0.055
Heating	9	25-03-2019	29-03-2019	RE	Control	0.7	0.041	0.060

### 6.2.3.2 Model Development

Detailed MPC was executed using a control agent to determine the result of the cost function in Equation 6-3, where  $J$  is the cost (an arbitrary and differentiable constant in this study),  $SP$  is the temperature setpoint provided by the occupant,  $T$  is the air temperature as observed by the BAS and predicted using modelling,  $\Delta U$  is the energy input from the heating and cooling system, and  $R$  is a matrix of penalty weights for different time-steps within the prediction horizon adapted from Wang et al. (Wang 2009). The control agent

collects system input and output data to determine what actions to take. A prediction horizon of two to three hours was found to be appropriate for buildings without large amounts of thermal mass, therefore a prediction horizon of three hours was chosen in this study (Gunay, Bursill et al. 2014, Bursill, O'Brien et al. 2019). The VAV damper positions and VAV reheat coil valves were controlled by the MPC in the cooling season, and the VAV damper positions, VAV reheat coil valve positions, and radiant heater valve positions were controlled in the heating season.

$$J = (SP - T)^T (SP - T) + \Delta U^T R \Delta U \quad (6-3)$$

In Equation 6-3, the weight matrix  $R$  was set as the identity matrix and was intended to provide a compromise between maintaining the setpoint and minimizing energy use for all prediction steps. Weight matrices may need to be adjusted in cases where time of use rates are applicable or there is any other reason to explicitly shift cooling and heating from one time-step to another. Time of use energy price rates were not considered in this study because the energy provided to the building was not subject to them, thus reducing the importance of adjusting the weights of the  $R$  matrix and improving the simplicity of the detailed MPC in this study. The linear time-invariant (LTI) state-space thermal model form used for the spaces shown in Equation 6-4, where  $t$  is the time-step index,  $u$  is the input array, and  $A$  and  $B$  are the constants for each respective space. State-space parameters were found using the MATLAB system identification toolbox which utilizes a subspace method (Ljung 1995). States and inputs were constrained by the limits of the sensor readings and equipment capacities. The parameter values of  $A$  and  $B$  were unconstrained. Because LTI

state space models were used, the models must be retrained when input ranges change significantly (e.g., seasonal transition, significant furniture addition or removal, envelope upgrade) to maintain prediction accuracy and account for differing environmental and occupant conditions. Unique LTI state-space models were used for each room and for the cooling and heating seasons to account for changing boundary conditions and to improve the accuracy of the linear models (which are modelling highly nonlinear physical processes). The input array for the model selected in this study consisted of five sensors from Table 6-1 adapted based on recommendations in the literature, where  $u_1$  is the room ceiling illuminance in lux,  $u_2$  is the VAV air side input,  $u_3$  is the outdoor air temperature (OAT),  $u_4$  is the occupancy state, and  $u_5$  is the radiant panel valve position (Gunay, Bursill et al. 2014, Gunay, O'Brien et al. 2016, Bursill, O'Brien et al. 2018). The VAV air side input ( $u_2$ ) was calculated using Equation 6-5 where  $V_{VAV}$  is the air flow rate,  $T_{VAV1}$  is the VAV supply air temperature, and  $T_{VAV2}$  is the VAV return air temperature. Equation 6-4 expands to Equation 6-6 when considering the five model inputs in this study. The ceiling illuminance sensor was used as a proxy for the solar gains of the space because building envelope surface irradiation correlates with ceiling illuminance, and the rooms studied each had illuminance sensors available (Bursill, O'Brien et al. 2018). Solar gains were relatively low because of the orientation and size of the windows, and it is acknowledged this approach might not work for highly glazed spaces. The use of illuminance sensors as a solar proxy was also found convenient and practical by Gunay et al., and was confirmed by Bursill et al. (2014, Bursill, O'Brien et al. 2019). The state-space model inputs were chosen based on the past success of other zone-level MPC studies with similar objectives

(Gunay, Bursill et al. 2014, Gunay 2016, Bursill, O'Brien et al. 2018, Bursill, O'Brien et al. 2019).

$$T(t + 1) = AT(t) + Bu(t) \quad (6-4)$$

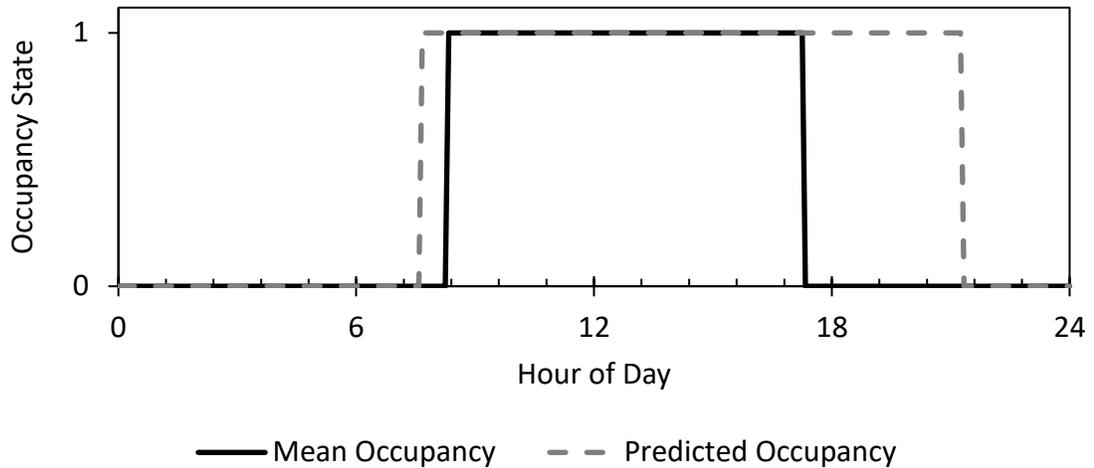
$$u_2(t) = V_{VAV}(t)(T_{VAV1}(t) - T_{VAV2}(t)) \quad (6-5)$$

$$T(t + 1) = AT(t) + B_1u_1(t) + B_2u_2(t) + B_3u_3(t) + B_4u_4(t) + B_5u_5(t) \quad (6-6)$$

The predictive algorithm used the previously developed indoor air temperature (IAT) models and local weather station predictions. Hourly weather predictions were updated at six-hour intervals for illuminance (converted from solar radiation, used as a solar proxy) and outdoor air temperature prediction from data provided by the CanMETEO program developed by NRCan (Candanedo 2018). The predicted incident solar radiation on the wall from CanMETEO was adjusted to an equivalent illuminance (lux) value using a linear fit to two weeks of hourly data (validated with an additional two weeks of data). The  $R^2$  values comparing recorded predictions to estimated values from the linear fits ranged from 0.27 to 0.72 for the cooling season and 0.17 to 0.48 for the heating season for the measured rooms. Both the site observations and CanMETEO predictions use different sensors in different orientations to measure different quantities, thus significant error was anticipated. Despite the low  $R^2$  value between the solar irradiation and illuminance predictions, these fits are considered adequate for this application due to the cost and difficulty of installing surface irradiation sensors in most commercial buildings and the

proxy usage of this inferred value. Using the MAE of the illuminance fits in the state-space models yielded an average sensitivity of  $0.06^{\circ}\text{C}$ . Degradation of the line fit can be attributed to differing shade positions, physical quantities, and local solar conditions between the airport weather station and building location (approximately 8km apart) among the previously described factors. The heating season fits were worse than the cooling season fits which could be attributed to solar geometry and dynamic reflective conditions from local snow properties introducing additional nonlinearity (Purdy and Beausoleil-Morrison 2001).

The timing of first occupant arrival and last occupant departure events were predicted recursively using logistic regression as described in Gunay et al. (Gunay, O'Brien et al. 2015). Arrival and departure events were assumed to be Gaussian for the applied model. The occupied schedule was predicted as the mean first arrival minus 1.5 standard deviations to the mean last departure plus 1.5 standard deviations. This accounts for roughly 90% of all arrivals and departures and is a conservative schedule estimate under the Gaussian assumption. A sample predicted schedule for a room is provided in Figure 6-3, where the standard deviations of the arrival and departure times are 0.42 and 2.64 hours respectively. The time between the estimated arrival and departure times on weekdays was assumed to be occupied without absences or breaks in the prediction algorithm as the input to future time-step iterations of  $u_4$  in Equation 6-6. Only a weekday schedule was generated for each room because the weekends were assumed to be unoccupied by both the building VAV schedule and this study (and were verified to have significantly lower occupancy ratios than the weekdays).



**Figure 6-3: Mean occupancy and predicted occupancy schedule for room 1. The predicted occupancy schedule calculates the recursive mean arrival and departure times and applies 1.5 standard deviations of the respective times to encompass 90% of all arrivals and departures under a Gaussian assumption. This follows the same format as used in (Gunay, O'Brien et al. 2016).**

### 6.2.3.3 Control Strategy

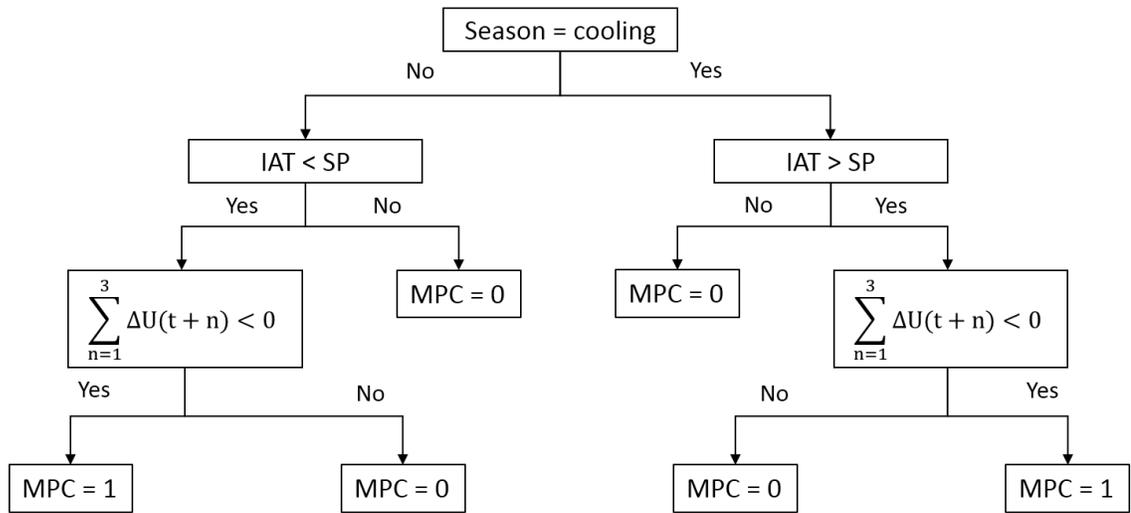
Control decisions to override reactive control operation during MPC testing were sent to the BAS controllers based on the energy balance of Equation 6-7, a differentiation and rearrangement of Equation 6-3 (Wang 2009). Equation 6-7 introduces  $\varphi$ , which is expanded in Equation 6-8, and  $H$ , a grouping of terms that remain nearly constant (Gunay 2016). The MPC algorithm is applied to each room with unique state-space parameters and illuminance fits for each room. The setpoint logic with a three-step (1.5 hours) control horizon is applied in Figure 6-4, where the zero or one output of the MPC algorithm is be used to determine whether cooling or heating should be provided to each individual room. The  $\Delta U(t+n)$  term in Figure 6-4 represents the energy required to reach the setpoint at time  $t$  plus time-step  $n$ , where a summation of three time-steps occurs over the control

horizon to predict the heating or cooling needs of the space. A one in Figure 6-4 indicates the HVAC equipment should operate normally, and a zero indicates it should stop operating to reduce setpoint overshoot. A control horizon of half of the prediction horizon was chosen as a compromise between receding horizon control (one time-step) and using all prediction steps in the control horizon.

Multi-room zones were controlled such that if more than one room in the zone was predicted to have a controllable deviation from the setpoint when running the MPC algorithm (e.g., if the setpoint was 1°C below the setpoint in the heating season for two of three rooms in a zone and at the setpoint in the third room), the HVAC system outputs would be allowed to operate using the base reactive control loops in a “business as usual” format. In cases where one room or less were predicted to require conditioning (e.g., if the setpoint was 1°C below the setpoint in the heating season for one of three rooms in a zone and at the setpoint in the other two rooms), the HVAC system outputs were overridden for the VAV unit to prevent predicted setpoint overshoot. The effect of this zone control adjustment when using the MPC and RE algorithms is discussed based on differences in actuator commands in the Control Mode Operational Comparison section. During the heating season, room-level radiant heaters were controlled based on room-level MPC algorithm outputs that did not require a multi-room zoning approach. In this study the carbon dioxide levels were monitored (but not controlled) and verified to have not exceeded 1000 ppm for more than 30 minutes at a time.

$$\Delta U = (\varphi^T \varphi + R)^{-1} \varphi^T (SP - T) = H(SP - T) \quad (6-7)$$

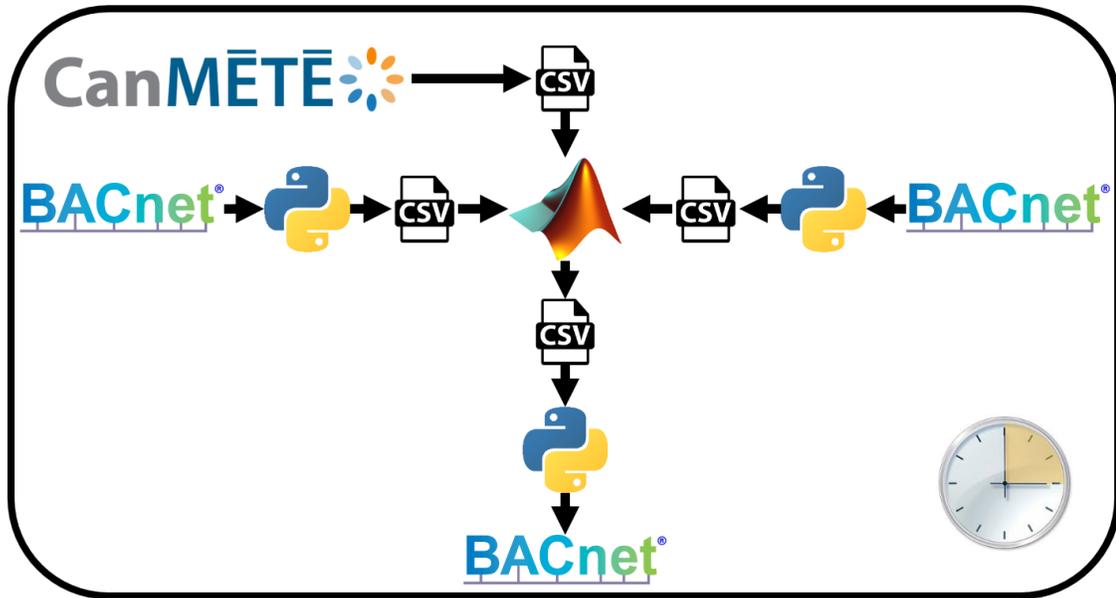
$$\varphi_{n \times n} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ A & 1 & 0 & 0 & 0 & 0 \\ A^2 & A & 1 & 0 & 0 & 0 \\ A^3 & A^2 & A & 1 & 0 & 0 \\ A^4 & A^3 & A^2 & A & 1 & 0 \\ A^5 & A^4 & A^3 & A^2 & A & 1 \end{bmatrix} \quad (6-8)$$



**Figure 6-4: Flow chart for reactive control suppression from the MPC algorithm and setpoint data where  $t$  and  $n$  are the current time-step and predicted time-step respectively. When  $MPC = 0$  it indicates the HVAC equipment should be overridden, and when  $MPC = 1$  it indicates the HVAC equipment should continue using reactive control.  $\Delta U$  is the array of predicted energy required by the HVAC system at each timestep to condition the IAT to the IAT setpoint.**

The previously described components of the control agent in the detailed MPC approach were interfaced with the BAS using the workflow in Figure 6-5. Data is read from the BAS inputs and outputs, and processed overrides are written back to the BAS for

each room. BACnet data is converted to and from comma delimited files (csv) using Python scripts with BACnet stack commands. Weather prediction data is provided by CanMETEO in comma delimited format. All comma delimited data is read by the agnostic control engine (ACE) and processed for predictive control based on the previously described algorithm. For this study MATLAB was used as the ACE, but in the workflow of Figure 6-5 any programming environment that can read and write from comma delimited files can be used. The equations of the Detailed Model Predictive Control Development section are applied to the BAS data, their results are processed by the logic in Figure 6-4, and the equations and logic are embedded in the ACE and interface with the building controllers (and reactive control sequences) via the paths in Figure 6-5. The resulting detailed MPC calculations and workflow are complicated and unreliable for large scale implementation by practitioners and are the motivation for developing a simpler rule extraction approach.



**Figure 6-5: Detailed MPC control agent workflow. BACnet reactive controller input, BACnet reactive controller output, and CanMETE weather prediction data are read into comma delimited files on a PC using Python scripts. The ACE then executes the MPC algorithm on the data, writes the predicted overrides into comma delimited files, and sends BACnet output override commands back to the BAS from the comma delimited files via a script. This is all executed on a at regular intervals using Windows Task Scheduler.**

#### **6.2.4 Rule Extraction Predictive Control Development**

The goal of simplified rule extraction is to find the minimum number of inputs, thresholds, and parameters required to maintain most of the MPC energy savings and similar levels of comfort for the occupant. This approach proposes and develops a commercially practical implementation by removing the optimization calculation from building automation controllers where matrix operations and computational power are limited. Matrix multiplication and inversion cannot be executed in database programs of most BACnet building controllers without manual addition of code, rendering commercial usage of database matrix operations virtually non-existent. Decision trees were generated in this

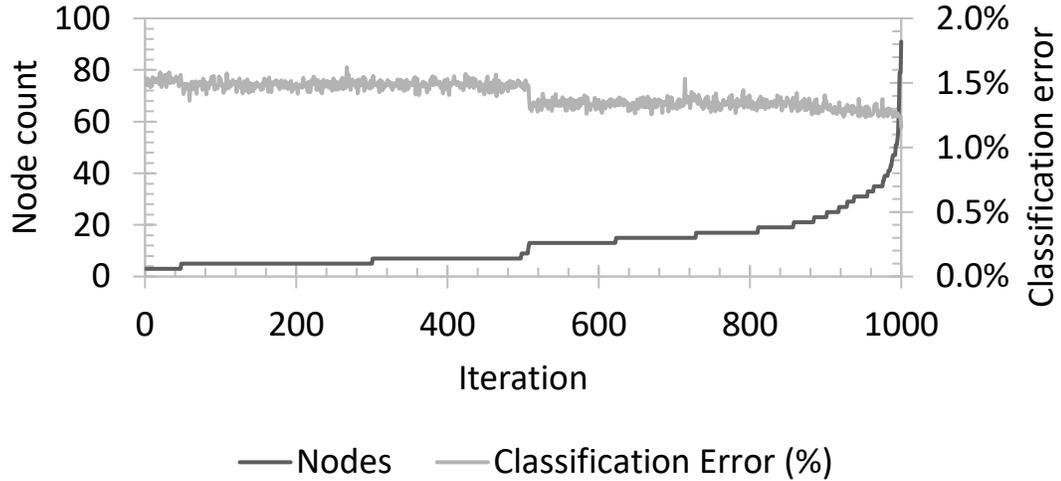
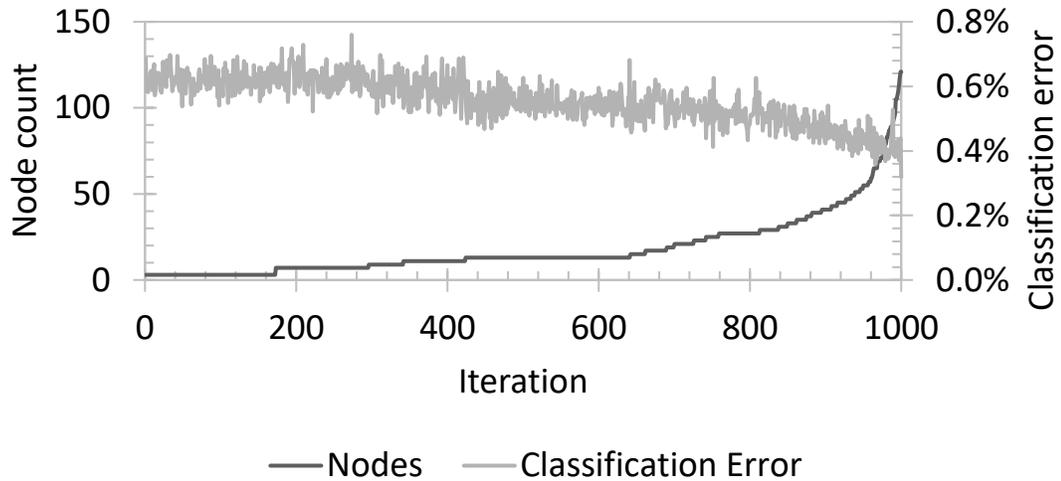
study using the MATLAB Classification Toolbox (Duin, Juszczak et al. 2000).

In this study, the prototype rule extraction control was developed based on the control decisions of the detailed MPC algorithm (i.e., the detailed algorithm was required to extract rules for this building). However, a long-term goal of future work is to generalize this approach by developing rules based on detailed MPC for several building archetypes, and utilize reinforcement learning to tune to individual buildings. Simulation studies already exist in the literature applying a similar approach to archetype buildings to extract rules (Coffey, Haghghat et al. 2010, Coffey 2013). The approach in this study is intended as a methodology to apply detailed MPC in a format that can easily interface with rule extraction for real BASs.

Data collection for each season of detailed MPC for the 27 rooms resulted in over 40,000 total time series samples (roughly 1,500 samples of each data point per room) containing all the physical input and outputs of the BAS controllers and the intermediary parameters and predictions from the detailed MPC algorithm. All 27 rooms had the same primary façade and similar construction; thus, a single decision tree was proposed for room control. This had two primary benefits: the aggregate sample size was larger, and the implementation of a single set of extracted rules per façade (with similar space type and use) is significantly simpler to implement within a BACnet BAS controller via the GCL+ programming object than unique rule for each room.

One important challenge introduced by aggregating the room data is the potential for overfitting to certain rooms that may be more prevalent in the training data set (particularly the rooms with higher window to wall ratios than the majority). To mitigate

this, the samples were shuffled randomly, and classification tree learning was executed 1,000 times to identify the different rulesets that would be extracted by the random forests. For training and validation, the shuffled data was divided into two equal sets. The 1,000 iterations were considered an ensemble, where multiple data-driven models are used to reduce the variation of selected model parameters. Junctions in the decision trees where classifications were made (nodes) were also recorded. The classification error progression and node counts are plotted with respect to ensemble iteration in Figure 6-6 (sorted by classification tree node count). Classification error had a stochastic element due to the random shuffling of samples.



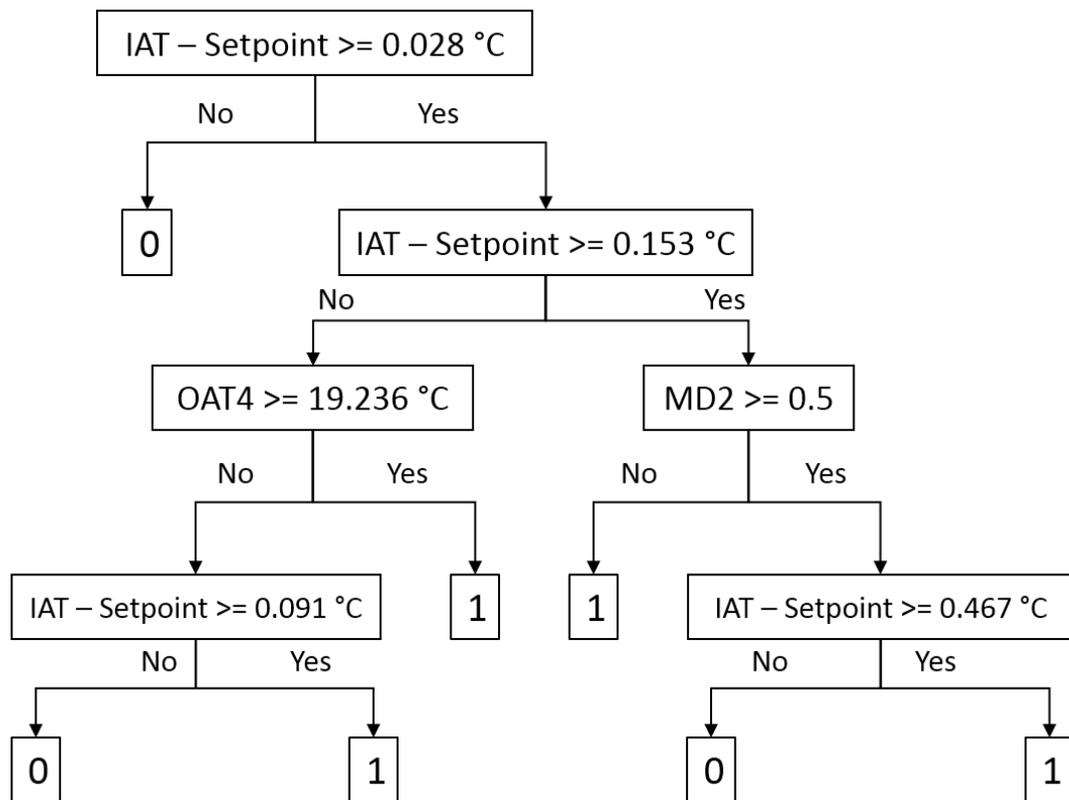
**Figure 6-6: Node vs. iteration for the cooling season (top) and heating season (bottom) classification tree fitting processes for 1000 sample-randomized scenarios. Results were sorted such that node counts increment from left to right. One node represents one split where two unique decision tree outcomes can occur.**

The rule extraction process resulted in 49 unique rule sets for the cooling season, and 28 unique rule sets for the heating season. Most of the additional unique rule sets for the cooling season occurred with larger classification trees. A compromise between

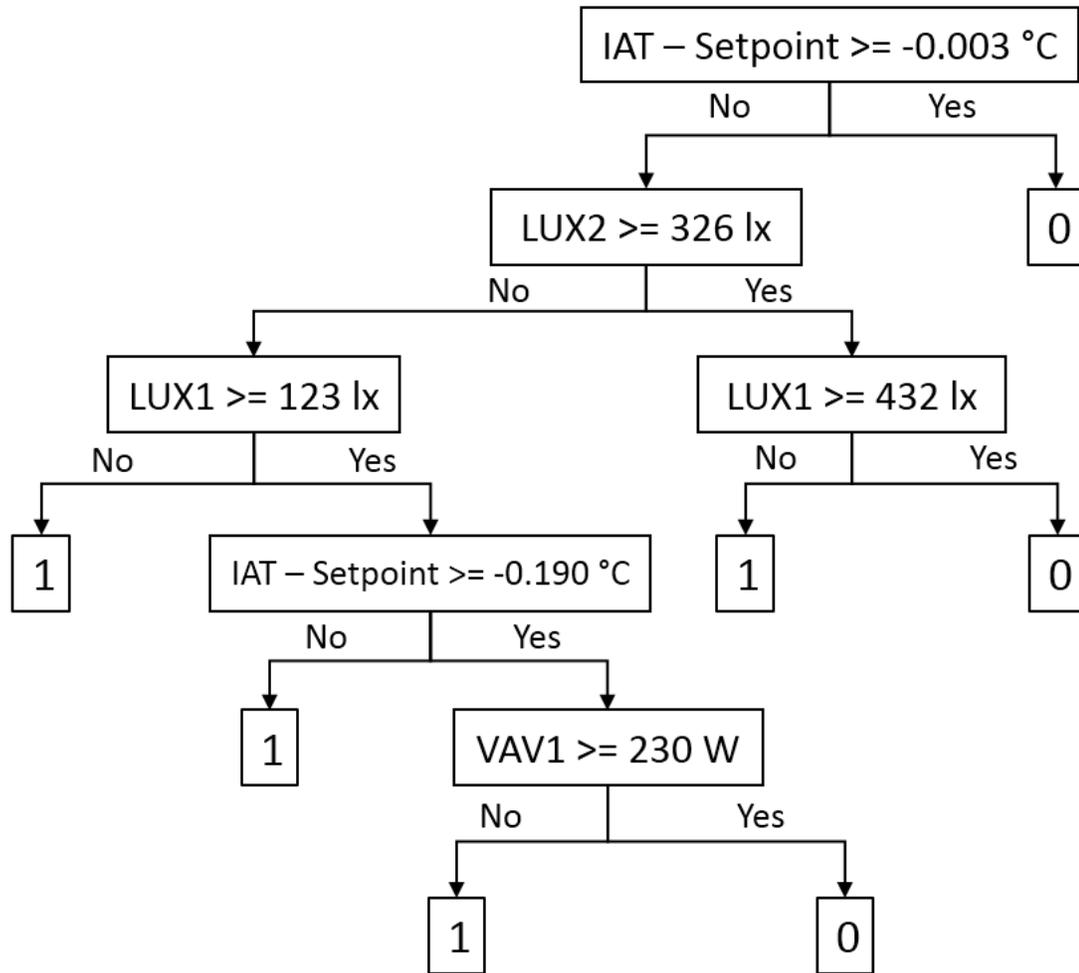
accuracy and complexity was considered when selecting the best rule set for the rooms where a distinct “jump” in classification accuracy or the median (if no distinct “jump” is present) were selected from the ensemble. As a result, the 13 node trees were selected for both the cooling and heating seasons. The selected decision trees correspond with the median iteration for the cooling season and where there was a distinct jump in accuracy in the heating season when the ensemble increased from 11 to 13 nodes. The selected 13 node cooling and heating season decision trees are shown in Figure 6-7 and Figure 6-8. Incorrect classification of cooling and heating season MPC control decisions occurred in 0.5-0.7% and 1.3-1.4% of the samples respectively for the selected rules.

The resulting decision trees from the extracted rules relied on significantly less data than the detailed MPC algorithm. In the cooling season room IAT, IAT setpoint, occupancy prediction for 30 minutes in the future (MD2), and OAT prediction for 2 hours in the future (OAT4) are required. In the heating season room IAT, IAT setpoint, VAV heating power (VAV1), center of room ceiling illuminance (LUX1), and center of room ceiling illuminance prediction for 30 minutes in the future (LUX2) are required. This approach only requires four and five control inputs to emulate the MPC algorithm in the measure rooms in the cooling and heating seasons respectively. Of the previously mentioned inputs, predictions are only required for occupancy, outdoor air temperature, and illuminance. Outdoor air temperature predictions can be acquired from weather station applications, and occupancy schedule predictions can be made either with an embedded model or linking to enterprise level schedule data (and possibly access card logs). Illuminance prediction is the most difficult prediction to acquire in the proposed decision trees and required an

individual linear fit for each room to convert from weather station solar radiation predictions. Other forms of modelling could be used to predict proxy solar gains in the rooms (e.g., moving average, autoregressive moving average), but room-level linear models were selected here to provide higher granularity and reduce algorithm complexity.



**Figure 6-7: Cooling season decision tree. A 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control.**



**Figure 6-8: Heating season decision tree. A 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control.**

The cooling season decision tree notably does not include any kind of solar proxy. This can be attributed to the measured rooms not receiving significant solar gains during most of the operating period due to their North-East orientation. Rooms with facades that receive more solar gains may include solar terms in the decision tree due to prevalence of radiative effects. Likewise, the heating season decision tree does not include occupancy or outdoor air temperature. These omissions are likely due to the relatively low occupancy

ratio of the offices (most were occupied by professors with erratic schedules) and the relatively constant and low ambient air temperatures of the local climate in the daytime of the heating season (Ottawa, Canada). In both decision trees, the selected inputs and threshold were the result of the random forest machine learning algorithm and were representative of the transient characteristics of the MPC data (Van Der Heijden, Duin et al. 2005).

Due to the nature of the inverse modelling approach applied to the rooms in the building studied, it is not possible to consider the generated decision trees as universal control algorithms. The presented decisions trees are a result of the methodology applied, and a proof of concept for utilizing classification trees for predictive control. The energy and comfort performance of the reactive control, MPC, and RE approaches for the measured rooms are compared in the following section.

### **6.2.5 Energy Use Analysis**

Energy use had to be quantified at the room level to determine if energy savings were achieved by MPC and RE control relative to the baseline reactive control. Data from the hydronic system and AHU sensors were used to estimate the room-level energy use.

Steady-state heat transfer was assumed over each 5-minute interval. The literature asserts that this is a valid assumption for 15-minute (or smaller) intervals (Bursill, O'Brien et al. 2018). The energy use at each sample is calculated as the average over the sample interval as shown in Equation 6-9, where  $E$  is the energy use at each time-step,  $Q$  is the average power at each time-step, and  $t$  is the index of the time-step. The heating power of

the radiant heaters and reheat coils was estimated with Equation 6-10, where  $V_h$  is the water flow rate,  $\rho_w$  is the density of water at 40°C,  $c_{pw}$  is the specific heat of water at 40°C,  $T_{1w}$  is the supply water temperature, and  $T_{2w}$  is the return water temperature. Thermo-physical properties were acquired from Incropera et al. (Incropera, Lavine et al. 2007). Heating and cooling energy supplied by the AHU was calculated using Equation 6-11, where  $Q_{SA}$  is the supply air energy,  $V_r$  is the estimated room-level air flow rate, and  $V_{AHU}$  is the AHUs air flow rate.

$$E(t) = Q(t)\Delta t \quad (6-9)$$

$$Q_h(t) = V_h(t)\rho_w c_{pw}(T_{1w}(t) - T_{2w}(t)) \quad (6-10)$$

$$Q_{SA}(t) = \frac{V_r(t)}{V_{AHU}(t)} \left( V_h(t)\rho_w c_{pw}(T_{1w}(t) - T_{2w}(t)) \right) \quad (6-11)$$

### 6.3 Results and Discussion

This section demonstrates the value of the proposed rule extraction approach by quantifying the energy savings from testing the three different control algorithms over two floors of similar offices of an institutional building over the cooling and heating seasons. Results are presented in one-day periods for one room and are followed by aggregated results for all offices. Highlighted results include control mode operational comparisons, energy use comparisons, and clear statements of the limitation of the approach presented in this paper.

### **6.3.1 Control Mode Operational Comparison**

The reactive control schedule-based IAT, IAT setpoint, VAV operation state (0/1 for OFF and ON), and heat supplied by the radiant heating panels ( $Q_{\text{Rad}}$ ), VAV reheat coil ( $Q_{\text{RHC}}$ ), and AHU coil ( $Q_{\text{AHU}}$ ) are shown for both the cooling and heating seasons in Figure 6-9 for a sample room over one day (room 3, from within a three-office zone with one exterior wall on the primary façade). The cooling and heating seasons will have different HVAC equipment requirements and loads due to ambient conditions (e.g., there is no radiant heater operation in the cooling season). Reactive control maintains VAV operation over the entire day with a relatively high minimum airflow rate (an average of 25 L/s for the single occupant offices) which can lead to over-ventilation (i.e., when more fresh air is supplied to a room than is required by code or for the occupants), a common problem in modern buildings (Chen 2009).

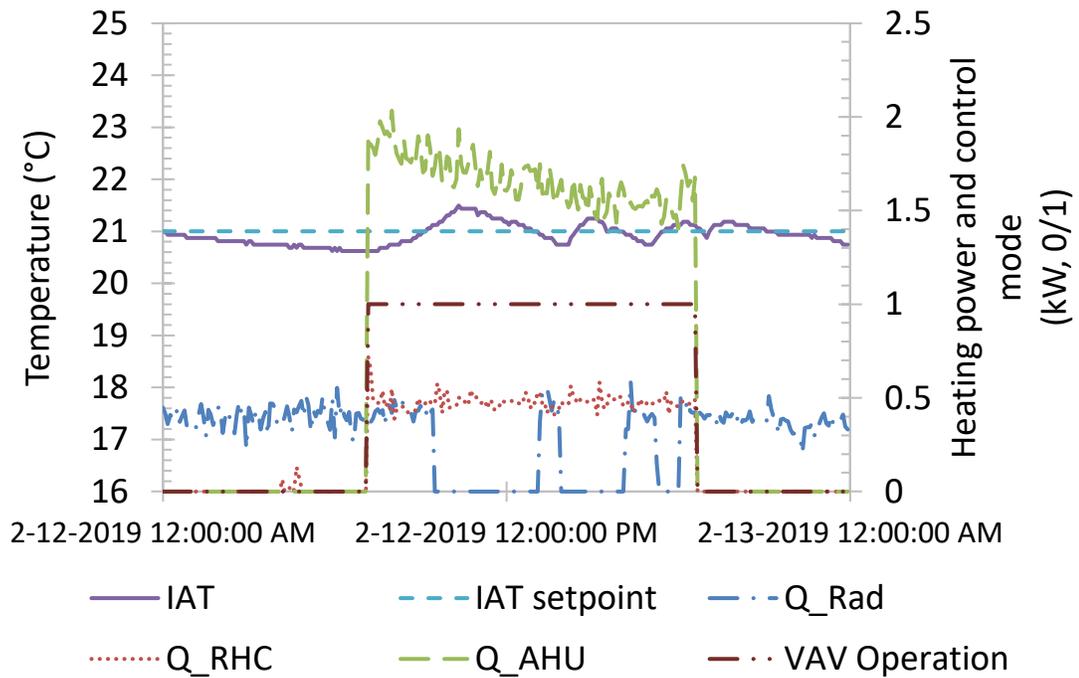
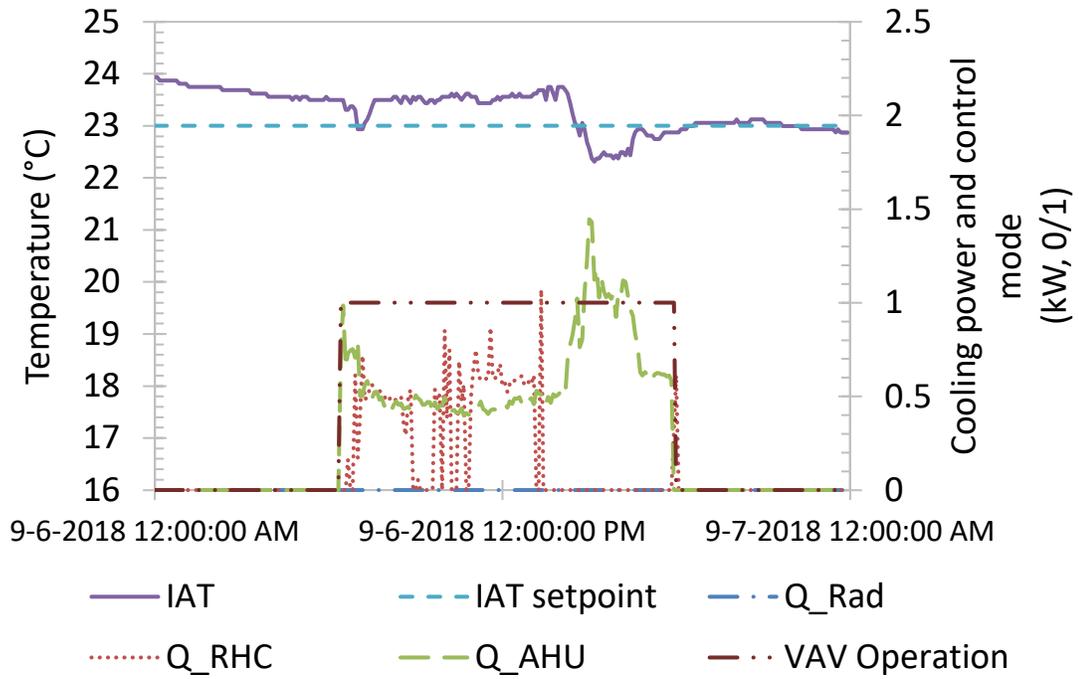


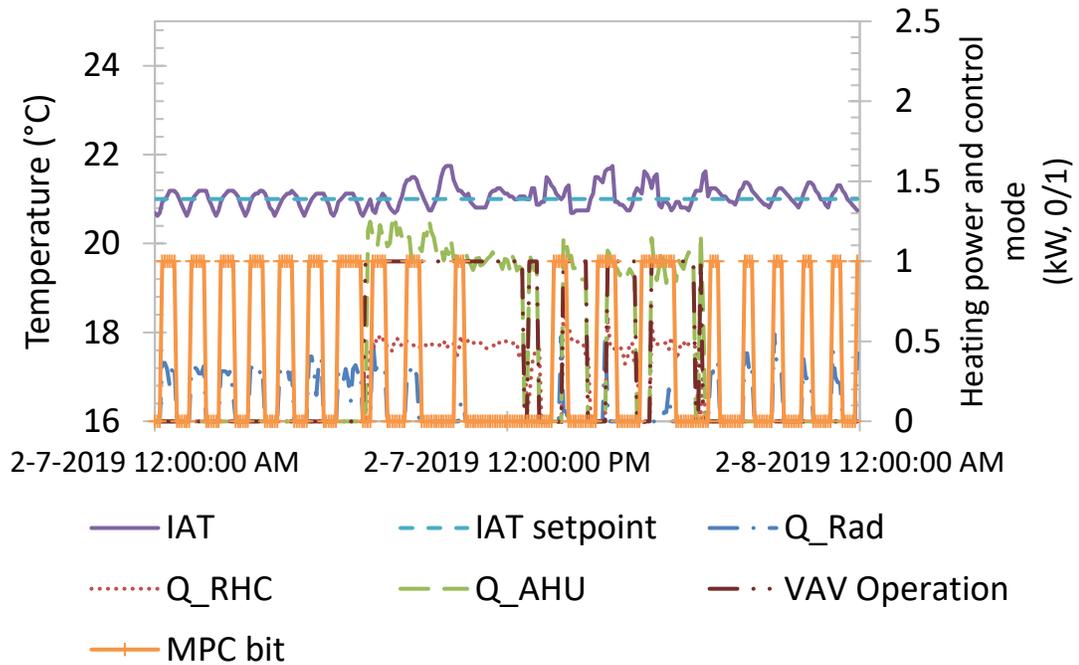
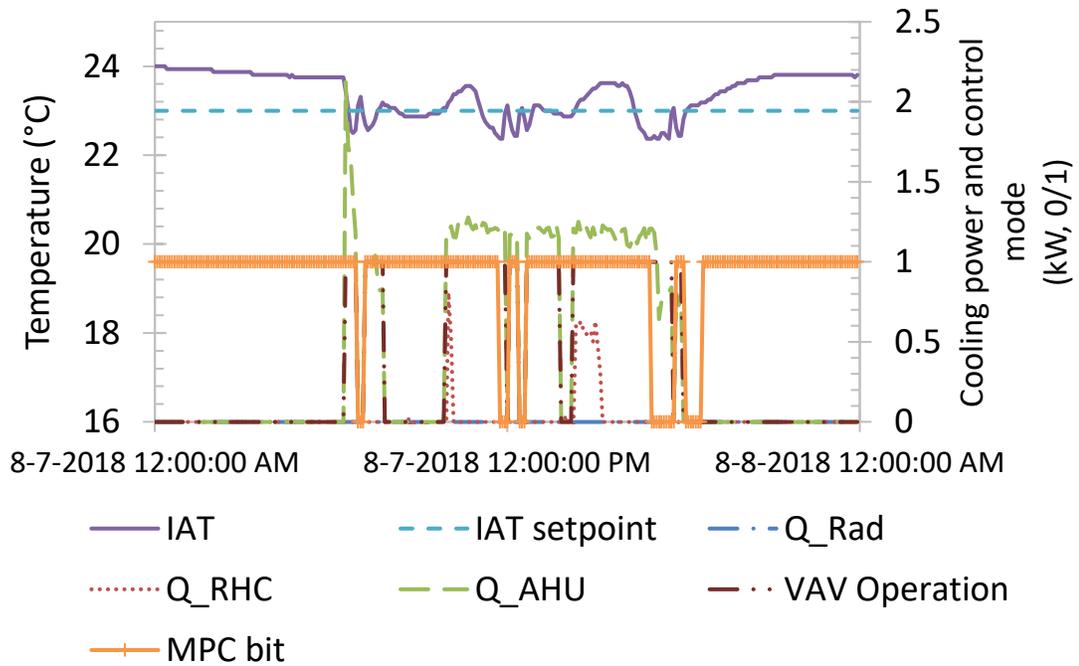
Figure 6-9: Reactive control operation for room 3 in the cooling season (top) and heating season (bottom)

Reactive control HVAC operation in the cooling and heating seasons are clearly distinguished from one another by the 24-hour availability of the radiant heater. The IAT unoccupied setback is most apparent in Figure 6-9 and Figure 6-10 (reactive control and MPC cases). There is no programmed coordination between the VAV heating and radiant panel heating units, which can also lead to additional overshoot of the setpoint in the heating season. Additionally, the deviation between cooling and heating requirements of the individual rooms within the zone can lead to sub-optimal setpoint control for at least one room within the zone.

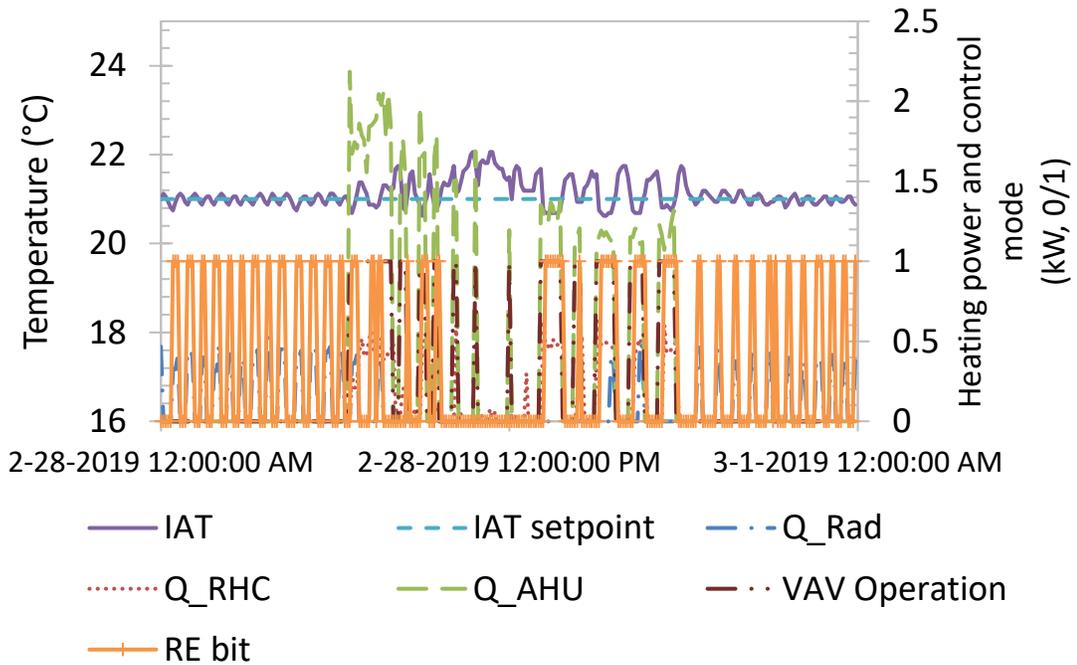
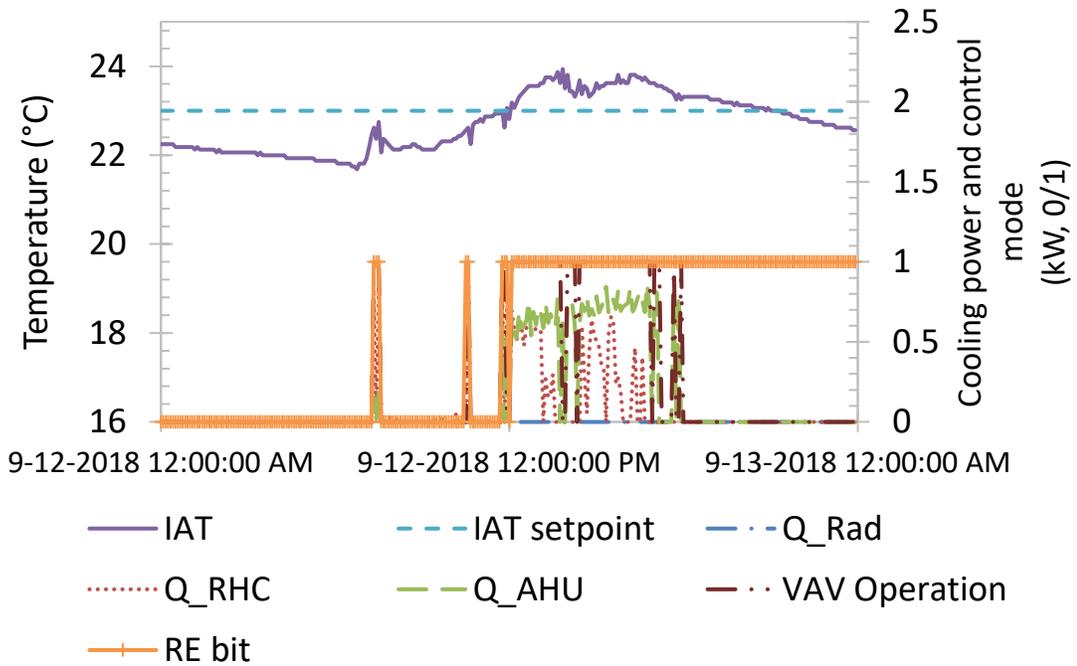
The MPC algorithm presented in this study used machine learning (state-space models), input predictions (illuminance, occupancy, OAT, and HVAC cooling/heating), and modified zone control (code to override reactive control) to reduce energy use while maintaining similar or better occupant comfort compared to the reactive control algorithm as defined by setpoint deviation. Typical MPC operation for one room over a sample day in the cooling and heating season is shown for room 3 in Figure 6-10, where the MPC bit and RE bit values are the 0 or 1 values (corresponding with overrides and normal reactive control respectively) from the MPC algorithm or RE decision trees. Zone-level VAV turndown in anticipation of setpoint overshoot can be seen in the MPC sample days. Setpoint overshoot is maintained to be similar or less than with the original reactive control algorithm, but the energy savings can be interpreted visually from the turndown of the VAV unit supply of AHU cooling or heating (and the turn down of the reheat coil and radiant heaters during the heating season). Further details of the setpoint deviations are discussed in the Energy Use section and weekly average setpoint deviations are presented

in Table 6-4.

One day of operation for room 3 of the simplified RE algorithm is shown in Figure 6-11 for the cooling and heating season in the same format as Figure 6-10. The IAT is maintained closer to the setpoint when the RE algorithm is active, and this graphical observation is verified in Table 6-4. However, this also produces some cycling that occurs at a maximum frequency of every five minutes (the loop delay of the decision trees within the BAS controller program). Setpoint deviation is smaller in the RE case largely due to the precise thresholds within the decision tree. Reactive control and MPC have some inherent deadbands due to PI (proportional integrating) loops and real-time optimization computational delays respectively.



**Figure 6-10: MPC operation for room 3 in the cooling season (top) and heating season (bottom). For the MPC bit, a 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control.**



**Figure 6-11: RE operation for room 3 in the cooling season (top) and heating season (bottom). For the RE bit, a 0 indicates the HVAC equipment should be overridden and a 1 indicates it should continue using reactive control**

MPC and RE overrides were compared to see how well the RE control emulated the MPC decisions in Figure 6-12 and Figure 6-13. MPC and RE overrides do not have priority simultaneously (i.e., the MPC override does not work during RE scheduled weeks and vice versa). Consider the following two examples:

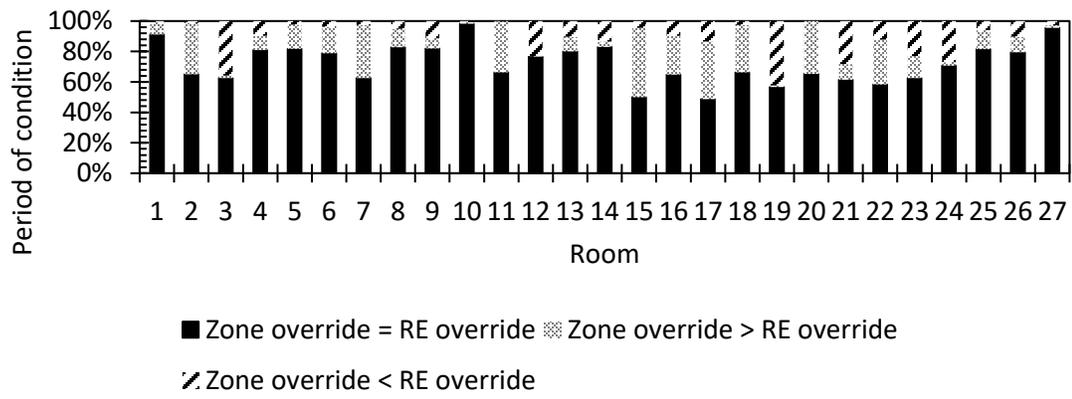
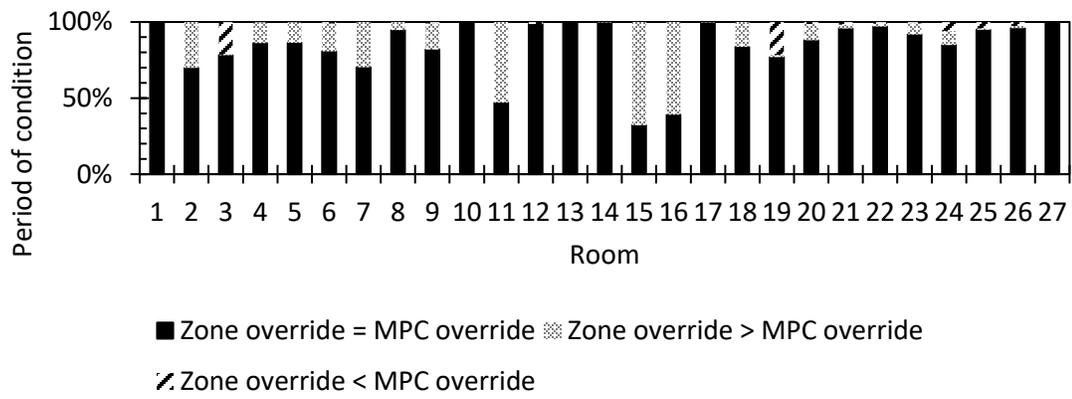
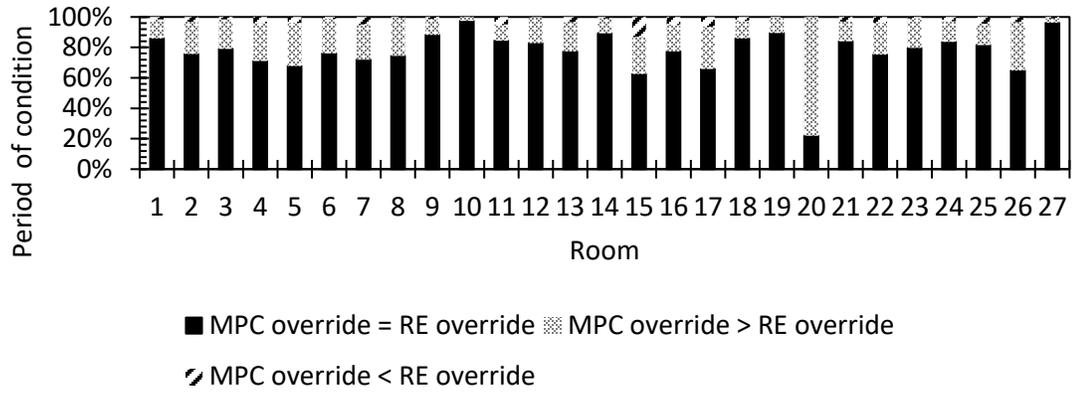
- When an override is equal to one, the radiant heaters and VAV units are commanded to operate using unmodified reactive control.
- When an override is equal to zero, the system is overridden to the minimum operating level (e.g., 0% valve position for radiant heaters and unoccupied mode for VAV units).

Zone-level overrides were compared to room-level overrides in Figure 6-12 and Figure 6-13 to illustrate the impact of modifications to zone control in testing weeks (where the average of room-level overrides in the zone had to be  $\geq 0.5$ ) on operation decisions. The unmodified zone-level reactive control used scaling and averaging of embedded room-level control loops to achieve a similar multi-room zone control effect in the absence of MPC or RE control. The scaling and averaging of room-level control loops in the reactive control algorithm yielded many heuristic (and often poorly tuned) parameters and rendered it relatively opaque.

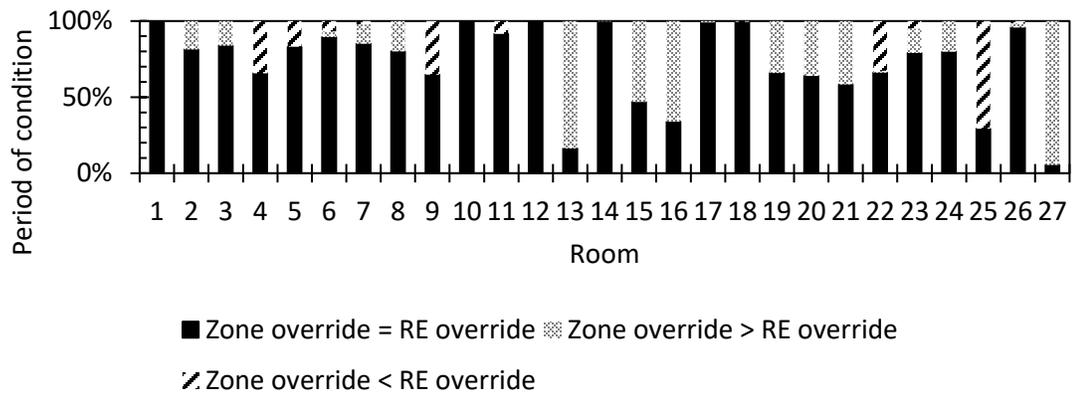
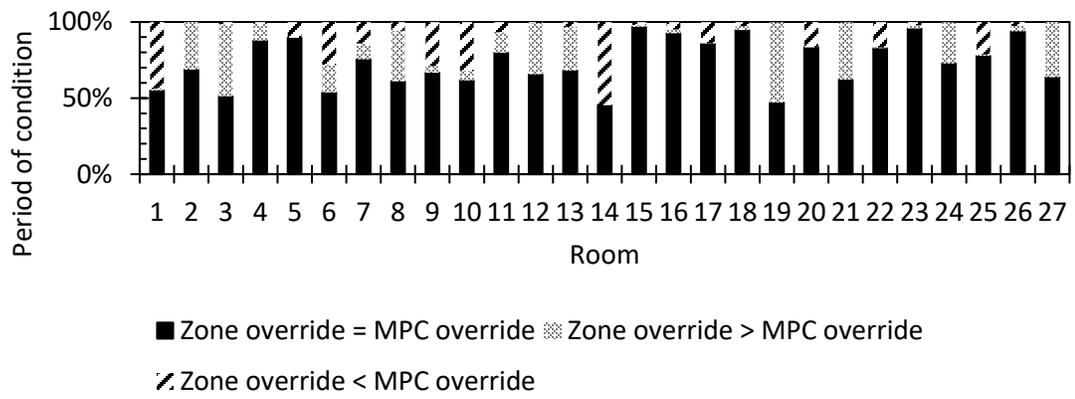
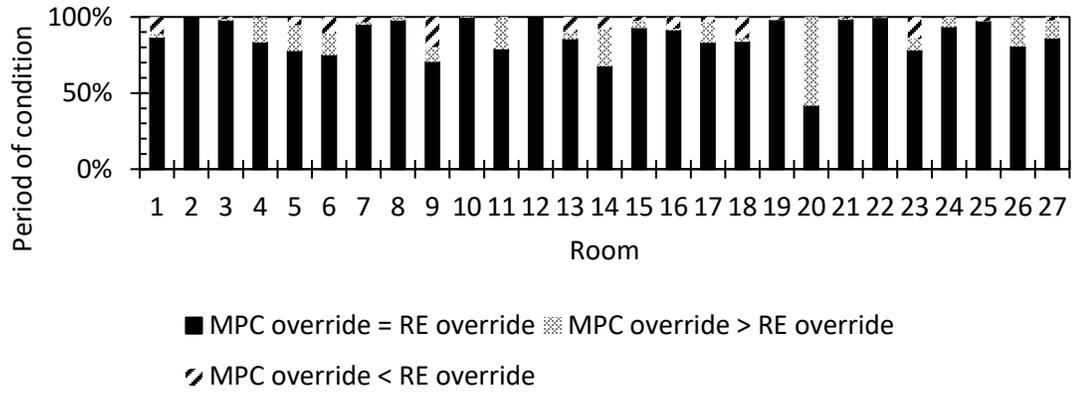
The RE overrides were retroactively calculated using the MPC algorithm input data and the selected decision trees for both seasons. If the RE and MPC overrides from the previously introduced algorithms both override for all operating time-steps (i.e., when the VAV unit was scheduled to be active) they will have a period of condition for MPC override = RE override of 100%. Consequently, the sum of the remaining possible

conditions (MPC override > RE override and MPC override < RE override) and MPC override = RE override must be one, or 100%. The same logic can be applied between the MPC and RE overrides and the zone state to determine the influence of each room on the zone. Low period of equal conditions (MPC override = RE override, zone override = MPC override, and zone override = RE override) can indicate excessive energy use or poor setpoint control due to zoning. The zoning of the 27 rooms is shown in Figure 6-14. Some rooms will inherently dominate zone control due to aggressive setpoints (e.g., 24°C in one room vs. 20°C in the other rooms in the heating season) or unique thermal conditions (e.g., an extra external wall). This is a condition that occurred in the reactive control algorithm and the MPC and RE algorithms.

Agreement can be seen between MPC and RE overrides in most cases, demonstrating the effectiveness of using one decision tree for all measured offices. The agreement between the two values could be attributed to either good model fit or relatively constant room cooling or heating requirements (i.e., the room almost always required heating or cooling). Room 20 had a low period of MPC overrides equaling RE overrides due missing data from a database malfunction. The MPC overrides were equal to the RE overrides for 77% and 86% of active VAV time-steps of the cooling and heating seasons respectively.



**Figure 6-12: Comparison of MPC and RE overrides (top), zone and MPC overrides (middle), and zone and RE overrides (bottom) in the cooling season. The period of condition compares the percent of operational time that the MPC override, RE override, and zone overrides meet the conditions in the legend of each bar graph. In ideal conditions the period of condition of the solid black “=” condition would be 100% for each room and each graph. Classification error and differences in room HVAC needs cause this to be less than 100%.**



**Figure 6-13: Comparison of MPC and RE overrides (top), zone and MPC overrides (middle), and zone and RE overrides (bottom) in the heating season. The period of condition compares the percent of operational time that the MPC override, RE override, and zone overrides meet the conditions in the legend of each bar graph. In ideal conditions the period of condition of the solid black “=” condition would be 100% for each room and each graph. Classification error and differences in room HVAC needs cause this to be less than 100%.**

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27
---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

**Figure 6-14: Zoning of rooms indicated by room numbers within the same box. Zones groupings varied in size in the measured building from one to four rooms per zone.**

Notable rooms that did not have much influence on zone-level overrides from room-level MPC and RE overrides in Figure 6-12 are rooms 11, 15, and 16 in the cooling season and rooms 9, 19, and 21 in the heating season. They had significantly different setpoints than the other rooms in the zone (e.g., 23 °C vs. 22 °C and 21.5 °C in one zone in the cooling season) and were often not able to simultaneously satisfy all spaces in the zone. The floor-level setpoint deviation is shown in Table 6-4 and is discussed further in the following section. The effect of room-level overrides on zone-level overrides was similar for both the MPC and RE algorithms. However, some rooms had significant increases or decreases. These differences are a result of the aggregation of room MPC control decisions to generate one decision tree for all of the rooms. The difference between the MPC and RE control overrides from the supervisory algorithm is another source of energy waste, but based on the energy use and setpoint deviation in Table 6-4 the inefficiency from using RE control is not large.

### **6.3.2 Energy Use and Indoor Air Temperature**

Energy use and IAT deviation from setpoint for the measured rooms was estimated at a weekly resolution and the floor-level to compare the effectiveness of the reactive control, MPC, and RE algorithms at saving energy and maintaining thermal comfort.

Aggregate energy use was only presented here for the 25 offices to maintain consistent space types and was normalized by office for each floor to account for the different number of offices (12 on one floor, 13 on the other). Aggregate office-averaged cooling MPC and RE energy savings were 42% and 27%, respectively. MPC outperformed RE in the cooling season, but the sample size was only one week for each control mode due to the AHU level cooling coil flow meters not being programmed by the contractor until the third week of testing. RE outperformed MPC in the heating season, where MPC and RE achieved aggregated office averaged MPC and RE energy savings of 18% and 33% respectively. Heating energy use for week 9 was added to Table 6-4 for context but was not used in total algorithm energy savings because it had a higher OAT than the other weeks.

The energy savings of the RE algorithm may have been higher than the MPC algorithm due to the significant setpoint deviation (and underheating) of the reactive control rooms in the first week of heating season testing. The reactive control rooms were not augmented by the test algorithm and operated as programmed by the control contractor during construction. Additionally, Floor A also had a higher setpoint deviation than the other test weeks in week 1, suggesting there was an issue with the hot water system at the building level that week. Floor B had higher setpoint deviations than Floor A throughout the heating season (particularly in the MPC weeks where the OAT and solar gains were lower), but this effect should be averaged out through the four weeks of alternating MPC and RE testing.

In the cooling season the average setpoint deviations for all offices over the entire testing period for the reactive control, MPC, and RE algorithms were 0.5 °C, 0.6 °C, and 0.2 °C respectively. In the heating season the average setpoint deviations for all offices over the entire testing period for the reactive control, MPC, and RE algorithms were -0.6 °C, -0.6 °C, and -0.2 °C respectively. Negative setpoint deviations indicate the IAT is below the setpoint. The setpoint deviation of the MPC algorithm may have been higher than the reactive control in the cooling season due to ambient conditions. The 0.1°C difference is still within the sensor bias error.

Some rooms had worse setpoint control than others due to setpoint differences between the rooms within the zone. Zones with three or four offices exhibited more setpoint deviation during active schedule hours, but the zones with two offices on the North-East corner of the building (and higher glazing ratios) exhibited the largest setpoint deviations due to a demand mismatch (one office had two external walls and the other had one). Weekly average setpoint deviations are provided for each room and algorithm in Table 6-4.

The heating season for Floor B generally also had higher setpoint deviation than the other offices (and even the same offices in the cooling season). In a single room test this would have been a significant issue, but because the test and control algorithms were each used to operate the floors for two weeks the bias from the setpoint deviation in Floor B was assumed to mostly cancel out. However, differences in ambient conditions for different weeks of testing could cause some imbalance in the setpoint and energy use bias. The similarity of ambient conditions during testing weeks is demonstrated in Figure 6-15, where the weekly average OAT, outdoor air relative humidity (OAH), and illuminance

near the window (as a proxy for solar gains that enter the rooms) are plotted with office-normalized test and reactive energy use.

**Table 6-4: Summary of weekly energy use of the floors and their corresponding algorithms. Average floor-level occupancy ratios and IAT setpoint deviation are also provided.**

Week	Algorithm - Floor A	Algorithm - Floor B	Fraction of time with at least one occupant - Floor A	Fraction of time with at least one occupant - Floor B	Average reactive control setpoint deviation (°C)	Average MPC setpoint deviation (°C)	Average RE control setpoint deviation (°C)	Average energy use per office – Floor A	Average energy use per office – Floor B	Energy savings of MPC/RE offices
Cooling 1	MPC	Control	0.031	0.034	0.43	0.51		N/A	N/A	N/A
Cooling 2	Control	MPC	0.030	0.036	0.56	0.56		N/A	N/A	N/A
Cooling 3	MPC	Control	0.020	0.023	0.60	0.60		N/A	N/A	N/A
Cooling 4	Control	MPC	0.033	0.041	0.54	0.60		0.61	0.35	42%
Cooling 5	RE	Control	0.018	0.034	0.31		0.44	0.65	0.52	-25%
Cooling 6	Control	RE	0.051	0.049	0.32		0.24	0.78	0.40	49%
Cooling 7	RE	Control	0.046	0.057	-0.03		-0.05	0.22	0.29	23%
Cooling 8	Control	RE	0.042	0.051	0.20		-0.01	0.45	0.22	51%
Heating 1	MPC	Control	0.022	0.049	-1.22	-0.43		0.72	0.63	-14%
Heating 2	Control	MPC	0.052	0.064	-0.14	-0.86		1.73	1.42	18%
Heating 3	MPC	Control	0.054	0.049	-0.83	-0.20		0.74	1.26	41%
Heating 4	Control	MPC	0.031	0.038	-0.02	-0.72		1.07	0.95	11%
Heating 5	RE	Control	0.046	0.059	-0.48		-0.06	0.78	1.21	36%
Heating 6	Control	RE	0.040	0.061	-0.04		-0.35	1.10	0.88	20%
Heating 7	RE	Control	0.035	0.056	-0.23		-0.03	0.34	0.68	50%
Heating 8	Control	RE	0.030	0.055	0.12		-0.40	0.72	0.54	24%
Heating 9	RE	Control	0.041	0.060	-0.26		-0.23	0.40	0.67	41%

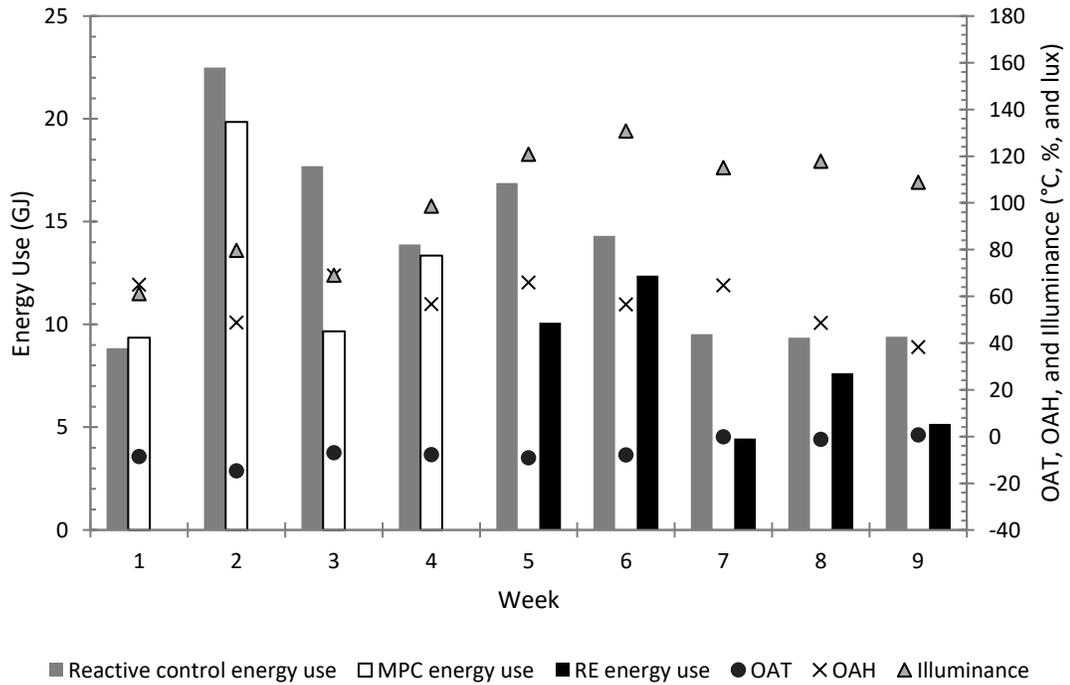
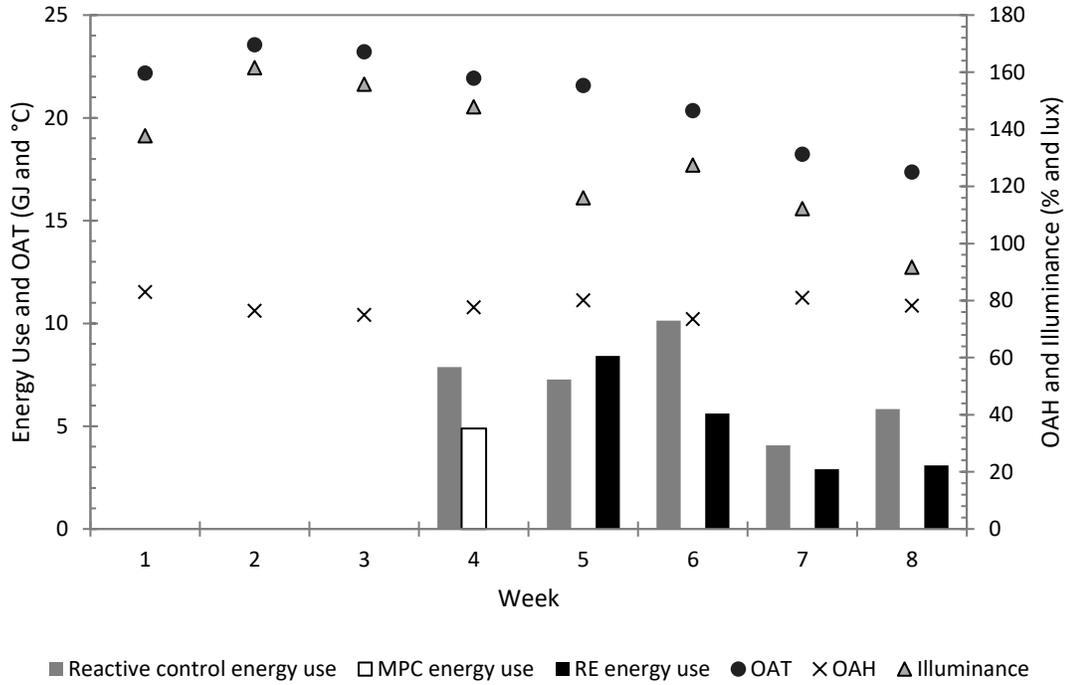


Figure 6-15: Weekly energy use and average ambient conditions for cooling season (top) and heating season (bottom)

### 6.3.3 Limitations of Study

This study has many limitations based on the scope, available hardware, and available information. These limitations include the following:

- The cost of electrical and thermal energy was assumed to be equal. Equation 6-3 would need to be adjusted to reflect the difference in value of energy used if electrical and thermal energy are assumed to have different costs.
- The measured building had a high outdoor air ratio (0.9). This may have increased the total energy savings due to the relatively high heating and cooling demands, but the percent savings should still be representative of the effectiveness of the algorithm.
- Other algorithms with different cost functions and  $R$  matrices could yield lower energy use or better occupant comfort. The MPC algorithm itself was not a target of optimization in this study. Several studies and tools exist that address this problem directly (Wetter, Haves et al. 2008, Hoes, Loonen et al. 2012, Afram and Janabi-Sharifi 2014, Žáčková, Váňa et al. 2014, Serale, Fiorentini et al. 2018, Wetter, Hu et al. 2018).
- A single decision tree for similar offices with the same façade boundary conditions was assumed to be adequate. Using one aggregated decision tree can result in sub-optimal control decisions for some rooms within the zone, but significantly simplifies the programming and troubleshooting effort.
- Energy was only quantified for two weeks of the four-week cooling season MPC period due to missing data in the BAS. Ideally, energy used at the cooling coils of

the AHUs would be quantified over the entire testing period. This limited testing period could reduce generalizability of the results when comparing the MPC and reactive control algorithms in the cooling season.

- Ambient conditions varying between the testing of MPC and RE algorithms could bias the results of the in-situ testing. The average outdoor air temperatures for the test weeks are provided in Table 5-2 and Figure 6-15, and remained relatively consistent. Effort was made in this study to mitigate the effect of ambient conditions on the test results by conducting testing as close to the middle of the heating and cooling seasons as possible.
- Occupancy fraction differences between the floors during testing weeks can also have an impact on the energy use results. These occupancy fractions are provided in Table 5-2. The impact of occupants on energy use is somewhat mitigated by the high outdoor air ratio of this building but could have a greater impact on other buildings.
- Floor-level energy use was normalized by office count (Floor A had 12 offices and Floor B had 13).
- Differing occupancy between test and control group floors was assumed to average out over four weeks of switching algorithms.
- Some sensors could have unverified bias or precision errors. This was mitigated via spot checks to verify reasonable values were returned during dynamic building operation.

- The selected decision trees developed here may not be applicable to other spaces. The process in this paper is presented in conjunction with the building specific results to facilitate the fitting of decision trees to other buildings.
- RE control is impossible currently without detailed MPC (and likely networked control from a remote computer) to train the decision tree.
- Due to the networked control from a remote computer, detailed MPC operation using linear time-invariant state-space models was limited in speed by the ability to read input, output, and variable data from the BAS controllers. The BAS controllers used in this study were new product and had better processor hardware to handle database queries than equipment installed in most existing buildings. Slower hardware could limit zone-level detailed MPC to updating at half-hour or hour intervals, limiting the ability to adjust to changing input conditions or prediction inaccuracy.
- The energy savings of the predictive control algorithm in the form proposed in this study are limited by the underlying reactive control algorithm installed by the control contractor. High performance sequences of operation could improve the reactive control and augment the energy savings of MPC and RE control.

#### **6.4 Summary**

This study described a process for generating simplified RE decision tree to perform control overrides and evaluated the energy use of both detailed MPC and the abstracted RE algorithm. It was found that aggregate office-averaged cooling MPC and RE energy savings were 42% and 27% respectively when compared to the reactive control over the

four-week testing period. RE outperformed MPC in the heating season, where MPC and RE achieved aggregated office averaged MPC and RE energy savings of 18% and 33% respectively. The unexpected increase in energy savings of RE compared to MPC in the heating season could have been due to underheating of the reactive control offices in the early (MPC) test weeks. The difference between the indoor air temperature and the setpoint was lower in the reactive control rooms in the later (RE) weeks of the heating season, increasing the baseline energy use. In the cooling season the average setpoint deviations (averaged over the entire test period of four weeks) for all offices over the entire testing period for the reactive control, MPC, and RE algorithms were 0.5 °C, 0.6 °C, and 0.2 °C respectively. In the heating season the average setpoint deviations for all offices over the entire testing period for the reactive control, MPC, and RE algorithms were -0.6 °C, -0.6 °C, and -0.2 °C respectively.

Overall, the MPC and RE algorithms both reduced energy use when compared to reactive control and did not have a negative impact on temperature setpoint control. Energy savings fluctuated weekly, and in some cases MPC saw more energy savings than RE (and vice versa). In this study MPC had a 1.5-hour control horizon, 15-minute calculation time for all 27 rooms, and required external calculations to actuate a BACnet BAS. Meanwhile, RE control had a five-minute control horizon, parallel room-level calculation time of less than one second and could directly actuate within the BACnet BAS. Due to the simplicity, tangibility, and comparable energy savings of the RE algorithm compare to the MPC algorithm, it is suggested that RE approaches be evaluated in-situ in future work. This approach relies heavily on inverse modelling, and the repeatability and reliability will

likely improve with a larger sample size.

## **Chapter 7: Conclusions and Future Work**

The ultimate goal of this thesis is to provide a reference for how zone control can be improved using proxy sensors and simple supervisory control overrides. Sensor selection and a simplified approach to MPC to reduce energy use and ease commercial implementation of supervisory algorithms for zone-level heating and cooling control were explored. Using the BAS infrastructure of an institutional building, the approaches were tested in-situ to verify performance and provide a reference for technicians and operators. The sensor selection and MPC approaches presented form a foundation from which the building control problem can be solved using proxy sensors and supervisory control.

The approaches presented in this thesis build upon over a century of innovation in automated heating and cooling of buildings. Major periods of historical technological breakthroughs have yielded fundamental changes in how buildings are controlled and operated. It is the author's belief at the time of writing this thesis that we are on the verge of another one of these breakthroughs. The technology sector is driving the global economy, and as such, there is an abundance of inexpensive processing and data storage that can be applied to building control. Combining the global growth of BACnet and data-driven approaches to energy use estimation (and control) can provide a promising path for improving building performance. This thesis attempts to harness this powerful combination to estimate energy use at the zone level, and use this estimation ability to provide a simple-to-implement approach to supervisory predictive control that achieves most of the energy savings of more detailed (but difficult to implement) algorithms.

Retrofits to measure energy use at the zone level are expensive and impractical. This thesis developed an approach to using proxy sensors to reduce the sensors required to estimate major energy end uses at the zone level using 27 rooms in the Health Sciences Building at Carleton University. The proxy sensing approach allows for temporary direct sensors to be used to inexpensively calibrate each proxy sensor model over a few days. Using this approach, the most valuable sensor from an energy use estimation perspective was the OAT sensor due to its ability to act as a proxy for estimating multiple energy end uses accurately over a validation period of two months. The next tier of useful sensors contained the flow and temperature sensors on the hydronic radiant heating system. After accounting for certain limitations (e.g., temporary measurements, model validation periods), proxy sensing can provide an inexpensive way to estimate zone level energy use and provide a baseline to compare supervisory control algorithms.

Before a wide-scale evaluation could be conducted, a simplified MPC algorithm had to be developed and monitored. The simplified MPC pilot study used a single node linear time-invariant state-space thermal model with five sensor inputs to characterize the indoor air temperature response of the space for an MPC cost function. Two approaches were developed and tested in a multi-occupant office, the former more detailed and cumbersome to implement than the latter. The first approach generated a binary override for the BAS based on the sum of the MPC predicted output over the control horizon and the setpoint deviation. The second approach used classification learning to generate a decision tree using a CART algorithm to emulate the MPC binary override from the sensor information and weather predictions for the office.

It was found that the MPC implemented in the pilot study achieved heating energy savings of 23% with respect to the reactive control installed during construction. Further, the simplified MPC approach was observed to achieve 94% of the heating energy savings of the MPC approach (21% with respect to reactive control) without using matrix inverse operations. Thermal comfort was assessed by occupant feedback via button presses to adjust setpoints and air temperature. No overrides were made over the entire testing period, and an average indoor air temperature of 21.1°C, with a standard deviation of 0.6°C and a minimum of 19.4°C were observed over all testing conditions. The developed decision tree for the simplified MPC approach contained six thresholds and required six linear time-invariant parameters from the original state-space model. Various paths and thresholds could be identified and justified using the simplified decision tree, and intuitive actions for the zone level BAS were identified. The simplified MPC algorithm provided a feasible retrofit solution for existing buildings where older control hardware, mathematical operator limitations in programming, and low network bandwidth would make detailed MPC application difficult or impossible without extensive hardware upgrades. The classification tree utilized for the pilot office could not be verified as a generic solution, and the logical sequence could vary from case to case. Thus, the proposed workflow was adapted to a larger office sample size of 27 rooms in the Health Sciences Building where testing was executed in parallel with a control group.

Prior to wide-scale implementation the supervisory algorithm underwent the following modifications: removal of the embedded indoor air temperature prediction model, additional RE testing time (compared to sMPC), parallel control (untested reactive

control) and test (MPC or RE control) groups, inclusion of both the heating and cooling season, modification of the override algorithm for zone control, improved automation of reading and writing between the detailed algorithm and the BAS. In principle, these modifications removed several parameters that were embedded within the supervisory algorithm that were required by the state-space model of the sMPC in the pilot implementation. It was assumed that the required information from the state-space model would appear in the extracted rules of the decision tree. It was found that aggregate office-averaged cooling MPC and RE energy savings were 42% and 27% respectively when compared to the reactive control over the four-week testing period. RE outperformed MPC in the heating season, where MPC and RE achieved aggregated office averaged MPC and RE energy savings of 18% and 33% respectively. The unexpected increase in energy savings of RE compared to MPC in the heating season could have been due to the underheating of the reactive control offices in the early (MPC) test weeks. The difference between the indoor air temperature and the setpoint was lower in the reactive control rooms in the later (RE) weeks of the heating season, increasing the baseline energy use. In the cooling season, the average setpoint deviations for all offices over the entire testing period for the reactive control, MPC, and RE algorithms were 0.5 °C, 0.6 °C, and 0.2 °C respectively. In the heating season the average setpoint deviations for all offices over the entire testing period for the reactive control, MPC, and RE algorithms were -0.6 °C, -0.6 °C, and -0.2 °C respectively. Thus, RE was able to produce comparable energy use savings compared to MPC with significantly less programming effort.

## 7.1 Contributions

The previously summarized body chapters of this thesis provide applied, transparent, and practitioner-understandable approaches to reducing sensor costs and improving control performance in commercial buildings. The main contributions of this thesis to the relevant bodies of work are explicitly stated in point form as follows:

- A BAS sensor and controller database specification that could facilitate wide-scale field implementation of experimental control algorithms. This specification acted as a framework to instrument the Health Sciences Building for the research conducted in Chapter 4 and Chapter 6.
- A methodology for quantitatively selecting sensors for indirect zone-level building energy use estimation.
- A field-tested approach to simplified MPC that achieves similar energy savings and occupant comfort to a detailed MPC algorithm and can be applied and understood by BAS technicians.
- An approach to RE that can simplify the implementation to a cluster of offices with similar thermal properties.
- A wide-scale field implementation of a simplified MPC algorithm in-situ that does not contain an embedded temperature model (i.e., rule extraction).

## 7.2 Future Work

The next steps to advance the work presented in this thesis include but are not limited to:

- Development of simple, accurate, and easy to use mobile sensing stations to record measurements of direct sensors for comparison to proxy sensor models.
- Extension of the proxy sensing approach presented in this thesis to different buildings and climates.
- Exploration of different proxy sensor inverse model forms to verify or improve energy use estimation accuracy (including grey box models).
- Inclusion of past sensor data as proxy sensor inputs to account for thermal lag.
- Evaluation and ranking of different proxy sensor estimated quantities (e.g., energy use). This thesis focused on energy use, but future work could focus on comfort metrics or other metrics important to building owners and operators.
- Inclusion of an energy cost model for the proxy sensors to demonstrate direct dollar values of proxy estimated primary energy use.
- Application of supervisory reinforcement learning (when a sufficient amount of data is collected to generate reasonable model archetypes) to tune archetype RE supervisory control models for specific building types and climates.
- Use of ASHRAE Guideline 13 and Guideline 36 for the baseline reactive control (ASHRAE 2000, ASHRAE 2018). The measured building in this study did not adhere to these new guidelines, thus the baseline control is not necessarily standardized due to the lack of a standard. Prior to the release of Guideline 36 in 2018, there was no consistently defined baseline control sequence for VAV systems

and sequences could vary significantly for every building.

- Expansion of control scenarios through simulation to develop base rules for different climates, building types, and MPC cost functions. Usage of the Open Building Control project from Lawrence Berkeley National Laboratory could automatically generate the GCL+ code for archetypal decision trees (Wetter, Hu et al. 2018).
- Testing of reinforcement learning to tune base models once archetypal decision trees are developed for RE control. To date, reinforcement learning has been evaluated in the literature for use with occupant-based control (Wei, Wang et al. 2017, Ding, Du et al. 2019, May 2019). The same adaptive tuning principles could be applied to decision trees.
- Application of the approach to simplified MPC presented in this thesis to pare down the input requirements of existing MPC algorithms in the absence of RE decision trees for direct control. Many inputs and parameters are not useful to the detailed algorithm and reducing the data acquisition and training effort could also improve the proliferation of detailed algorithms as well.
- Testing to determine the effect of sub-optimal zoning (as seen in some of the zones with 3-4 offices in Chapter 6) by comparing the energy savings and occupant comfort of single-room and multi-room zones. In the field study, sub-optimal zoning was found to contribute to higher energy use and higher setpoint deviation (a proxy for thermal comfort) for the offices within the zone.
- Adaptation and testing of the proxy sensor models and supervisory control

algorithms in other buildings and more types of buildings (e.g., residential buildings).

- Implementation of the proxy sensor models and supervisory control algorithms in a commercial BAS product.

## Appendices

### Appendix A : Manuscript Publishing Agreements

This appendix contains the publisher's agreements for the manuscripts contained within this thesis. The statements that pertain to the use of manuscripts in this thesis are underlined.

#### A.1 Elsevier (Journal of Building Engineering and Energy and Buildings) – Personal Use

Authors can use their articles, in full or in part, for a wide range of scholarly, non-commercial purposes as outlined below:

- Use by an author in the author's classroom teaching (including distribution of copies, paper or electronic).
- Distribution of copies (including through e-mail) to known research colleagues for their personal use (but not for Commercial Use).
- Inclusion in a thesis or dissertation (provided that this is not to be published commercially).
- Use in a subsequent compilation of the author's works.
- Extending the Article to book-length form.
- Preparation of other derivative works (but not for Commercial Use).
- Otherwise using or re-using portions or excerpts in other works.

These rights apply for all Elsevier authors who publish their article as either a subscription

article or an open access article. In all cases we require that all Elsevier authors always include a full acknowledgement and, if appropriate, a link to the final published version hosted on Science Direct.

## **A.2 Taylor & Francis (Science and Technology for the Built Environment) - Free Materials Definition**

All content that is made available by Us for free. This may include specific content made available: (i) to an individual through a periodic offer allowing free access for a limited period of time using either an access token or (for one session only) a session token; (ii) to the general public as free access content; or (iii) to an individual through referrer access from a relevant third party referring website.

If you register on the Site, this allows you to personalise the service we provide to you in relation to the Materials. Registration is required in order to access Free Materials using an access token and to purchase Premium Materials.

## **A.3 Taylor & Francis (Science and Technology for the Built Environment) - Usage rights for Free Materials**

We grant you a limited and revocable license to access and make personal use (including permission to print off one copy, and to download extracts) of any page containing Free Materials for your personal reference and you may draw the attention of others within your organization to the Free Materials.

You must not (i) modify the paper or digital copies of any Free Materials you have printed off or downloaded in any way; (ii) use any illustrations, photographs, video or audio

sequences, or any graphics separately from any accompanying text; and (iii) systematically download and store complete issues or volumes of any one journal.

If you print off, copy, or download any part of the Site or the Free Materials in breach of these Terms, your right to use the Site will cease immediately and you must, at our option, return or destroy any copies of the Free Materials you have made.

This license does not include any derivative use of the Site or the Materials, any collection and use of any product listings, descriptions, or prices; any downloading or copying of account information for the benefit of another merchant; or any use of data-mining, robots, or similar data-gathering and extraction tools. In addition, you may not use meta tags or any other "hidden text" utilizing our name or the name of any of our group companies without our express written consent.

You must not use any part of the Free Materials for commercial purposes without first obtaining a license in writing to do so from us.

Our status (and that of any identified contributors) as the authors of Free Materials on our site must always be acknowledged.

## References

- Afram, A. and F. Janabi-Sharifi (2014). "Theory and applications of HVAC control systems—A review of model predictive control (MPC)." Building and Environment **72**: 343-355.
- Afram, A. and F. Janabi-Sharifi (2015). "Black-box modeling of residential HVAC system and comparison of gray-box and black-box modeling methods." Energy and Buildings **94**: 121-149.
- Ahmad, A., M. Hassan, M. Abdullah, H. Rahman, F. Hussin, H. Abdullah and R. Saidur (2014). "A review on applications of ANN and SVM for building electrical energy consumption forecasting." Renewable and Sustainable Energy Reviews **33**(1): 102-109.
- Ahmad, M. W., M. Mourshed, D. Mundow, M. Sisinni and Y. Rezgui (2016). "Building energy metering and environmental monitoring – A state-of-the-art review and directions for future research." Energy and Buildings **120**: 85-102.
- Ali, A. S., Z. Zanzinger, D. Debose and B. Stephens (2016). "Open Source Building Science Sensors (OSBSS): A low-cost Arduino-based platform for long-term indoor environmental data collection." Building and Environment **100**: 114-126.
- Amasyali, K. and N. M. El-Gohary (2018). "A review of data-driven building energy consumption prediction studies." Renewable and Sustainable Energy Reviews **81**: 1192-1205.
- Anderson, M., M. Buehner, P. Young, D. Hittle, C. Anderson, J. Tu and D. Hodgson (2008). "MIMO robust control for HVAC systems." IEEE Transactions on Control Systems Technology **16**(3): 475-483.
- ASHRAE. (1999). "Vendor Gallery." Retrieved November 14, 2019, from [www.bacnet.org/Gallery/](http://www.bacnet.org/Gallery/).
- ASHRAE (2000). Guideline 13-2000 Specifying Direct Digital Control Systems, American Society of Heating, Refrigerating and Air-Conditioning Engineers.
- ASHRAE (2002). Guideline 14-2002 Measurement of Energy and Demand Savings, American Society of Heating, Refrigerating and Air-Conditioning Engineers.
- ASHRAE (2009). 2009 ASHRAE Handbook: Fundamentals, American Society of heating, refrigerating and air-conditioning engineers.
- ASHRAE (2013). 2013 ASHRAE Handbook: Fundamentals, American Society of heating, refrigerating and air-conditioning engineers.
- ASHRAE (2016). ASHRAE/IES Standard 90.1-2016. Energy Standard for Buildings Except Low-Rise Residential Buildings, American Society of Heating, Refrigerating and Air-Conditioning Engineers.
- ASHRAE (2017). ASHRAE Standard 55-2017. Thermal Environmental Conditions for Human Occupancy, American Society of Heating, Refrigerating and Air-Conditioning Engineers.
- ASHRAE (2017). Standard 135-2016, BACnet - A Data Communication Protocol for Building Automation and Control Networks, American Society of Heating, Refrigerating and Air-Conditioning Engineers.

- ASHRAE (2018). Guideline 36-2018 High Performance Sequences of Operation for HVAC Systems, American Society of Heating, Refrigerating and Air-Conditioning Engineers.
- Aswani, A., N. Master, J. Taneja, D. Culler and C. Tomlin (2012). "Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control." Proceedings of the IEEE **100**(1): 240-253.
- Athienitis, A. (1994). "Building thermal analysis." Electronic Book Mathcad, MathSoft, Boston, USA.
- Bacher, P. and H. Madsen (2011). "Identifying suitable models for the heat dynamics of buildings." Energy and Buildings **43**(7): 1511-1522.
- BACnet. (2019). "Vendor identifiers " Retrieved November 20th, 2019, from <http://www.bacnet.org/VendorID/BACnet%20Vendor%20IDs.htm>.
- Baeg, S.-H., J.-H. Park, J. Koh, K.-W. Park and M.-H. Baeg (2007). Building a smart home environment for service robots based on RFID and sensor networks. 2007 International Conference on Control, Automation and Systems, IEEE.
- Bajer, M. (2018). IoT for smart buildings-long awaited revolution or lean evolution. 2018 IEEE 6th International Conference on Future Internet of Things and Cloud (FiCloud), IEEE.
- Bălan, R., J. Cooper, K.-M. Chao, S. Stan and R. Donca (2011). "Parameter identification and model based predictive control of temperature inside a house." Energy and Buildings **43**(2-3): 748-758.
- Bengea, S., A. Kelman, F. Borrelli, R. Taylor and S. Narayanan (2012). Model predictive control for mid-size commercial building HVAC: Implementation, results and energy savings. Second international conference on building energy and environment.
- Biswas, K., S. S. Shrestha, M. S. Bhandari and A. O. Desjarlais (2016). "Insulation materials for commercial buildings in North America: An assessment of lifetime energy and environmental impacts." Energy and Buildings **112**: 256-269.
- Boss (2016). Case Study: City of Pittsburgh Plug Load Energy Management. Pittsburgh, PA, Boss Controls.
- Braun, J. E. and N. Chaturvedi (2002). "An inverse gray-box model for transient building load prediction." HVAC&R Research **8**(1): 73-99.
- BSRIA (2018). Global Market Research Report.
- Bursill, J., W. O'Brien and I. Beausoleil-Morrison (2018). An approach to virtual sensing of cooling systems using state-space inverse modelling and gradient descent. eSim, Montreal.
- Bursill, J., W. O'Brien and I. Beausoleil-Morrison (2019). "Experimental application of classification learning to generate simplified model predictive controls for a shared office heating system." Science and Technology for the Built Environment **25**(5): 615-628.
- Bursill, J., W. O'Brien and I. Beausoleil-Morrison (2018). "Experimental application of classification learning to generate simplified model predictive controls for a shared office heating system." Science and Technology for the Built Environment(just-accepted): 1-38.

- Bursill, J., W. O'Brien and I. Beausoleil-Morrison (2019). Development and Analysis of Simplified Control-oriented Models for a Group of Institutional Offices. Building Simulation 2019 Rome **16**.
- Bursill, J., W. O'Brien and I. Beausoleil-Morrison (2020). "Multi-zone field study of rule extraction control to simplify implementation of predictive control to reduce building energy use." Energy and Buildings **Accepted for publication**.
- Candanedo, J. (2018). CanMETEO Background Information.
- Candanedo, J. and A. Athienitis (2011). "Predictive control of radiant floor heating and solar-source heat pump operation in a solar house." HVAC&R Research **17(3)**: 235-256.
- Carrier. (2019). "About Willis Carrier." Retrieved December 8, 2019, from <http://www.williscarrier.com/about.php>.
- Carrier, W. H. (1926). Refrigerating system and method of refrigeration, US Patents.
- Causone, F., S. Carlucci, M. Ferrando, A. Marchenko and S. Erba (2019). "A data-driven procedure to model occupancy and occupant-related electric load profiles in residential buildings for energy simulation." Energy and Buildings: 109342.
- Chae, Y. T., R. Horesh, Y. Hwang and Y. M. Lee (2016). "Artificial neural network model for forecasting sub-hourly electricity usage in commercial buildings." Energy and Buildings **111**: 184-194.
- Chen, K.-Y., L.-S. Chen, M.-C. Chen and C.-L. Lee (2011). "Using SVM based method for equipment fault detection in a thermal power plant." Computers in industry **62(1)**: 42-50.
- Chen, Q. (2009). "Ventilation performance prediction for buildings: A method overview and recent applications." Building and Environment **44(4)**: 848-858.
- Chen, T. and A. Athienitis (2003). "Investigation of practical issues in building thermal parameter estimation." Building and Environment **38(8)**: 1027-1038.
- Cigler, J., J. Siroky, M. Korda and C. Jones (2013). On the selection of the most appropriate MPC problem formulation for buildings. 11th REHVA World Congress CLIMA 2013.
- Coffey, B. (2013). "Approximating model predictive control with existing building simulation tools and offline optimization." Journal of Building Performance Simulation **6(3)**: 220-235.
- Coffey, B., F. Haghighat, E. Morofsky and E. Kutrowski (2010). "A software framework for model predictive control with GenOpt." Energy and Buildings **42(7)**: 1084-1092.
- Cole, W. J., K. M. Powell, E. T. Hale and T. F. Edgar (2014). "Reduced-order residential home modeling for model predictive control." Energy and Buildings **74**: 69-77.
- Cotrufo, N. and R. Zmeureanu (2016). A new algorithm for sensor verification and correction in air handling units. eSim, Hamilton.
- Cunningham, M. (1981). "Measurement errors and instrument inaccuracies." Journal of Physics E: Scientific Instruments **14(8)**: 901.
- Cuzzocrea, A., I.-Y. Song and K. C. Davis (2011). Analytics over large-scale multidimensional data: the big data revolution! Proceedings of the ACM 14th international workshop on Data Warehousing and OLAP, ACM.

- D'Oca, S. and T. Hong (2014). "A data-mining approach to discover patterns of window opening and closing behavior in offices." Building and Environment **82**: 726-739.
- Deng, H., D. Fannon and M. J. Eckelman (2018). "Predictive modeling for US commercial building energy use: A comparison of existing statistical and machine learning algorithms using CBECS microdata." Energy and Buildings **163**: 34-43.
- Deng, Z. and Q. Chen (2019). "Simulating the impact of occupant behavior on energy use of HVAC systems by implementing a behavioral artificial neural network model." Energy and Buildings **198**: 216-227.
- Dermardiros, V., Y. Chen and A. K. Athienitis (2015). "Modelling of an Active PCM Thermal Energy Storage for Control Applications." Energy Procedia **78**: 1690-1695.
- Dietrich, D., D. Bruckner, G. Zucker and P. Palensky (2010). "Communication and computation in buildings: A short introduction and overview." IEEE transactions on industrial electronics **57**(11): 3577-3584.
- Ding, X., W. Du and A. Cerpa (2019). "OCTOPUS: Deep Reinforcement Learning for Holistic Smart Building Control." ACM BuildSys.
- Dong, B., Z. Li, S. M. M. Rahman and R. Vega (2016). "A hybrid model approach for forecasting future residential electricity consumption." Energy and Buildings **117**: 341-351.
- Dong, B., Z. O'Neill and Z. Li (2014). "A BIM-enabled information infrastructure for building energy Fault Detection and Diagnostics." Automation in Construction **44**: 197-211.
- Dong, B., V. Prakash, F. Feng and Z. O'Neill (2019). "A Review of Smart Building Sensing System for Better Indoor Environment Control." Energy and Buildings.
- Duin, R., P. Juszczak, P. Paclik, E. Pekalska, D. De Ridder, D. Tax and S. Verzakov (2000). "A matlab toolbox for pattern recognition." PRTools version 3: 109-111.
- Eltgrowth, D. R. (2009). "Best evidence and the Wayback Machine: toward a workable authentication standard for archived Internet evidence." Fordham L. Rev. **78**: 181.
- Epstein, J. (2015). "A brief history of building automation and controls." Retrieved December 8, 2019, from <http://www.harcourtbrown.com/a-brief-history-of-building-automation-and-controls/>.
- Fanger, P. O. (1970). "Thermal comfort. Analysis and applications in environmental engineering."
- Ferrari, S. and M. Beccali (2017). "Energy-environmental and cost assessment of a set of strategies for retrofitting a public building toward nearly zero-energy building target." Sustainable cities and society **32**: 226-234.
- Florita, A. and G. Henze (2009). "Comparison of Short-Term Weather Forecasting Models for Model Predictive Control." HVAC&R Research **15**(5): 835-853.
- Fu, Y., Z. Li, H. Zhang and P. Xu (2015). "Using Support Vector Machine to Predict Next Day Electricity Load of Public Buildings with Sub-metering Devices." Procedia Engineering **121**: 1016-1022.
- Fuertes, G. and S. Schiavon (2014). "Plug load energy analysis: The role of plug loads in LEED certification and energy modeling." Energy and Buildings **76**: 328-335.
- Gilani, S. and W. O'Brien (2018). "A preliminary study of occupants' use of manual

- lighting controls in private offices: A case study." Energy and Buildings **159**: 572-586.
- Gilliland, D. (2016). A Field Evaluation of Cloud-Based Predictive HVAC Control Software—The Next Step in Energy Efficiency, or Marketing Mumbo Jumbo? 2016 ACEEE Summer Study on Energy Efficiency in Buildings. Washington, DC: 1-13.
- Gouda, M., S. Danaher and C. Underwood (2000). "Low-order model for the simulation of a building and its heating system." Building Services Engineering Research and Technology **21**(3): 199-208.
- Gouda, M., S. Danaher and C. Underwood (2002). "Building thermal model reduction using nonlinear constrained optimization." Building and Environment **37**(12): 1255-1265.
- Goyal, S., P. Barooah and T. Middelkoop (2015). "Experimental study of occupancy-based control of HVAC zones." Applied Energy **140**: 75-84.
- Grant, C. K., P (2016). Smart Thermostats: An Alternative to Load Control Switches? Arlington, VA.
- GSA (2012). Advanced Power Strips. Washington, DC, General Services Administration.
- Gunay, B., W. O'Brien and I. Beausoleil-Morrison (2016). "Control-oriented inverse modeling of the thermal characteristics in an office." Science and Technology for the Built Environment **22**(5): 586-605.
- Gunay, H. B. (2016). Improving energy efficiency in office buildings through adaptive control of the indoor climate, Carleton University.
- Gunay, H. B., J. Bursill, B. Huchuk, W. O'Brien and I. Beausoleil-Morrison (2014). "Shortest-prediction-horizon model-based predictive control for individual offices." Building and Environment **82**: 408-419.
- Gunay, H. B., W. O'Brien and I. Beausoleil-Morrison (2013). "A critical review of observation studies, modeling, and simulation of adaptive occupant behaviors in offices." Building and Environment **70**: 31-47.
- Gunay, H. B., W. O'Brien and I. Beausoleil-Morrison (2015). "Development of an occupancy learning algorithm for terminal heating and cooling units." Building and Environment **93**: 71-85.
- Gunay, H. B., W. O'Brien, I. Beausoleil-Morrison and J. Bursill (2016). "Implementation of an adaptive occupancy and building learning temperature setback algorithm." ASHRAE Transactions **122**: 179.
- Gunay, H. B., W. O'Brien, I. Beausoleil-Morrison and J. Bursill (2018). "Development and implementation of a thermostat learning algorithm." Science and Technology for the Built Environment **24**(1): 43-56.
- Gunay, H. B., W. O'Brien, I. Beausoleil-Morrison and S. Gilani (2016). "Modeling plug-in equipment load patterns in private office spaces." Energy and Buildings **121**: 234-249.
- Gwerder, M., D. Gyalistras, C. Sagerschnig, R. Smith and D. Sturzenegger (2013). "Final Report: Use of Weather and Occupancy Forecasts for Optimal Building Climate Control—Part II: Demonstration (OptiControl-II)." Automatic Control Laboratory, ETH Zurich: Zug, Switzerland: 156.

- Hagan, M. T., H. B. Demuth, M. H. Beale and O. De Jesús (1996). Neural network design, Pws Pub. Boston.
- Hargadon, A. B. and Y. Douglas (2001). "When innovations meet institutions: Edison and the design of the electric light." Administrative science quarterly **46**(3): 476-501.
- Hecht-Nielsen, R. (1992). Theory of the backpropagation neural network. Neural networks for perception, Elsevier: 65-93.
- Henze, G. P., M. J. Brandemuehl, C. Felsmann, H. Cheng, A. R. Florita and C. E. Waters (2008). "Optimization of Building Thermal Mass Control in the Presence of Energy and Demand Charges." ASHRAE Transactions **114**(2).
- Henze, G. P., C. Felsmann and G. Knabe (2004). "Evaluation of optimal control for active and passive building thermal storage." International Journal of Thermal Sciences **43**(2): 173-183.
- Hilliard, T., L. Swan, M. Kavgić, Z. Qin and P. Lingras (2016). "Development of a whole building model predictive control strategy for a LEED silver community college." Energy and Buildings **111**: 224-232.
- Hoes, P., R. Loonen, M. Trcka and J. Hensen (2012). "Performance prediction of advanced building controls in the design phase using ESP-r, BCVTB and MATLAB." BSO12 (IBPSA-England). Loughborough, UK.
- Incropera, F. P., A. S. Lavine, T. L. Bergman and D. P. DeWitt (2007). Fundamentals of heat and mass transfer, Wiley.
- InvisiShade. (2016). "FAQ." Retrieved December 11, 2019, from [www.invisishade.com/](http://www.invisishade.com/).
- ISO (2012). Information technology — Control network protocol — Part 1: Protocol stack, International Organization for Standardization
- Jain, R. K., K. M. Smith, P. J. Culligan and J. E. Taylor (2014). "Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy." Applied Energy **123**: 168-178.
- Johnson, W. S. (1883). Electric tele-thermoscope, US Patent
- Jung, H. C., J. S. Kim and H. Heo (2015). "Prediction of building energy consumption using an improved real coded genetic algorithm based least squares support vector machine approach." Energy and Buildings **90**: 76-84.
- Kalman, R. E. (1960). "A new approach to linear filtering and prediction problems." Journal of Basic Engineering **82**(1): 35-45.
- Kalogirou, S. A. (2000). "Applications of artificial neural-networks for energy systems." Applied energy **67**(1): 17-35.
- Karava, P., A. Athienitis, T. Stathopoulos and E. Mouriki (2012). "Experimental study of the thermal performance of a large institutional building with mixed-mode cooling and hybrid ventilation." Building and Environment **57**: 313-326.
- Katipamula, S. and M. R. Brambley (2005). "Review Article: Methods for Fault Detection, Diagnostics, and Prognostics for Building Systems—A Review, Part II." HVAC&R Research **11**(2): 169-187.
- Killian, M. and M. Kozek (2016). "Ten questions concerning model predictive control for energy efficient buildings." Building and Environment **105**: 403-412.
- King, J. P., Christopher (2017). Smart Buildings: Using Smart Technology to Save Energy

- in Existing Buildings. Washington, DC, American Council for an Energy-Efficient Economy: 55.
- Kintner-Meyer, M., M. R. Brambley, T. A. Carlon and N. N. Bauman (2002). Wireless sensors: technology and cost-savings for commercial buildings. Teaming for Efficiency: Proceedings of the 2002 ACEEE Summer Study on Energy Efficiency in Buildings.
- Kitchin, R. (2014). The data revolution: Big data, open data, data infrastructures and their consequences, Sage.
- Koller, D. and N. Friedman (2009). Probabilistic graphical models: principles and techniques, MIT press.
- Konar, P. and P. Chattopadhyay (2011). "Bearing fault detection of induction motor using wavelet and Support Vector Machines (SVMs)." Applied Soft Computing **11**(6): 4203-4211.
- Kottek, M., J. Grieser, C. Beck, B. Rudolf and F. Rubel (2006). "World map of the Köppen-Geiger climate classification updated." Meteorologische Zeitschrift **15**(3): 259-263.
- Kramer, R., J. van Schijndel and H. Schellen (2012). "Simplified thermal and hygric building models: A literature review." Frontiers of Architectural Research **1**(4): 318-325.
- Kumar, P., C. Martani, L. Morawska, L. Norford, R. Choudhary, M. Bell and M. Leach (2016). "Indoor air quality and energy management through real-time sensing in commercial buildings." Energy and Buildings **111**: 145-153.
- Lee, K.-h. and J. E. Braun (2008). "Model-based demand-limiting control of building thermal mass." Building and Environment **43**(10): 1633-1646.
- Li, B. and L. Xia (2015). A multi-grid reinforcement learning method for energy conservation and comfort of HVAC in buildings. 2015 IEEE International Conference on Automation Science and Engineering (CASE).
- Li, H., D. Yu and J. E. Braun (2011). "A review of virtual sensing technology and application in building systems." HVAC&R Research **17**(5): 619-645.
- Li, X. and J. Wen (2014). "Review of building energy modeling for control and operation." Renewable and Sustainable Energy Reviews **37**: 517-537.
- Li, X., J. Wen, R. Liu and X. Zhou (2016). "Commercial building cooling energy forecasting using proactive system identification: A whole building experiment study." Science and Technology for the Built Environment **22**(6): 674-691.
- Ljung, L. (1995). "System Identification Toolbox for use with MATLAB."
- Lobato, C., S. Pless, M. Sheppy and P. Torcellini (2011). Reducing plug and process loads for a large scale, low energy office building: NREL's research support facility, National Renewable Energy Lab.(NREL), Golden, CO (United States).
- Lutron. (2014). "Lutron Energy Savings Claims." from [www.lutron.com/TechnicalDocumentLibrary/Lutron\\_Energy\\_Savings\\_Claims.pdf](http://www.lutron.com/TechnicalDocumentLibrary/Lutron_Energy_Savings_Claims.pdf).
- Ma, J., J. Qin, T. Salsbury and P. Xu (2012). "Demand reduction in building energy systems based on economic model predictive control." Chemical Engineering Science **67**(1): 92-100.
- Ma, Y., A. Kelman, A. Daly and F. Borrelli (2012). "Predictive control for energy efficient

- buildings with thermal storage: Modeling, stimulation, and experiments." IEEE control systems magazine **32**(1): 44-64.
- Maddalena, E. T., Y. Lian and C. N. Jones (2020). "Data-driven methods for building control — A review and promising future directions." Control Engineering Practice **95**: 104211.
- Maji, S., A. C. Berg and J. Malik (2008). Classification using intersection kernel support vector machines is efficient. Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE.
- May-Ostendorp, P., G. P. Henze, C. D. Corbin, B. Rajagopalan and C. Felsmann (2011). "Model-predictive control of mixed-mode buildings with rule extraction." Building and Environment **46**(2): 428-437.
- May-Ostendorp, P. T. (2012). Offline model predictive control of mixed mode buildings for near-optimal supervisory control strategy development, University of Colorado at Boulder.
- May, R. (2019). The reinforcement learning method: A feasible and sustainable control strategy for efficient occupant-centred building operation in smart cities, Dalarna University.
- May, R., G. Dandy and H. Maier (2011). "Review of input variable selection methods for artificial neural networks." Artificial neural networks-methodological advances and biomedical applications **10**: 16004.
- McDonald, E., R. Zmeureanu and D. Giguère (2014). Virtual flow meter for chilled water loops in existing buildings. Conference eSim.
- McNeil, A., C. Kohler, E. S. Lee and S. Selkowitz (2017). High Performance Building Mockup in FLEXLAB, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).
- Minoli, D., K. Sohraby and B. Occhiogrosso (2017). "IoT Considerations, Requirements, and Architectures for Smart Buildings—Energy Optimization and Next-Generation Building Management Systems." IEEE Internet of Things Journal **4**(1): 269-283.
- Mirieli, J., L. Serres and A. Trombe (2002). "Radiant ceiling panel heating-cooling systems: experimental and simulated study of the performances, thermal comfort and energy consumptions." Applied Thermal Engineering **22**(16): 1861-1873.
- MODBUS (2018). MODBUS/TCP Security, Modbus Organization.
- Moon, J. W. and J.-J. Kim (2010). "ANN-based thermal control models for residential buildings." Building and Environment **45**(7): 1612-1625.
- Nesler, C. G. (1985). Experiences in applying adaptive control to thermal processes in buildings. 1985 American Control Conference, IEEE.
- North, Z. (2018). "Things to Keep in Mind When Considering Workplace Sensors." Retrieved July 23, 2019, from <https://www.officespacesoftware.com/blog/things-to-keep-in-mind-when-considering-workplace-sensors>.
- NRCan (2017). National energy code of Canada for buildings, Canadian Commission on Building and Fire Codes; National Research Council of Canada.
- Oldewurtel, F., A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann and M. Morari (2012). "Use of model predictive control and weather forecasts for energy efficient building climate control." Energy and Buildings **45**: 15-27.

- Page, E. R. (2013). Movable illuminance sensors for fixture light sources, Google Patents.
- Peksenar, F. (2013). "Designing Smart Thermostats that Provide Wireless Connectivity."
- Penman, J. (1990). "Second order system identification in the thermal response of a working school." Building and Environment **25**(2): 105-110.
- Peters, A. (2014). "Smart Sensors For Offices That Do It All: Save Energy, Track Employees, And Create A Happier Workplace." from <https://www.fastcompany.com/3034225/smart-sensors-for-offices-that-do-it-all-save-energy-track-employees-and-create-a-happier-wo>.
- Popescu, D., S. Bienert, C. Schützenhofer and R. Boazu (2012). "Impact of energy efficiency measures on the economic value of buildings." Applied Energy **89**(1): 454-463.
- Privara, S., J. Široký, L. Ferkl and J. Cigler (2011). "Model predictive control of a building heating system: The first experience." Energy and Buildings **43**(2): 564-572.
- Purdy, J. and I. Beausoleil-Morrison (2001). The significant factors in modelling residential buildings. Canmet Center for Technology, 7th International IBPSA Conf., Rio de Janeiro.
- Rahimi, M. and A. Sabernaemi (2010). "Experimental study of radiation and free convection in an enclosure with a radiant ceiling heating system." Energy and Buildings **42**(11): 2077-2082.
- Reddy, T. A. (2011). Applied data analysis and modeling for energy engineers and scientists, Springer Science & Business Media.
- Riederer, P., D. Marchio and J. C. Visier (2002). "Influence of sensor position in building thermal control: criteria for zone models." Energy and buildings **34**(8): 785-798.
- Roth, K. W., P. Llana, D. Westphalen, L. Quartararo and M. Y. Feng (2006). "Advanced controls for commercial buildings: Barriers and energy savings potential." Energy engineering **103**(6): 6-36.
- Rouchier, S. (2018). "Solving inverse problems in building physics: An overview of guidelines for a careful and optimal use of data." Energy and Buildings **166**: 178-195.
- Royapoor, M., A. Antony and T. Roskilly (2018). "A review of building climate and plant controls, and a survey of industry perspectives." Energy and Buildings **158**: 453-465.
- Ruparathna, R., K. Hewage and R. Sadiq (2016). "Improving the energy efficiency of the existing building stock: A critical review of commercial and institutional buildings." Renewable and Sustainable Energy Reviews **53**: 1032-1045.
- Ryzhov, A., H. Ouerdane, E. Gryazina, A. Bischi and K. Turitsyn (2019). "Model predictive control of indoor microclimate: Existing building stock comfort improvement." Energy conversion and management **179**: 219-228.
- Saberi Derakhtenjani, A., J. A. Candanedo, Y. Chen, V. R. Dehkordi and A. K. Athienitis (2015). "Modeling approaches for the characterization of building thermal dynamics and model-based control: A case study." Science and Technology for the Built Environment **21**(6): 824-836.
- SageGlass. (2016). "SageGlass." Retrieved December 11, 2019, from [www.sageglass.com/products](http://www.sageglass.com/products).

- Salakij, S., N. Yu, S. Paolucci and P. Antsaklis (2016). "Model-Based Predictive Control for building energy management. I: Energy modeling and optimal control." Energy and Buildings **133**: 345-358.
- Salsbury, T. I. (2005). "A survey of control technologies in the building automation industry." IFAC Proceedings Volumes **38**(1): 90-100.
- Schachinger, D. and W. Kastner (2016). Semantics for smart control of building automation. 2016 IEEE 25th International Symposium on Industrial Electronics (ISIE), IEEE.
- Serale, G., M. Fiorentini, A. Capozzoli, D. Bernardini and A. Bemporad (2018). "Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities." Energies **11**(3): 631.
- Shahin, M. A., H. R. Maier and M. B. Jaksa (2004). "Data division for developing neural networks applied to geotechnical engineering." Journal of Computing in Civil Engineering **18**(2): 105-114.
- Shaikh, P. H., N. B. M. Nor, P. Nallagownden, I. Elamvazuthi and T. Ibrahim (2014). "A review on optimized control systems for building energy and comfort management of smart sustainable buildings." Renewable and Sustainable Energy Reviews **34**: 409-429.
- Shoemaker, S. (2015). "Innovative, Lower Cost Sensors and Controls Yield Better Energy Efficiency." from [www.ornl.gov/news/innovative-lower-cost-sensors-and-controls-yield-better-energy-efficiency](http://www.ornl.gov/news/innovative-lower-cost-sensors-and-controls-yield-better-energy-efficiency).
- Široký, J., F. Oldewurtel, J. Cigler and S. Prívára (2011). "Experimental analysis of model predictive control for an energy efficient building heating system." Applied Energy **88**(9): 3079-3087.
- Smith, R. C. (2013). Uncertainty quantification: theory, implementation, and applications, Siam.
- Solomatine, D. P. and A. Ostfeld (2008). "Data-driven modelling: some past experiences and new approaches." Journal of Hydroinformatics **10**(1): 3-22.
- Song, L., I.-s. Joo and G. Wang (2012). "Uncertainty analysis of a virtual water flow measurement in building energy consumption monitoring." HVAC&R Research **18**(5): 997-1010.
- Spengler, J. D., J. M. Samet and J. F. McCarthy (2001). "Indoor air quality handbook."
- Stluka, P., G. Parthasarathy, S. Gabel and T. Samad (2018). Architectures and Algorithms for Building Automation—An Industry View. Intelligent Building Control Systems, Springer: 11-43.
- Sutton, R. S. and A. G. Barto (2018). Reinforcement learning: An introduction, MIT press.
- Szepesvári, C. (2010). "Algorithms for reinforcement learning." Synthesis lectures on artificial intelligence and machine learning **4**(1): 1-103.
- Thieblemont, H., F. Haghghat, R. Ooka and A. Moreau (2017). "Predictive control strategies based on weather forecast in buildings with energy storage system: A review of the state-of-the art." Energy and Buildings **153**: 485-500.
- Thompson, M. (1982). "Regression methods in the comparison of accuracy." Analyst **107**(1279): 1169-1180.

- Touretzky, C. R. and M. Baldea (2014). "Integrating scheduling and control for economic MPC of buildings with energy storage." Journal of Process Control **24**(8): 1292-1300.
- Tsanas, A. and A. Xifara (2012). "Accurate quantitative estimation of energy performance of residential buildings using statistical machine learning tools." Energy and Buildings **49**: 560-567.
- Tzempelikos, A. and H. Shen (2013). "Comparative control strategies for roller shades with respect to daylighting and energy performance." Building and Environment **67**: 179-192.
- Van Der Heijden, F., R. P. Duin, D. De Ridder and D. M. Tax (2005). Classification, parameter estimation and state estimation: an engineering approach using MATLAB, John Wiley & Sons.
- Wagner, A., T. Lützkendorf, K. Voss, G. Spars, A. Maas and S. Herkel (2014). "Performance analysis of commercial buildings—results and experiences from the german demonstration program ‘energy optimized building (EnOB)’." Energy and Buildings **68**: 634-638.
- Wang, B. and X. Xia (2015). "Optimal maintenance planning for building energy efficiency retrofitting from optimization and control system perspectives." Energy and Buildings **96**: 299-308.
- Wang, L. (2009). Model predictive control system design and implementation using MATLAB®, Springer Science & Business Media.
- Wang, S. and Z. Ma (2008). "Supervisory and optimal control of building HVAC systems: A review." HVAC&R Research **14**(1): 3-32.
- Wang, Z., X. Duan, D. Yi, J. Yao and D. Gu (2016). Measurement data modeling and parameter estimation, CRC Press.
- Wang, Z. and C. J. Ong (2017). "Distributed Model Predictive Control of linear discrete-time systems with local and global constraints." Automatica **81**: 184-195.
- Wei, T., Y. Wang and Q. Zhu (2017). Deep reinforcement learning for building hvac control. Proceedings of the 54th Annual Design Automation Conference 2017, ACM.
- Wei, Y., X. Zhang, Y. Shi, L. Xia, S. Pan, J. Wu, M. Han and X. Zhao (2018). "A review of data-driven approaches for prediction and classification of building energy consumption." Renewable and Sustainable Energy Reviews **82**: 1027-1047.
- Wetter, M., P. Haves and B. Coffey (2008). Building controls virtual test bed, Lawrence Berkeley National Laboratory.
- Wetter, M., J. Hu, M. Grahovac, B. Eubanks and P. Haves (2018). OpenBuildingControl: Modeling feedback control as a step towards formal design, specification, deployment and verification of building control sequences. Building Performance Modeling Conference and SimBuild.
- Widagdo, S. (2006). "Incandescent light bulb: Product design and innovation." Industrial & engineering chemistry research **45**(25): 8231-8233.
- Winston, A. (2016). "Sensor technology can improve office design and create happier employees, says Haworth." Retrieved July 23, 2019, from <https://www.dezeen.com/2016/06/13/haworth-white-paper-sensor-technology->

[happier-employees-future-office-design-research/](#).

- Wu, X., V. Kumar, J. R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. J. McLachlan, A. Ng, B. Liu and S. Y. Philip (2008). "Top 10 algorithms in data mining." Knowledge and information systems **14**(1): 1-37.
- Xuan, Z., Z. Xuehui, L. Liequan, F. Zubing, Y. Junwei and P. Dongmei (2019). "Forecasting performance comparison of two hybrid machine learning models for cooling load of a large-scale commercial building." Journal of Building Engineering **21**: 64-73.
- Yang, J., H. Rivard and R. Zmeureanu (2005). "On-line building energy prediction using adaptive artificial neural networks." Energy and buildings **37**(12): 1250-1259.
- Yang, L., Z. Nagy, P. Goffin and A. Schlueter (2015). "Reinforcement learning for optimal control of low exergy buildings." Applied Energy **156**: 577-586.
- Yoon, S. and Y. Yu (2017). "Extended virtual in-situ calibration method in building systems using Bayesian inference." Automation in Construction **73**: 20-30.
- Yoon, S. and Y. Yu (2017). "A quantitative comparison of statistical and deterministic methods on virtual in-situ calibration in building systems." Building and Environment **115**: 54-66.
- Yoon, S. and Y. Yu (2018). "Strategies for virtual in-situ sensor calibration in building energy systems." Energy and Buildings **172**: 22-34.
- Žáčková, E., Z. Váňa and J. Cigler (2014). "Towards the real-life implementation of mpc for an office building: Identification issues." Applied Energy **135**: 53-62.
- Zeitoun, B. (2011). Environmental housekeeping handbook: energy efficiency investment opportunities in commercial buildings. Arab Forum for Environment and Development, Beirut, Lebanon.
- Zhang, F., C. Deb, S. E. Lee, J. Yang and K. W. Shah (2016). "Time series forecasting for building energy consumption using weighted Support Vector Regression with differential evolution optimization technique." Energy and Buildings **126**: 94-103.
- Zhang, Z. and K. P. Lam (2017). An implementation framework of model predictive control for HVAC systems: A Case Study of Energyplus Model-Based Predictive Control. 2017 ASHRAE Annual Conference.
- Zhao, J., K. P. Lam and B. E. Ydstie (2013). EnergyPlus model-based predictive control (EPMPC) by using MATLAB/SIMULINK and MLE+. Proceedings of 13th Conference of International Building Performance Simulation Association.
- Zhao, J., K. P. Lam, B. E. Ydstie and O. T. Karaguzel (2015). "EnergyPlus model-based predictive control within design–build–operate energy information modelling infrastructure." Journal of Building Performance Simulation **8**(3): 121-134.
- Zhao, L., Y. Chen and D. W. Schaffner (2001). "Comparison of logistic regression and linear regression in modeling percentage data." Appl. Environ. Microbiol. **67**(5): 2129-2135.
- Zirnhelt, H. E. and R. C. Richman (2015). "The potential energy savings from residential passive solar design in Canada." Energy and Buildings **103**: 224-237.