

**Development, implementation, and industry reception of a
novel multi-source, data-driven building energy management
toolkit**

by

Andre A. Markus

A thesis submission to the

Faculty of Graduate and Postdoctoral Affairs

in partial fulfillment of the requirements for the degree of

Masters of Applied Science

in

Civil Engineering

Department of Civil and Environmental Engineering

Carleton University

Ottawa, Canada

April 2022

© Copyright

Andre A. Markus, 2022

Abstract

Building energy performance is often negatively impacted by inefficient and uninformed operations, leading to wide disparities between predicted and actual energy use in commercial and institutional buildings. Though ample research in data-driven building operations and maintenance analytics has derived various methodologies for extracting energy-saving insights that can supplement operating practice, these approaches have traditionally remained disparate, limiting their application, and are exclusively prevalent in academia. Furthermore, operations personnel who regularly manage controls and maintenance of HVAC equipment can benefit from these novel approaches in augmenting their duties and optimizing building energy efficiency. This research explores the development of a novel multi-source, data-driven building energy management toolkit as a synthesis of established data-driven approaches in the literature comprising inverse energy modelling, anomaly detection and diagnostics, load disaggregation, and occupancy and occupant complaint analytics methods. The toolkit inputs various data types to detect hard and soft faults, optimize sequences of operation settings, and monitor energy flows, occupancy patterns, and occupant satisfaction. The toolkit's unique multi-source analytical approach was used to pinpoint operational deficiencies stemming from inappropriate zone temperature overheating thresholds and perimeter heating devices. Energy-saving insights were generated using data from four separate case study buildings to demonstrate the utility of the toolkit's web-based application platform. Finally, interviews with building operators and facility managers to their interpretations of insights from data-driven approaches were conducted; possible barriers were identified which inhibited industry professionals from effectively deriving and utilizing insights from the visualizations and KPIs.

Preface

The integrated thesis herein consists of two journal papers and one conference paper, either published or undergoing review. However, readers who wish to refer to materials from this document should cite this thesis. The following articles are contained within this thesis:

- **Article 1:** Markus, Andre A.; Hobson, Brodie W.; Gunay, H. Burak; and Bucking, Scott. A framework for a multi-source, data-driven building energy management toolkit. *Energy and Buildings*. 2021; 250: 111255.
- **Article 2:** Markus, Andre A.; Hobson, Brodie W.; Gunay, H. Burak; and Bucking, Scott. FRAMEWORK: A multi-source, web-based application to identify suboptimal energy use management. *Proc. of IBPSA-Canada's eSim 2022 Conf.* [Accepted]
- **Article 3:** Markus, Andre A.; Hobson, Brodie W.; Gunay, H. Burak; and Bucking, Scott. Does a knowledge gap contribute to the performance gap? Interviews with building operators to identify how data-driven insights are interpreted. *Energy and Buildings*. [Under review]

Slight alterations were made to the articles for cohesion of this document. Use of copyrighted material from the published articles is acknowledged as per the corresponding publisher's permission guidelines with respect to the authors' rights.

Andre A. Markus is the primary contributor to the research methodology, data analysis, and preparation of written material and figures presented in the aforementioned articles, under the supervision of Dr. H. Burak Gunay. Brodie W. Hobson contributed through data acquisition (particularly in the case studies of section 3.0), source code refactoring, and, alongside Dr. Scott Bucking, assisted through critical review and feedback of the articles.

Acknowledgements

I express my most sincere appreciation to my supervisors, Professor Burak Gunay and Professor Scott Bucking. Your guidance, support, and insight, alongside your commitment to the research field, have fostered for me a genuinely enjoyable and gratifying research experience, and were instrumental in enabling me to meet the expectations of others and myself. I extend my thanks to my co-author, Brodie Hobson, and to Narges Zaeri, Saptak Dutta, and Hendrik Poernama for their technical expertise.

My gratitude extends to members of my committee, Professor Liam O'Brien and Professor Jean Duquette, for their time in providing feedback on the manuscript of this dissertation. I wish to thank members of the Building Performance Research Centre for, despite unprecedented circumstances, projecting an essence of comradery and community through regular virtual seminars, discussions, and events; I am thankful to be part of such a tight-knit, diverse, and knowledgeable group.

This research was funded by Natural Resources Canada, CopperTree Analytics, Bentall GreenOak, Sensible Building Science, Delta Controls, and National Research Council of Canada. I acknowledge the support of Brandon Malleck of Bentall GreenOak, Gavin Symonds of Carleton's Facilities Management team, and Ricardo Moromisato of CopperTree Analytics.

Not least of all, I extend warm thanks to my friends and family, foremost my parents, for their constant and unconditional love, support, and kindness, and to the congregation of the Indonesian Reformed Evangelical Church in Toronto (IRECT). Your prayers have encouraged me through trials and tribulations, and continue to ensure the keeping of my faith; for these, I am indebted.

Table of Contents

Abstract.....	i
Preface.....	ii
Acknowledgements.....	iii
Table of Contents.....	iv
List of Figures.....	vii
List of Tables.....	ix
1.0 Introduction.....	1
1.1 Motivation and research questions.....	3
1.2 Document structure.....	4
2.0 A novel building energy management toolkit’s framework.....	8
2.1 Introduction.....	8
2.2 Background and previous work.....	9
2.2.1 Energy meter data in anomaly detection, benchmarking, and end-use.....	9
2.2.2 HVAC controls network data in fault detection and diagnostics.....	11
2.2.3 Wi-Fi device count data in occupant-centric controls.....	13
2.2.4 Computerized maintenance management system data in benchmarking.....	14
2.2.5 Data-driven approaches to building energy management.....	15
2.3 Motivation and objectives.....	18
2.4 Methodology.....	18
2.4.1 Operation data input.....	21
2.4.2 Intended use of the functions.....	22
2.4.3 Case study.....	24
2.5 Results.....	25
2.5.1 Energy performance benchmarking.....	26
2.5.2 Metadata inference for AHU and zone-level BAS labels.....	30

2.5.3	Detection of AHU anomalies.....	32
2.5.4	Detection of zone-level anomalies.....	39
2.5.5	End-use disaggregation of meter data.....	42
2.5.6	Analysis of hot/cold complaints.....	45
2.5.7	Occupant count estimation.....	48
2.6	Discussion and limitations	51
2.7	Closing remarks	54
3.0	FRAMEWORK: A web-based application to derive insights.....	57
3.1	Introduction.....	57
3.2	Software platform architecture	58
3.2.1	Backend.....	59
3.2.2	Web interface	62
3.3	Demonstration of the toolkit	64
3.4	Results.....	66
3.4.1	Baseline energy performance.....	66
3.4.2	AHU anomaly detection	69
3.4.3	Zone anomaly detection.....	72
3.4.4	Occupancy using motion detection.....	75
3.5	Discussion.....	76
3.6	Closing remarks	77
4.0	Interviews with building operators & facility managers.....	80
4.1	Introduction.....	80
4.2	Background and previous work	81
4.2.1	Data-driven operation analytics	81
4.2.2	Interviews with building operators and managers	83
4.3	Motivation and objectives.....	86
4.4	Methodology.....	87

4.4.1 Questionnaire development	87
4.4.2 A multi-source approach to extracting energy-saving insights.....	89
4.4.3 Data collection and processing	95
4.4.4 Report generation and distribution.....	97
4.4.5 Participant recruitment and interview administration.....	97
4.5 Interview results.....	99
4.5.1 Interpretation of AHU faults and anomalies	99
4.5.2 Interpretation of zone-level faults and anomalies	100
4.5.3 Interpretation of building-level energy use deficiencies.....	101
4.5.4 Interpretation of occupancy and occupant complaints-based KPIs	102
4.6 Discussion.....	102
4.7 Closing remarks	107
5.0 Conclusions.....	111
5.1 Summary and contribution.....	111
5.1.1 A framework for a multi-source, data-driven building energy	111
5.1.2 FRAMEWORK: A web-based application to derive insights.....	112
5.1.3 Interviews with building operators and facility managers.....	114
5.2 Research contributions.....	115
5.2.1 A framework for a multi-source, data-driven building energy	115
5.2.2 FRAMEWORK: A web-based application to derive insights.....	116
5.2.3 Interviews with building operators and facility managers.....	116
5.3 Recommendations for future work	117
Appendix A.....	120
References.....	121

List of Figures

Figure 2.1: Overview of the BEM toolkit with the distribution of data	19
Figure 2.2: Flowchart depicting the BEM toolkit's intended order of operation	24
Figure 2.3: An illustration of the (a) modified three-parameter change point models	28
Figure 2.4: An illustrative example of the Levenshtein edit distance-based	31
Figure 2.5: Schematic visualization of (a) the coldest zone, average zone	35
Figure 2.6: Example of simplified AHU diagram snapshot with AHU dampers	37
Figure 2.7: Visual representation of relationship between outdoor air damper.....	38
Figure 2.8: Clusters of zones, represented by the black boxes, with similar airflow	41
Figure 2.9: Weekly EUI profiles for (a) electricity, disaggregated	44
Figure 2.10: Example of first set of visualization generated by the complaint	47
Figure 2.11: Example visualization of a regression decision tree model	47
Figure 2.12: Example visualization which is intended to depict hot and cold	48
Figure 2.13: 25 th , 50 th , and 75 th percentile occupant count profile	49
Figure 2.14: 75 th percentile superimposed floor-level occupancy count profile	50
Figure 3.1: Backend workflow for user data storing and processing	60
Figure 3.2: Workflow between users, web interface, and the server.....	63
Figure 3.3: Annotated screenshot of upload page.....	64
Figure 3.4: Heating (left column) and cooling (right column) energy use	68
Figure 3.5: Comparison of split-range controller diagrams for AHUs.....	70
Figure 3.6: A split-range controller diagram for an AHU in organization B	72
Figure 3.7: Heating season zone cluster diagram for organization D.....	74
Figure 4.1: Workflow for generating reports.....	89
Figure 4.2: Example outputted visualization of modified AHSRAE Guideline 14.....	91
Figure 4.3: Example outputted visualization of a SAT diagram	92
Figure 4.4: Example outputted visualization of zone cluster diagram.....	94
Figure 4.5: Outputted exemplar pseudocode for real-time adjusting zone-level VAV	95

Figure 4.6: Distribution of available building data and organization to the functions96

List of Tables

Table 2.1: Schedule effectiveness and after-hours energy use fraction.....	28
Table 2.2: AHU health index and summary of detected soft and hard faults.....	38
Table 2.3: Zone health index and summary of airflow controls errors.....	42
Table 2.4: Total annual EUIs of major end-uses for electricity, heating, and cooling.....	45
Table 2.5: Earliest arrival times, latest departure times, and peak occupancy.....	50
Table 3.1: Availability of building data by organization.....	65
Table 3.2: Number of AHUs and thermal zones analysed by organization.....	66
Table 3.3: Number of identified hard faults by organization.....	70
Table 3.4: Number of identified soft faults by organization.....	71
Table 3.5: Summary of zone indices by organization.....	73
Table 3.6: Earliest arrival and departure, latest departure, and longest break duration.....	75
Table 4.1: List of interview questions.....	88
Table 4.2: Example outputted earliest and latest arrival time, latest departure time.....	95
Table 4.3: Breakdown of reports distributed to organization's participating facilities.....	97
Table 4.4: Breakdown of participants.....	99

1.0 Introduction

Commercial and institutional buildings command an ample amount of resources to support its operations [1]. Over the course of a building's lifetime, operations alone accounts for up to 90% of its total energy use [2], with heating, ventilation, and air conditioning (HVAC) services accounting for up to half of that [3]. The management of these HVAC services constitutes frequent maintenance of air handling units (AHU) and variable air volume (VAV) terminals, and implementation of controls which ultimately impact building energy performance. Advanced sequencing settings are prescribed in ASHRAE Guideline 36-2018 [4] which outline optimal AHU modes of operation alongside suggested settings for advanced controls strategies such as an automatic supply air temperature reset. However, there remains a significant discrepancy between the actual amount of energy used in existing buildings and the amount predicted in its design phase; this is known as the "performance gap". With growing awareness of its environmental impact [5], there is an increasing need to improve building energy efficiency. Furthermore, poor building energy performance is often attributed to ill-informed and inefficient operations, including but not limited to a lack of optimal controls strategies or oversight of hard faults at various HVAC levels; poor controls have been reported to attribute as much as a 30% increase to overall energy use [6], [7]. However, improving building energy efficiency extends further than simply identifying and correcting hard and soft faults. Refined estimations of occupancy patterns can be used to optimize building- and zone-level schedules and reduce excess ventilation, and disaggregating bulk energy use into its constituent end-uses within a building can be used to identify and diagnose abnormal energy use at the component-level.

In response to this “performance gap”, data-driven approaches have secured prevalence in the literature in the past couple decades. These approaches capitalize on mainly process-history data by utilizing machine learning (ML) and optimization algorithms to derive energy-saving insights [6], [8]. Among established methods, inverse energy modeling, fault detection and diagnosis (FDD), end-use disaggregation, and occupancy and occupant complaint analytics were derivations of existing and emerging process-history data types such as energy meter and Wi-Fi device count data. These approaches are software based and can be implemented in a wide variety of existing buildings [9]. However, despite its prevalence in academia and demonstrated benefits of multi-source analytics [10], [11], such methods to improve energy efficiency remain disparate. Furthermore, they are not widely implemented in practice, unlike methods intended for the design phase. Its methodologies are inherently esoteric and its presumed novelty can be unconvincing in a risk-adverse environment such as building operations [12]. Thus, these powerful and proven methods to improve building energy performance remain as untapped potential.

Individual data-driven approaches, while similar in principle, input various types of data and derive vastly different types of energy-saving insights. Should multi-source, data-driven approaches take initiative in practice and make a practical contribution to addressing the “performance gap”, novel applications to identify and diagnose operational deficiencies through the consolidation of these discrete insights should be demonstrated, and their capabilities should be further disseminated. In the first of its kind, a toolkit must encompass various established data-driven approaches in the literature and embrace the public domain.

This multi-source toolkit should facilitate the user input of multiple types and quantity of

data, perform a selection of its included data-driven analyses, and output actionable key performance indicators (KPI) and visualizations.

1.1 Motivation and research questions

In light of a need for a versatile and accessible means to improve building energy performance, a multi-source, data-driven toolkit encompassing several established data-driven approaches in the literature is the subject of this thesis. Its framework, development, functionalities, and practical interpretations of its outputs form the cores in the subsequent chapters. Accordingly, each chapter attempts to resolve the following key inquiries:

- **Chapter 2 (A novel building energy management toolkit’s framework)**
 - What are the applications, input data types, adopted methodologies, and generated KPIs and visualizations of the toolkit’s functions?
 - What are the benefits of combining multiple data-driven approaches to identifying operational deficiencies or opportunities for improving energy efficiency?
 - What novel applications can be derived from consolidating various insights from multiple data-driven approaches in identifying and diagnosing operational deficiencies?

- **Chapter 3 (FRAMeWORK: A web-based application to derive insights)**
 - What features of user data input and results retrieval are enhanced through a web-based application platform of the toolkit?
 - How do different building configurations of HVAC and submetering infrastructure affect the toolkit’s ability to derive energy-saving insights?
 - What are the data limitations of the toolkit?

- **Chapter 4 (Interviews with building operators and facility managers)**
 - What are the interpretations of building operators and facility managers to the generated outputs (i.e., visualizations, KPIs, reports) of data-driven methods?
 - What barriers inhibit building operators and facility managers from effectively deriving energy-saving insights from the outputs and utilizing the insights to inform improved controls or administrative decisions?
 - How can these barriers be addressed or mitigated through workplace training?

1.2 Document structure

Hereafter, this integrated thesis is composed of three core chapters which describe the framework of a multi-source, data-driven building energy management toolkit, the implementation and utility of the toolkit through a web-based application platform, and the distribution and reception of the toolkit's generated outputs to its intended audience (i.e., building operators and facility managers). The following summarizes the contents of each chapter:

- **Chapter 2:** This chapter presents the framework of a novel multi-source, data-driven building energy management toolkit as a synthesis of various established data-driven approaches in the literature; these approaches include inverse energy modeling, FDD, end-use disaggregation, and occupancy and occupant complaint

analytics. A toolkit was compiled as an open-source functions library using Python containing seven discrete functions which each input one or a combination of types of process-history data. The inputs, methodology, and resulting outputted visualizations and KPIs of each function are described, along with its intended use case. The functionality of each function was demonstrated on a case study building in Ottawa, Canada, and the insights gathered from the function's outputs were discussed.

- **Chapter 3:** This chapter presents and demonstrates the toolkit's web-based application platform which hosts and provides user interaction with the toolkit's functions and disseminates the inputs, methodologies, and resulting visualizations and KPIs to a global audience. The backend server-based handling of user data and front-end web-based user interface of the toolkit was developed to facilitate user data input and results retrieval of the toolkit's functions. The web-based application, aptly named "FRAMEWORK", was demonstrated using various types of process-history data from four separate buildings to invoke up to five functions in the toolkit; these functions produced multiple reports containing the corresponding function's generated visuals and KPIs. The insights gathered from these reports were discussed and compared.
- **Chapter 4:** This chapter presents the interpretations of data-driven insights to the toolkit's intended recipients (i.e., building operators and facility managers). Interviews were solicited with building operators and facility managers regarding the data-driven outputs of the toolkit. Data from the participant's individual building portfolio was retrieved and analyzed using up to five of the toolkit's

functions. The resulting reports containing visuals and KPIs was distributed to participants who were subsequently interviewed regarding what insights they gathered from the reports and what changes they would make, either administratively or through controls, after considering the inefficiencies or faults realized in the reports.

The final chapter (Chapter 5.0) summarizes the three preceding chapters and includes their findings, contributions, and recommendations for future work described.

Chapter 2

This chapter has been published as:

Markus, Andre A.; Hobson, Brodie W.; Gunay, H. Burak; and Bucking, Scott. A framework for a multi-source, data-driven building energy management toolkit. *Energy and Buildings*. 2021; 250: 111255

2.0 A novel building energy management toolkit's framework

2.1 Introduction

Buildings often exhibit significant discrepancies between predicted and actual energy performance, largely attributable to suboptimal and uninformed operations, with reports of buildings consuming substantially more energy than predicted during the design phase [13]–[18]. The lack of accessible operations management solutions that derive insight and feedback contributes to this ‘performance gap’. There is great potential for operational improvements as operations account for 80% to 90% of a building’s life-cycle energy consumption [2], with over half of that consumed by HVAC systems [19], [20]. Despite ample resources and tools to support energy optimization in the design phase, methods to supplement energy optimization efforts during existing buildings’ operational phase lack substantiation and standardization. Improvements to the operation of buildings can result in overall annual energy savings of up to 30% [7], [21], [22], and presents considerable opportunities for energy-savings in a wide variety of existing buildings [9]. Furthermore, energy audits – which are traditionally performed to inform energy-saving measures – are often time-consuming, scarcely executed, and can ultimately be cost prohibitive to maintain continual energy savings [23], while a software-based operation-centric solution that is integrated into the existing building controls and operation data infrastructure can be implemented with relatively low effort. Thus, methods to improve buildings performance through better operation are promising for addressing the performance gap in existing buildings. This chapter presents a multi-source data-driven energy management toolkit as a synthesis of established data-driven approaches to operational analytics in the

literature. The capabilities of the toolkit were demonstrated on a case study building in Ottawa, Canada.

2.2 Background and previous work

Data-driven approaches for building operation are prevalent in the literature. Moreover, the growing practice of archiving long-term operational data presents potential for operation analytics [24]. Common examples of these data sources include energy meter data, HVAC control network data from the building automation system (BAS), Wi-Fi-based device count data from IT networks, and computerized maintenance management system (CMMS) data containing occupant complaints and work order logs.

2.2.1 Energy meter data in anomaly detection, benchmarking, and end-use disaggregation

Research using meter data have gravitated towards developing data-driven models to establish an energy use baseline, detect energy use anomalies, and make energy use forecasts. Though typically used by energy service companies (ESCO) to quantify energy savings achieved through a retrofit, baseline energy modelling can also be used to detect and interpret energy use anomalies and act as an energy performance benchmark. Zhang *et al.* [25] compared four popular baseline modelling approaches and their ability to predict HVAC heating energy consumption, noting the benefit and comparable predictive accuracy of simple change-point regression models to that of physics-based models. Gunay *et al.* [26] compared three inverse modelling approaches using heating and cooling loads, finding that nearly half the buildings tested in the study did not have a weekly AHU operating schedule. The complex and multi-faceted behaviour of building energy often solicits the use of such regression models to accurately establish energy baselines and

forecast energy consumption [25], [27]. Meter data have also been used to develop methods for energy end-use disaggregation. Most meter networks in existing buildings lack the resolution to understand major energy end-uses and their flows within a building. Though prevalent in the residential setting [28]–[30] to disaggregate energy use by appliances, its application and insight have seldom extended to large-scale commercial buildings. For example, Akbari [31] developed an end-use disaggregation model using meter data, temperature dependent load regression coefficients, and simulated end-use models to disaggregate hourly building electricity consumption into its major end-uses. Ji *et al.* [32] presented a Fourier series-based model to disaggregate electricity consumption into light and plug-in, HVAC, and other miscellaneous loads such as elevators in office buildings and shopping malls. Doherty and Trenbath [33] used smart plug submetering to assess the viability of a model to disaggregate individual plug-in loads of typical office equipment and its correlation to the building’s plug load submeter data. Operators of commercial buildings with access to limited submetering data are unable to discern the distribution of energy within the building. This insight can be useful for isolating energy use anomalies to specific end-uses from bulk meter data.

Recent advancements in energy anomaly detection and end-use disaggregation can be attributed to ML techniques which can detect operational change points and optimize parameters for inverse models. Miller *et al.* [34] reviewed several publications utilizing unsupervised ML techniques to detect anomalous energy use for whole buildings, noting the frequent use and variations of novelty detection and clustering techniques. A similar review by Himeur *et al.* [35] highlighted the use of metaheuristic optimization algorithms such as the genetic algorithm (GA) and particle swarm optimization as supplement to ML

techniques. Touzani *et al.* [36] employed a penalized change point detection method to detect irregularities in energy use independent of existing energy-saving measures, thus improving reliability of baseline energy models; this method can be extended to energy use anomaly detection applications.

2.2.2 HVAC controls network data in fault detection and diagnostics and model-based predictive controls

Research using HVAC control network data has been focussed on developing methods for automated fault detection and diagnosis (AFDD) and predictive controls, and efforts to improve the usability of HVAC control network data in AFDD and predictive controls at scale has led to the development of new metadata schemas and automated metadata inference algorithms. AFDD approaches have been intended to identify common hard faults in a building's HVAC system (e.g., stuck AHU dampers and valves, etc.) and common soft faults (e.g., a lack of an AHU operating schedule [37], [38], etc.). Kim and Katipamula [39] categorized several AFDD studies by their methods, noting that process-history based AFDD approaches, whereby archived performance data is used to estimate the parameters of inverse models, were the most popular. These studies were further subdivided into two groups, grey-box and black-box model-based methods, with the latter often deriving parameters with no physical implication [39]. A recent review article by Shi and O'Brien [38] noted that black-box methods were especially appealing to researchers due to their flexibility and ease of development, though the challenge becomes how to present AFDD results to building operators in such a way to intuitively foster energy-saving decisions [8]. Model-based predictive control (MPC) is another subfield of operational data analytics using HVAC controls data. Gunay *et al.* [40] reviewed several

indoor climate control strategies in the literature including MPC, noting its ability to generate optimal controls outputs given constraints over a prediction horizon and its applications of different forms of MPC to zone-, system-, and plant-level HVAC equipment. Afram and Janabi-Sharifi [41] concluded that most MPC approaches are based on steady-periodic linear models of the system, and can incorporate various factors related to energy cost and consumption as well as external factors such as weather and occupancy, allowing for a robust implementation in HVAC systems. Simulated and experimental system-level applications have been demonstrated to reduce energy consumption by ~20% [42]–[44]. There are, however, notable disadvantages to data-driven AFDD and MPC models derived from such large datasets. They require large quantities of high quality and high temporal resolution data under a wide range of operating conditions to function optimally, and their accuracy and reliability is heavily influenced by the quality of data [45]. Thus, the implementation of such data-driven approaches constitutes cleaning the data, such as removing stagnant values and outliers due to faulty sensors and temporary network communication issues, for example.

Though metadata models such as Project Haystack and Brick exist, the seldom implementation of a standardized data labelling terminology or format [46] also presents a major obstacle for data-driven approaches using HVAC controls data. As BASs in most existing buildings do not follow a standard metadata model, BAS metadata need to be inferred by human experts from labels, graphics, and as-built drawings prior to the deployment of data-driven energy management algorithms for AFDD, predictive controls, etc. As this is a manual and labour-intensive process, an active research area centers on

developing metadata inferencing algorithms which can automatically classify labels by type and associate them by their functional relationship [47]–[50].

In recent years, ML techniques have contributed to advancements in AFDD applications, particularly for their efficiency in analyzing large amounts of data [51]. Yan *et al.* [52] employed a density-based clustering algorithm to detect sensor faults in AHUs. Similarly, Narayanaswamy *et al.* [53] developed a cluster analysis-based anomaly detection method, which identified 78 anomalies in 237 thermal zones. Metaheuristic optimization techniques, such as the GA, have been particularly effective in estimating model parameters, subject to their practical and physical bounds. For example, Wang *et al.* [54] developed a model-based AFDD method for VAV terminals which used the GA to dynamically adjust parameters of a model, resulting in higher overall accuracy.

2.2.3 Wi-Fi device count data in occupant-centric controls

In recent years, Wi-Fi device count data have been recognized as a promising proxy to estimate occupancy levels in commercial buildings [55], [56]. Research using Wi-Fi device count data has been used in guiding occupant-centric control (OCC) algorithms. Wi-Fi data present a non-intrusive approach to infer occupancy metrics with relatively flexible applicability [57]. Wang *et al.* [57] developed and tested a Wi-Fi based occupancy model employing an ensemble modeling technique, specifically random forest, noting its high accuracy and ability to integrate with existing Wi-Fi infrastructure. Zhao *et al.* [58] developed a real-time occupancy detection model at the zone- and room-level using a Bayesian belief network (BBN) which incorporated a fusion of physical occupancy sensors with Wi-Fi and GPS data. Similarly, Longo *et al.* [59] developed an occupancy estimation model using Wi-Fi data and Bluetooth connectivity. Wi-Fi data has been used in concert

with Bluetooth or GPS data to refine its estimation accuracy, though arguably, archived Wi-Fi data is more prevalent in commercial buildings and would not require additional sensors. Implementation of such occupancy-based models have generally gravitated towards HVAC controls due to their demonstrated effectiveness to reduce energy consumption [60]–[66]. Park *et al.* [67] reviewed several occupancy-driven approaches to building controls, noting key occupancy indicators such as count, arrival time, and departure time, albeit few used Wi-Fi data. Zou *et al.* [68] implemented a predictive light dimming algorithm which derived occupants' presence and occupant count using Wi-Fi data to dynamically adjust lighting conditions based on occupant count and location, resulting in a 93% lighting energy saving from traditional lighting schedules in a commercial building. Balaji *et al.* [69] presented a predictive model for AHU outdoor air damper operations using Wi-Fi for high resolution occupancy data, resulting in HVAC electricity savings of ~18%, and Alishahi *et al.* [70] proposed a framework for extracting occupancy indicators using Wi-Fi data and various ML techniques. Though it should be noted that Wi-Fi data, as with any signal propagating medium, can suffer from dead-zones which will inhibit estimation accuracy [71]. Privacy is another concern which may cause reluctance in implementation of these models.

2.2.4 Computerized maintenance management system data in benchmarking operational performance

Albeit scarce in the practice, CMMS data in the form of complaint logs and work orders have been analyzed and used to identify anomalous zones and malfunctioning equipment. Occupant temperature complaints are commonly symptomatic of faulty HVAC operation though often the corrective course of action for operators is to change temperature setpoints

[72]. While this change may satisfy the individual complainant, it may dissatisfy the larger group. In this case, the problem may be more suitably attributed to an anomaly with the particular zone. Dutta *et al.* [73] employed association rule mining to extract occupant hot/cold complaints, and explored the spatiotemporal relationship of the complaints by using decision trees to quantify complaint frequencies and detect anomalous zone conditions; the same method was used to infer building performance metrics from tenant surveys [74]. Gunay *et al.* [75] proposed a complaints-based algorithm to extract insights into maintenance performance of building systems and components, and Assaf and Srour [76] developed a forecasting model for predictive maintenance employing artificial neural networks (ANN) to predict number of thermal complaints. The greatest challenge in using CMMS data is the interpretation of such a large and unstandardized dataset. Thus, use of this type of data solicits ample preprocessing.

2.2.5 Data-driven approaches to building energy management

In the past two decades, research fields using operational data have gained popularity with the increased practice of archiving operational data. However, these data sources have been largely treated as disparate. Thus, previous attempts at developing a data-driven building energy management (BEM) solution often focused on a single domain of data analytics. For instance, the Commercial Building Energy Saver (CBES) presented by Hong *et al.* [77] is intended to assess building performance pre- and post-retrofit and help owners make optimal retrofit decisions. Besides geometrical and construction parameters, the toolkit relies exclusively on electricity and gas usage for baseline energy benchmarking. Costa *et al.* [78] formulated a toolkit built upon previous efforts in data-driven strategies using HVAC controls data to derive operational modes, though ultimately focused on AFDD.

Zhang *et al.* [79] developed an open-source toolkit that synthesized existing occupant detection devices and analyzed occupancy patterns which can be used to tailor operating schedules; this toolkit was limited to OCC. BETTER by Li *et al.* [80] and FirstView by the New Buildings Institute [81] both offer a data-driven toolkit employing inverse modelling to benchmark building energy consumption, disaggregate end-uses, and identify energy-saving measures, though both limit derivations of energy-saving insight from energy use data. The disadvantage of restricting analysis to one operation data source is that it effectively constricts the generated insight from the toolkit to a single approach to energy savings. Ultimately, there are multiple avenues for energy-saving that would go unnoticed and ample energy-saving potential untapped.

Ample benefits exist in integrating various operational data sources for operational analytics [10]. The intricate and interdependent nature of building operations is such that inefficiencies or anomalies in one domain may translate to others [38]. Multi-source operational analytics can be used to derive additional or refine existing energy-saving insight by establishing functional relationships between multiple domains. For example, combining energy meter data with HVAC controls network data can be used to disaggregate end-uses and identify faulty equipment. Gunay *et al.* [11] employed linear regression models using HVAC controls data to characterize the operational state of major energy systems, which enabled high accuracy disaggregation of end-uses at low temporal resolution meter data. Disaggregation produced an accurate representation of the specific components' power draw and the insights generated can be used to focus fault detection efforts into specific equipment rather than an entire system. Integrating occupancy data with energy meters can disaggregate energy loads to individual occupants' plug-in loads.

Rafsanjani and Ahn [82] presented a cluster-based approach to associate occupancy events with power fluctuations using Wi-Fi and electricity load meters. Their approach was tested in two commercial buildings and was used to infer occupancy using power fluctuations from a predetermined baseline power consumption and entry/exit events. Approaches featuring this blend of data can be used to garner a more granular awareness of occupant behaviour and disaggregate electricity loads to occupant's plug-in and lighting loads. HVAC and occupancy data can be used to identify inefficiencies in outdoor airflow control and scheduling. Hobson *et al.* [83] developed and tested a cluster analysis-based approach focusing on occupancy via Wi-Fi data; classification trees intended to project day-ahead occupancy based on day type were employed using motifs from lighting and plug-in loads. The case study on the academic building exposed chronic overventilation, well above what was required for the projected occupancy. This insight was used to tailor occupant-centric schedules and ventilation rates. Combining CMMS and HVAC data show promise in identifying system conditions that trigger hot/cold complaints. Gunay *et al.* [84] modelled the frequency of complaint-driven temperature setpoints adjustments with concurrent indoor and outdoor temperature data. Their investigation revealed two distinct optimal temperature setpoints that minimize thermal complaints during the heating and cooling seasons separately. Evidently, optimal building operations stand to benefit from the integration of multiple data sources as it relies not only in the awareness of suboptimalities but also the targeting of suboptimality causes. A BEM toolkit incorporating multiple operational data sources would address multiple domains of energy savings and increase the flexibility of energy-saving decisions.

2.3 Motivation and objectives

While there have been major developments in the past two decades in data-driven energy monitoring and management methods, the majority are derived from a single domain of operational data, despite a variety of existing and emerging system performance data [85], [86]. The tendency to focus on a single stream of data means existing research and case studies in BEM toolkits tend to gravitate towards a single domain of performance data analytics such as energy benchmarking, AFDD, or OCC to name a few. There is a lack of studies and discussion within the building energy community regarding the applicability of a multi-source BEM toolkit. This chapter presents a preliminary multi-source, data-driven BEM toolkit that synthesizes established data-driven approaches from the literature. The toolkit is compiled as a Python function library and is publicly available through a GitHub repository along with documentation, example visualizations, and sample data; the link to the repository is found in Appendix A. The framework of the toolkit is intended to prompt discussion and stimulate further advancements in multi-source BEM toolkits. In addition to presenting the toolkit, the present capabilities of the toolkit are demonstrated using a six-storey academic building as a case study.

2.4 Methodology

This section presents an overview of the function library with relationship to the publications forming its structure, the input data types with the relevant parameters, and its intended uses. The algorithms, literature sources, specific inputs and generated KPIs of each function are presented in greater detail, as well as results from the case study, in Section 2.5. Figure 2.1 illustrates the input data types, distribution of the data, and the KPIs that each function output.

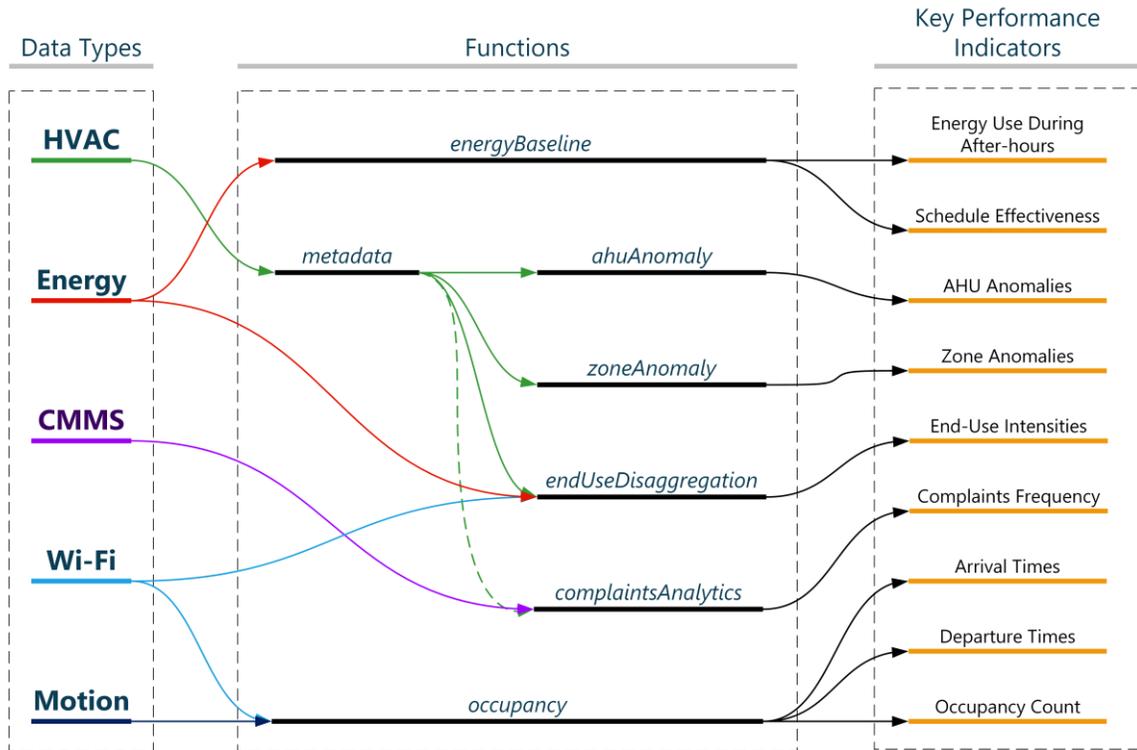


Figure 2.1: Overview of the BEM toolkit with the distribution of data and outputted KPIs of each function.

The corresponding data types are inputted into the functions which derive their own unique KPIs. Energy meter data are typically archived and available in commercial and institutional buildings. In the past couple decades, there has been an increasing practice in the collection and storage of HVAC controls network data [8], [39]; these data, as well as motion detector sensor data, would be extracted from a building’s BAS. Wi-Fi device count data are commonly available in institutional buildings with a network of Wi-Fi access points (i.e., university campuses), however, these data may be more difficult to acquire in commercial buildings since each tenant commonly has a separate Wi-Fi network; dedicated Wi-Fi enabled device counting sensors would need to be deployed throughout the building. Work order logs are often available within a CMMS database. Outdoor air temperature is used by the *energyBaseline* and *complaintsAnalytics* functions, though is treated as a discrete hourly weather data source (i.e., separate from HVAC controls network data) and

not explicitly shown. Note also that HVAC controls data are optional for the *complaintAnalytics* function.

Development of the BEM toolkit began with understanding the needs and challenges of the industry and the application venues for multi-source data analytics. A stakeholder's workshop titled "Designing Future Building Energy Management Tools" was organized with over 50 expert participants in Ottawa, Canada. The participants included representatives from eleven companies that specialize in data analytics for BEM, two large building portfolio management companies, and government and university researchers. The workshop discussed the potential for additional insights from incorporating multiple data sources. For example, complaint analytics can be derived from investigating the spatiotemporal relationship between CMMS and HVAC controls data. Energy meter, occupancy, and HVAC controls data can be used to further disaggregate energy use into the building's major end-uses. Key publications addressing multi-source data analytics applications were identified and gaps in such publications were addressed as additional contributions to the literature. Subsequently, interviews with building operators were conducted by Afroz *et al.* [87] to reveal the current functionalities of existing software tools. The findings of the workshop and interviews formed the structure of the function library which was subsequently presented to the industry stakeholders. The functions were developed to bridge the gap between information and action by providing operators with a means to quickly extract energy performance metrics and generate a brief and intuitive interface to inform energy-saving decisions.

The function library was prototyped in MATLAB and subsequently translated into Python, an open-source programming language. A qualitative analysis was conducted (described

later in subsection 2.4.3) to validate the accuracy of the translation. A link to the repository can be found in Appendix A.

2.4.1 Operation data input

Hourly HVAC controls network data at the AHU- and zone-level served as input for five of the seven functions in Figure 2.1. These data are trend logs of AHU- and zone-level sensor and actuator measurements which were extracted from the BAS. The relevant data types for AHUs are supply, return, and outdoor air temperature (T_{sa} , T_{ra} , T_{oa}), supply air pressure (P_{sa}), outdoor air damper position (S_{oa}), heating coil valve position (S_{hc}), cooling coil valve position (S_{cc}), supply fan state (S_{fan}), supply air temperature setpoint (T_{sasp}), and supply air pressure setpoint (P_{sasp}). The relevant parameters for zones are the indoor air temperature (T_{in}), VAV terminal device airflow rate (Q_{flo}) and airflow setpoint (Q_{flosp}), and perimeter heater state (S_{rad}).

Hourly building-level energy meter data for electricity, cooling, and heating energy use served as input for three of the functions in Figure 2.1. The electricity data typically represents end-uses for lighting, plug loads, fans, pumps, chiller consumption, and other equipment (e.g., elevators). The heating and cooling data represents the thermal energy transferred to the building at hourly intervals.

Two CMMS data fields, the operator description of the work order and the report time, are used as input for one of the functions in Figure 2.1.

Floor-level hourly Wi-Fi device count data acquired from an institutional IT network are used as input for two of the functions in Figure 2.1.

2.4.2 Intended use of the functions

The BEM toolkit is intended to provide building operators with a versatile method to address operational deficiencies by interpreting metadata, fixing hard faults, fixing soft faults and upgrading sequences, and monitoring KPIs. With the exception of the metadata inference function (*metadata*) which is intended to automate the selection of the appropriate sensor and actuator labels used to extract HVAC controls network data, each function employs a unique methodology from the existing literature and generates KPIs and visualizations which can be used to inform energy-saving decisions or monitor energy use.

The baseline energy function (*energyBaseline*) serves to illustrate the building's heating, cooling, and electricity energy use as a function of outdoor air temperature during work-hours and after-hours. The generated KPIs are intended to help building operators assess the effectiveness of current operating schedules in reducing energy consumption outside of the AHUs' operating hours and assess the effects that fixes to hard and soft faults and changes to sequences of operation have on energy consumption.

Soft and hard AHU and VAV terminal faults may result in unintended excessive energy consumption. The AHU anomaly detection function (*ahuAnomaly*) is intended to detect anomalous operating conditions in AHUs which may be symptomatic of faults such as stuck heating coil valves, stuck cooling coil valves, leaky or stuck dampers, excessive use of perimeter heating in the economizer mode, absence of a schedule, and/or absence of an economizer mode. The generated insight is intended to help building operators focus fault correction efforts to specific AHUs and components. Similarly, the zone anomaly detection function (*zoneAnomaly*) is intended to detect anomalous zones based on temperature and

airflow control errors which can help building operators isolate fault correction efforts to specific zones and VAV terminal units.

An understanding of how energy is distributed within a building may aid in fostering energy-saving decisions and may be used to isolate energy use anomalies in major end-uses. The end-use disaggregation function (*endUseDisaggregation*) further separates bulk electricity, heating, and cooling consumption into six end-uses: electricity consumption into lighting and plug-in equipment loads, distribution, and chillers, heating consumption into perimeter heating and AHU heating coils, and cooling consumption into AHU cooling coils. This function is intended to calculate energy use intensities (EUIs) at a higher end-use resolution and help operators understand the flow of energy within the building's systems should end-use submetering be unavailable in their buildings.

Occupant-count patterns may be used to derive ventilation schedules that mitigate excessive energy use from overventilation, arising from an assumed constant and often overestimated occupant-count in the HVAC design phase. The occupancy function (*occupancy*) estimates occupancy-centric metrics per floor and is intended to inform building operators in refining ventilation schedules to appropriate amounts for the estimated occupancy.

Occupant feedback may be used to estimate optimal indoor temperature setpoints for certain outdoor conditions that reduce the frequency of manual setpoint changes. Setpoint changes can impact a building's energy performance [88] which can be exacerbated with increased frequency of user-solicited changes. The complaint analytics function (*complaintAnalytics*) extracts hot/cold related complaints and models conditions which

trigger those complaints, and is intended to aid building operators in understanding what conditions give rise to thermal complaints.

Figure 2.2 summarizes the toolkit’s intended order to address operational deficiencies and the functions’ intended uses. It should be noted that the functions’ KPIs and visualizations, though developed with its intended use in mind, should not be limited to its intended use. For example, though the end-use disaggregation function is intended to monitor the flow of energy within a building, it can also be used to inform fault evaluation. However, the order of operation is critical. Detecting faults without first interpreting the metadata is difficult and understanding the flow of energy in a building without first correcting hard and soft faults is a futile effort.

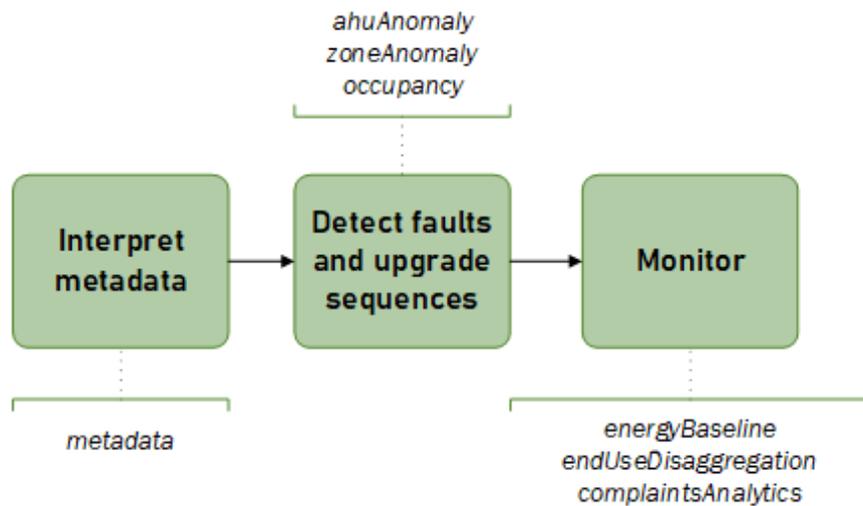


Figure 2.2: Flowchart depicting the BEM toolkit's intended order of operation to address operational deficiencies with the corresponding functions.

2.4.3 Case study

The capabilities of the toolkit were demonstrated on an academic building in Ottawa, Canada. Each function was employed using one or multiple streams of archived operational data accumulated over a year. The building is six-storeys high with one floor below-grade, has a gross area of 6650 m², accommodates a suite of classrooms, offices, meetings rooms,

and labs, and is equipped with two AHUs which serve 42 thermal zones. Each zone is equipped with at least one type of perimeter heating device – VAV reheat coils and/or ceiling-mounted hydronic heaters. Cooling is provided using two chillers located on site. A district plant provides steam which passes through the building’s heat exchanger and heats hot water used in the AHUs’ heating coils, VAV terminal reheat coils, and perimeter hydronic heating devices. The building has a switchover period between the heating and cooling season whereby mechanical cooling is unavailable during the heating season and space heating is unavailable during the cooling season. Two to five Wi-Fi access points are located on each of the seven floors. Hourly or sub-hourly Wi-Fi-enabled device count, HVAC control network trend, and energy meter network data were extracted, along with CMMS data, from four disparate operational databases. The case study dataset is provided along with the functions in the public GitHub repository; the link is provided in Appendix A. These data were also used to validate the accuracy of the Python translation of the MATLAB source code; the outputted visualizations and KPIs from both sets of code were qualitatively analyzed and was determined to be identical or of strong likeness to each other.

2.5 Results

This section presents the inputs, generated KPIs and visualizations, algorithms used, assumptions and limitations, and accompanying literature which forms the basis of each function. The implementation of each function on data from the case study building is discussed along with the resultant visualizations and KPIs, as well as any potential insights.

2.5.1 Energy performance benchmarking (energyBaseline function)

Inputting energy meter data and outdoor weather data, the baseline energy function adopts the methodology described by Gunay *et al.* [26] and Afroz *et al.* [89]. The KPIs generated are the schedule effectiveness (SE) and after-hours energy use ratio (AEUR), calculated separately for heating, cooling, and electricity use. The SE is calculated per Eqn. 1 whereby $k_{afterhours}$ represents the change in energy use with respect to outdoor air temperature outside scheduled AHU operating hours, and $k_{workhours}$ represents the change in energy use with respect to outdoor air temperature during scheduled AHU operating hours. Lower values (approaching 0%) correspond to similar rates of change of energy use during workhours and after-hours with respect to outdoor air temperature which may be symptomatic of an ineffective or outright absence of a schedule where there is minimal reduction to energy use during after-hours. Greater positive values indicate a greater rate of change of energy use during workhours than after-hours while negative values indicate a greater rate of change of energy use during after-hours than workhours. The AEUR is calculated per Eqn. 2, which computes the ratio of energy use during after-hours to the total energy use over the reference period. Lower values indicate minimal after-hours energy use while values approaching and exceeding 50% indicate similar or greater after-hours energy use. The energy baseline function generates two modified ASHRAE Guideline 14 [90] three-parameter univariate change point models - one representing workhours energy use and another representing after-hours energy use - for heating, cooling, and electricity use separately. Each model is defined using three parameters – the rate of change of energy use, the change point temperature, and the energy use associated with the change point

temperature, with outdoor air temperature as the regressor. The after-hours and workhours separation of the modified three-parameter change point model is governed by a steady-periodic weekly AHU operating schedule. The model parameters and the start and stop times associated with the weekly AHU operating schedule were iteratively estimated using the GA which minimized the root-mean-square-error (RMSE) between the measured and modeled energy use, and were subject to upper and lower bounds. The GA was performed with 15 generations and a population size of 5000. Note that workhours and after-hours consumption is not representative of energy use during and outside occupant work schedules, but rather differentiates energy use during scheduled AHU operating hours and outside scheduled AHU operating hours, respectively.

$$\text{Schedule Effectiveness (SE)} = 1 - \frac{k_{\text{afterhours}}}{k_{\text{workhours}}} \quad \text{Eqn. 1}$$

$$\text{Afterhours Energy Use Ratio (AEUR)} = \frac{\sum E_{\text{afterhours}}}{\sum E_{\text{afterhours}} + \sum E_{\text{workhours}}} \quad \text{Eqn. 2}$$

Collected meter data and outdoor temperature weather data from January 1st, 2019 to December 31st, 2019 from the building's three dedicated heating, cooling, and electricity meters. These data were used to demonstrate the capability of the baseline energy function. Figure 2.3(a) depicts the building's baseline heating energy consumption as a function of outdoor air temperature; two other similar plots were produced which depicted cooling and electricity use. The primary slope which represents the energy use rate during scheduled AHU operating hours is the steeper of the two slopes, while the shallower slope represents the energy use rate outside scheduled AHU operating hours. Figure 2.3(b) represents the building's predicted heating load; another two plots were generated for cooling and

electricity load in the same format. The predicted loads include the average steady-periodic variation during a weekday, resulting in the minor hourly variations, and are presented at select outdoor air temperature. They can be used to assess how energy use fluctuate during the day and how outdoor air temperature affects the magnitude of the fluctuations. Since the predicted energy loads are generated using the same datasets, alterations to the baseline consumption are reflected in the resulting predicted load. Table 2.1 presents the generated KPIs.

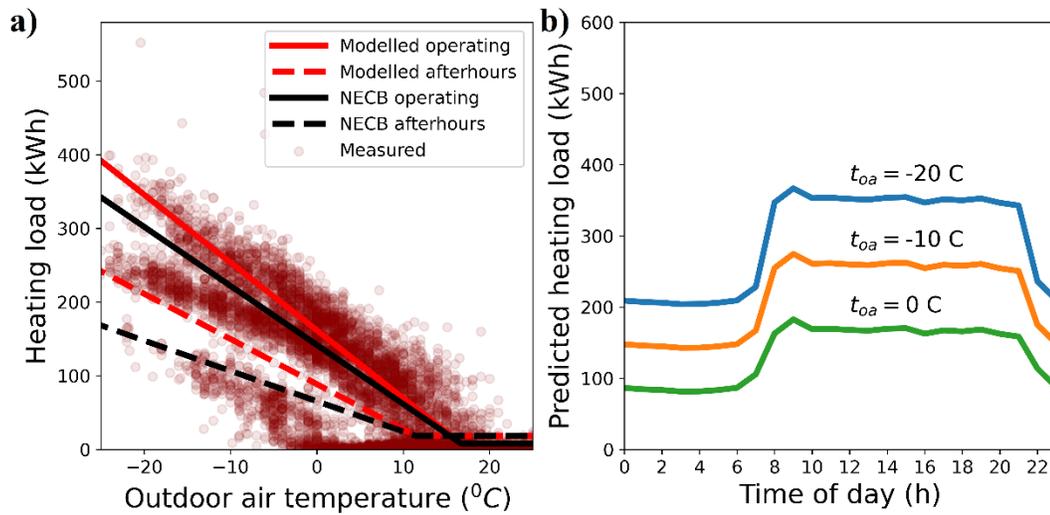


Figure 2.3: An illustration of the (a) modified three-parameter change point models for workhours and after-hours heating energy use compared to workhours and after-hours heating energy use defined in the National Energy Code of Canada (NECB) [91], and (b) predicted heating energy consumption at representative outdoor air temperatures (t_{oa}) with added weekly average residuals.

Table 2.1: Schedule effectiveness and after-hours energy use fraction (energyBaseline function KPIs).

KPIs	Schedule effectiveness (SE)	After-hours energy use ratio (AEUR)
Heating	33%	40%
Cooling	90%	8%
Electricity	2%	50%

Although the high SE and low AEUR for cooling suggest an effective reduction in cooling energy use during after-hours in the case study building, it is unusual that any cooling energy is consumed outside the AHUs' scheduled operating hours. In this case, an

overheating IT room during the cooling season was identified to have woken up the AHUs overnight (i.e., night-cycling). The issue was discussed with the facility managers and the unoccupied mode temperature setpoint was increased from 27°C to 28°C; this behaviour has since stopped.

Electricity use exhibited the lowest SE and highest AEUR, suggesting that electricity energy use was not effectively reduced outside of the AHUs' scheduled operating hours. This is not ideal. Although a fraction of the lights remains on in corridors and multi-purpose rooms during after-hours for safety purposes, the majority operate based on motion detection and turn off automatically if no motion from occupants is detected. Other electricity end-uses such as plug loads and elevators are influenced by occupancy. The most likely cause for the high after-hours electricity use may be from chiller activity and subsequent AHU fan power usage outside the AHU's scheduled operating hours as a result of the previously mentioned overheating IT room (i.e., inappropriate night-cycling). Plug loads outside AHU operating hours would also contribute to the high after-hours electricity use (i.e., computers left on overnight, office appliances such as mini-fridges, etc.).

Heating energy use is less effectively reduced during after-hours than cooling energy use. It should be considered that once the AHUs shut off, ventilation stops and there is no need to heat or cool outdoor air. Thus, the SE for heating is indicative of the role of ventilation on the overall heating demand. An abnormally high heating SE can indicate excessive heating demand during operating hours resulting from overventilation, at least beyond ASHRAE Standard 62.1 [92]. In contrast, an abnormally low heating SE indicates similar heating demands when ventilation is present during operating hours to when there is no ventilation outside operating hours. Such a case may be symptomatic of underventilation

or poor envelope thermal insulative performance. In the case study building, the overheating IT room coupled with a disparate issue with the AHUs economizing in the heating mode, discussed later in subsection 2.5.3, caused the building to intake sub-zero outdoor air during the heating season which exacerbated perimeter heater loads outside scheduled hours.

2.5.2 Metadata inference for AHU and zone-level BAS labels (metadata function)

The metadata inference function which adopts the methodology described by Chen *et al.* [47] automates the metadata mapping process by classifying BAS labels into groups of the relevant sensor and actuator types, and then associating the labels to the corresponding AHU or zone [48], [50], [93]. The function classifies labels based on a pre-defined dictionary of common tag abbreviations which was developed by extracting multiple naming conventions from over 20 buildings' BAS labels commissioned by varying control vendors. For example, labels representing supply air temperature trend log data commonly contain the abbreviations 'tsa' or 'sat'. Thus, labels which contain one or more of these tag abbreviations are included in the label group for supply air temperature data. However, should the labels contain unsuitable tags such as 'sp', 'stp' or 'set', which may suggest the labels represent temperature setpoint data rather than temperature data, they are subsequently excluded. The edit distance for each label was computed and the least dissimilar label of each unique point type was associated to AHUs or zones. Should the labels be insufficiently descriptive and too dissimilar to be associated, the function resorts to applying the same analysis to the labels' controller address which identifies the type of device and control points. A modified Levenshtein-distance analysis as depicted in Figure

2.4 is employed for zone-level labels where numerical changes are weighed more heavily than alphabetical changes.

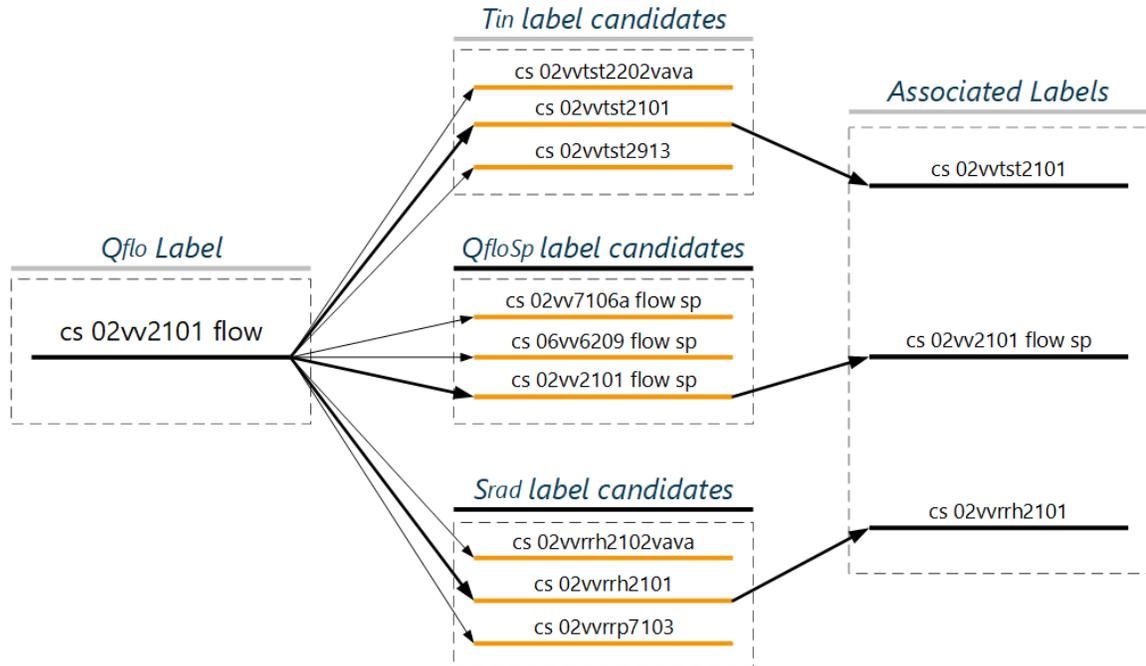


Figure 2.4: An illustrative example of the Levenshtein edit distance-based label association process for zone-level metadata labels. Indoor air temperature (T_{in}), airflow rate setpoint (Q_{floSp}), and fraction of active perimeter heaters (S_{rad}) label candidates are compared to the airflow rate (Q_{flo}) label, and the label candidates resulting in the least edit-distance is selected and associated with the Q_{flo} label.

The metadata inference function assumes that every AHU has one cooling coil with a unique label and that every zone has one airflow sensor with a unique label; the number of AHUs and zones are determined through this assumption. It is also assumed that HVAC controls data labels are associated with the AHU or zone so long as the label name or controller address are alphabetically or numerically similar (i.e., the cooling coil label for a particular AHU is similar to all other point type labels associated with that AHU). With the case study building, these assumptions resulted in inaccuracies in the metadata mapping process where one point label would be applied to multiple AHUs; these inaccuracies are expected. As such, this function is intended to be part of a semi-automated process administered by an analyst whereby review and manual revisions are allowed afterwards.

Though the metadata inference function is intended to automate BAS metadata label classification in buildings without a metadata ontology, it is expected that this function may work just as well with metadata models like Brick and Project Haystack. The tag library which is used to classify metadata labels into groups of the relevant sensor and actuator types could be updated to specifically reflect the metadata ontology employed by such initiatives. This might even improve reliability of the function due to the greater level of labeling consistency that these initiatives provide.

It should be noted that this function does not output any KPIs and serves purely to identify the labels which correspond to the required data types for HVAC controls network trend data. Though its purpose is inconsistent with the other functions in the toolkit, its inclusion was necessary as it enables the user to quickly identify the relevant HVAC controls network trend data required for the AHU anomaly detection, zone anomaly detection, end-use disaggregation, and complaint analytics functions.

2.5.3 Detection of AHU anomalies (ahuAnomaly function)

The AHU-level anomaly detection function inputs AHU- and zone-level HVAC controls network trend data to detect common soft and hard faults in AHUs and is based on the methodology described in by Gunay and Shi [94] and Darwazeh *et al.* [95]. The function outputs an AHU health index KPI for each AHU which yields a value of 100% when the AHU is free from the six fault types referenced in Table 2.2 (three hard and three soft fault categories) and 0% when all six faults are present. The function also develops two visualizations for each AHU in a building. The first presents the AHU actuator positions and the fraction of active perimeter heater devices on a schematic demonstrating the four modes of operation described in ASHRAE Guideline 36 [4] (i.e., heating, economizer,

economizer with cooling, and cooling) over outdoor air temperature, see Figure 2.5. The coldest, warmest, and average zone temperatures as well as the supply air temperature are superimposed. The second visualization is a set of simplified AHU diagrams which capture AHU actuators and supply, return, outdoor, and mixed air temperatures at four to six distinctive periods of operations. The fraction of the captured actuator and temperature values to the total operating duration of the AHU and average perimeter heater valve state is also displayed as in Figure 2.6.

The AHU-level anomaly detection function develops an inverse greybox model for each AHU and compares the estimated model parameters to detect common hard faults. The model predicts the supply air temperature of the AHU as a function of the outdoor and return air temperatures, outdoor air damper position, and the heating and cooling coil valve positions. The inverse model for each AHU is defined using three parameters – the outdoor air fraction bias, the increase in air temperature across the heating coil when the heating coil valve is open, and the increase in air temperature across the cooling coil when the cooling coil valve is open. The parameters were iteratively estimated using the GA which minimized the RMSE between the measured and modeled supply air temperature. The GA was performed with 20 generations and a population size of 5000. Detailed information about the model form is described by Darwazeh *et al.* [95]. The hard faults detected by the function were outdoor air damper stuck (open or closed), and heating and cooling coil stuck (open or closed).

In addition to hard faults, three soft faults were detected by employing the cluster analysis-based anomaly detection approach described by Gunay and Shi [94]. The k-means, Gaussian mixture, and agglomerative clustering algorithms were employed to group AHU

operational periods of statistically similar damper, valve, and temperature measurements into four to six clusters for each AHU. Of the three algorithms, the best algorithm and its resultant number of clusters is selected using the Calinski-Harabasz index. The three soft faults detected by the function are excessive perimeter heating in the economizer mode, inappropriate economizer mode settings, and absence of a weekly schedule. The first soft fault is usually associated with an inappropriate supply air temperature setpoint reset approach. The second soft fault can be associated with the incorrect configuration of the split range AHU supply air temperature controller. The third fault can be due to the building not having a schedule or frequent cycling caused by a few rooms overheating beyond unoccupied indoor temperature setpoints; the former can be fixed by implementing an AHU operating schedule and the latter can be fixed by adjusting airflow and/or temperature setpoints of the VAV terminals of the overheating zones during unoccupied periods.

HVAC controls network data from January 1st, 2019 to December 31st, 2019 were used to demonstrate the function's ability to detect faults and energy use anomalies in the case study building's AHUs. The schematic in Figure 2.5 is the first generated visualization for one of the two AHUs in the building. The supply air temperature remains relatively constant at around 16°C and dips slightly from an outdoor air temperature of -3°C to 10°C. This is indicative of a supply air temperature setpoint reset where the supply air temperature is modulated with respect to outdoor air temperature to minimize concurrent cooling and heating. Ideally, supply air temperature is modulated to reflect the cooling and heating demands of the warmest and coldest zone respectively [96], [97]. The GA determined the change point temperatures between the heating and economizer mode,

economizer and economizer with cooling mode, and economizer with cooling and cooling mode to be about -5°C , 12°C , and 21°C , respectively.

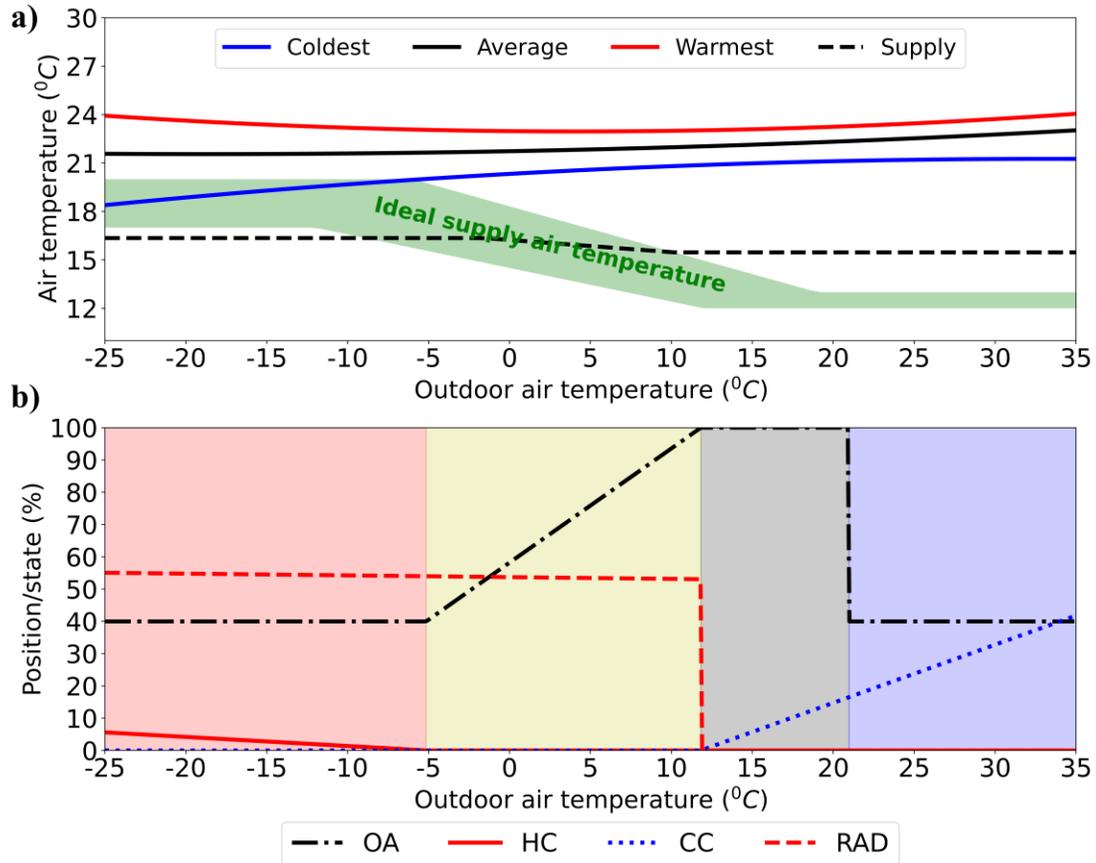


Figure 2.5: Schematic visualization of (a) the coldest zone, average zone, warmest zone, and supply air temperature with respect to outdoor air temperature, and (b) outdoor air damper (OA), heating coil valve (HC), cooling coil valve (CC) positions, and fraction of active perimeter heater devices (RAD) with four modes of operation over outdoor air temperature. The AHUs' four modes of operation are represented as heating (red zone), economizer (yellow zone), economizer with cooling (grey zone), and cooling (blue zone).

Two points of concern were observed in the case study building; the first point of concern was the unusually large fraction of operating perimeter heaters while the AHU was in the economizer mode. Ideally, there is minimal to no use of the perimeter heaters in this mode since the building is trying to induce cooling; hence the outdoor air damper modulating above its minimum position in the economizer mode. Second, the valve position of the heating coil, even at the peak of the heating season, was minimal; this suggests underutilization of the AHUs' heating coils. At -25°C , the heating coil valve position is only at

about 4% open. The heating coil valve positions are expected to be significantly more open (i.e., near/at 100% [4]) at extreme cold outdoor air temperatures than what was observed. These abnormal trends suggested an identified conflict between the economizer mode and perimeter heating devices. To prevent overheating zones during the heating season, the economizer mode is engaged and subsequently the AHUs' heating coils shut off if the average zone temperature exceeded an overheating zone temperature threshold. However, both the overheating threshold and zone temperature setpoints were 22°C. The identical overheating threshold and the default zone temperature setpoints resulted in this threshold being frequently exceeded since the perimeter heaters, which are controlled through the zone temperature setpoints, operate on reactive-based controls and allow the zone temperature to slightly exceed their setpoints. This resulted in the AHUs operating in the economizer mode in an effort to cool the seemingly overheating building but with perimeter heaters working excessively to combat the intake of sub-zero outdoor air and maintain the zone temperature setpoint. Since the default zone temperature setpoints are 22°C, the threshold for overheating in the heating season to engage the economizer mode was increased to 23.5°C from 22°C to allow the average zone temperature to slightly exceed its setpoint without engaging the economizer mode and differentiate from true overheating.

The second generated visualization includes six diagrams for each of the case study building's two AHUs (i.e., 12 diagrams in total) which provides a visual summary of common AHU operating periods and average actuator and damper positions and temperature measurements. The k-means clustering algorithm was selected based on the Calinski-Harabasz index and was used to group AHU operating periods of similar damper,

valve, and temperature measurements into six clusters for each AHU. Figure 2.6 is one of six snapshots of one of the AHUs which characterizes 25% of the AHU's cumulative operation. These summary diagrams can assist operating staff in interpreting the detected anomalies or in detecting other anomalies that were not detected by the function.

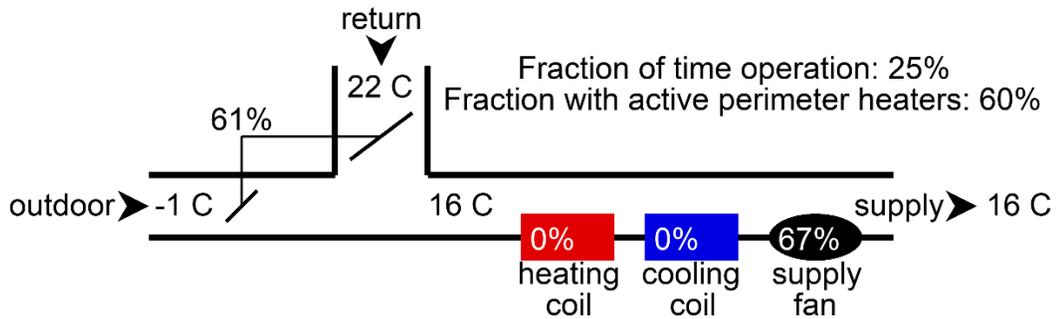


Figure 2.6: Example of simplified AHU diagram snapshot with AHU dampers, valve states, air temperatures, and fraction of time of operation. The fraction of active perimeters for the period of operation is also displayed.

Table 2.2 show the detected hard and soft faults and AHU health index. Low outdoor air was flagged in the second AHU which is typically symptomatic of an outdoor air damper that is stuck closed. However, this is more likely a result of a high fault detection threshold for outdoor air fraction bias. For the second AHU, lower outdoor damper positions tend to produce lower than expected outdoor air fractions, as indicated in Figure 2.7; though not a fault, this resulted in an outdoor air fraction bias that is symptomatic of a stuck closed outdoor air damper. Even so, this can still prove problematic for indoor air quality purposes as outdoor air damper positions as high as 55% only supply about 20% outdoor air flow. The bias threshold is not suitable and should be reduced to avoid raising this false-positive.

Table 2.2: AHU health index and summary of detected soft and hard faults (ahuAnomaly function KPIs)

AHU	Health index (%)	Cooling coil	Economizer	Heating coil	Outdoor air damper	Supply air temperature	Schedule
1	67%	Normal	Normal	Stuck	Normal	Check supply air temperature reset logic	Check mode of operation logic
2	50%	Normal	Normal	Stuck	Low outdoor air	Check supply air temperature reset logic	Check mode of operation logic

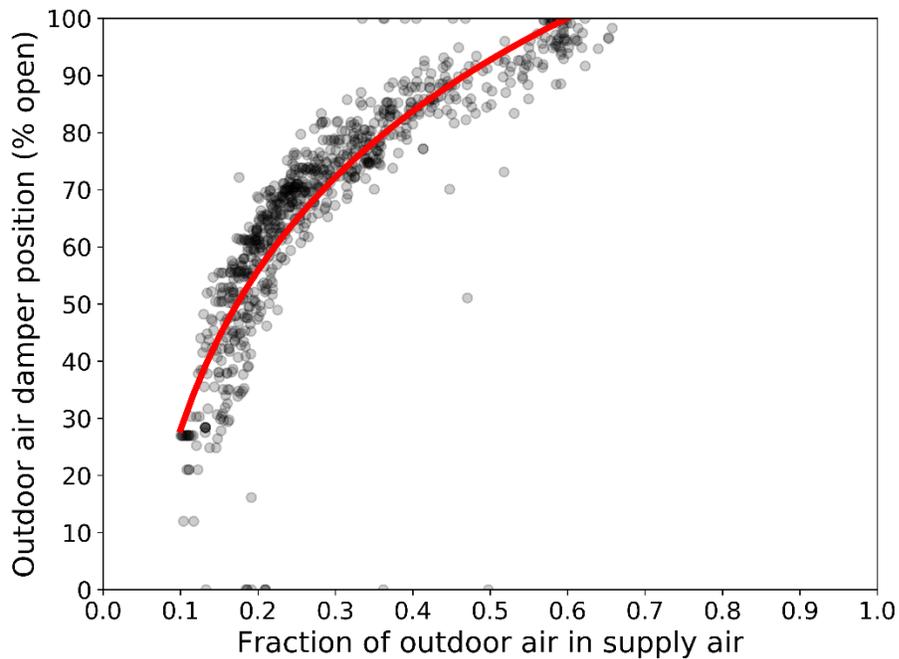


Figure 2.7: Visual representation of relationship between outdoor air damper position and outdoor air fraction for AHU 2. Lower outdoor air damper positions produce notably lower outdoor air fractions in the supply air, resulting in a low outdoor air fraction bias.

An absence of a schedule is detected in both AHUs which indicated operation of the AHUs exceeding 100 hours per week (i.e., approximately 14 hours per day); 100 hours per week is the detection threshold for this fault for this case study. This was a result of an

overheating IT room causing night-cycling. This issue has since been resolved as the unoccupied mode temperature setpoint was increased from 27°C to 28°C.

Both AHUs were flagged for an inappropriate supply air temperature reset logic. The change point temperature between the heating and economizer with cooling mode for both AHUs was estimated to be about -4°C; it is suggested to increase the supply air temperature setpoint as an inappropriately low setpoint may exacerbate the use of perimeter heaters in the economizer mode.

Both AHUs were also flagged for a stuck heating coil valve. This was expected as the conflict between the perimeter heating devices, which operate based on zone temperature setpoints, and the overheating zone temperature threshold to engage the economizer mode deterred the use of the AHUs heating coils.

2.5.4 Detection of zone-level anomalies (zoneAnomaly function)

Per the methodology presented by Gunay and Shi [94], the zone-level anomaly function inputs the zone-level HVAC controls network trend data and groups zones with similar average seasonal errors for airflow and temperature setpoint control. The function outputs a zone health index KPI for the heating and cooling seasons separately which ranges from 0% if all zone clusters exhibit anomalous airflow control errors to 100% if no zone clusters exhibit anomalous conditions. Visualizations depicting the clustered zones with the corresponding indoor air temperature and acceptable ranges for airflow control error accompany the KPI. The function employs k-means, Gaussian mixture, and agglomerative clustering algorithms to group zones into three to five clusters, separately for the heating and cooling seasons. Of the three algorithms, the best algorithm and its resultant number of clusters is selected using the Calinski-Harabasz index and is used to produce the

visualizations and KPIs. Faults are inferred should the mean zone temperature and/or mean airflow control error of a zone cluster fall beyond the pre-defined threshold for normal error. Note that this function may not be suitable for buildings with large core spaces and few perimeter heating devices; as the fraction of active perimeter heater devices is a required input for clustering for the heating season, this function is intended for building's where the majority of zone contains some form of zone-level heating.

Zone-level HVAC controls network data from January 1st, 2019 to December 31st, 2019 were used to demonstrate the zone-level anomaly detection function's ability to detect anomalous zones in the case study building. Figure 2.8 is the resulting visualization for the heating season analysis portion of the function; the cooling season analysis is depicted in a separate plot. The figure indicates which zone clusters suffer from airflow control and zone temperature anomalies and the severity of the anomalies. The airflow control error is the ratio of the error between the actual and the setpoint airflow rate over the setpoint airflow rate. In this analysis, an airflow control error of $\pm 20\%$ and an indoor air temperature of 20°C to 25°C was deemed normal. For the heating season, four clusters of zones were identified using the k-means clustering algorithm. Zone clusters C0, C1, C2, and C3 were comprised of 17, 12, 12, and 1 zone respectively. Clusters C0, C1, and C2 all fall within acceptable airflow setpoint controls errors. The zone cluster C3 deviated from the other clusters and exhibited an airflow control error of -20.4% which is marginally below the threshold for normal airflow control errors. This may be symptomatic of a stuck damper inhibiting airflow to the zone, causing a deviation between the measured airflow and airflow setpoint. If a large number of zones exhibited this issue, it may be indicative of a low supply air pressure setpoint at the AHU-level. The health index, which is the ratio

of zones within the threshold of acceptable airflow control error and indoor air temperature over the total number of zones, remains relatively high at 97.6% since only cluster C3 which contains only one zone is considered anomalous. For the summer, five clusters of zones were identified with the k-means clustering algorithm. The clusters comprised of 27, 3, 2, 4, and 6 zones, none of which exhibited anomalous conditions (within what is deemed normal). Thus, the health index remains at 100%. The breakdown of the plots and the zone health index KPIs are shown in Table 2.3.

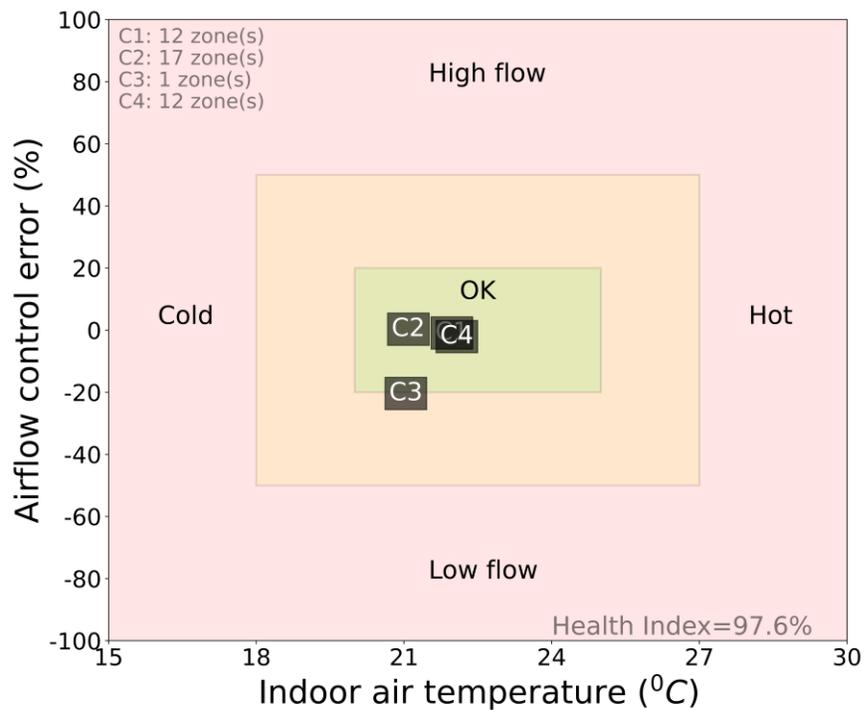


Figure 2.8: Clusters of zones, represented by the black boxes, with similar airflow control errors over indoor air temperatures for the heating season. The airflow control error is defined as the ratio of the difference between the airflow rate and airflow setpoint over the airflow setpoint. The top-left of the plot indicates the number of zones attributed to each cluster, and the bottom-left indicates the zone health index which is the ratio of zones within acceptable conditions over the total number of zones.

Table 2.3: Zone health index and summary of airflow controls errors of zone clusters (zoneAnomaly function KPIs).

Cluster	Heating season					Cooling season			
	# of zones	Mean T_{in} (°C)	Mean Q_{no} error (%)	Mean S_{rad} (%)	Health index (%)	# of zones	Mean T_{in} (°C)	Mean Q_{no} error (%)	Health index (%)
0	17	21.1	0.30	86.4	97.6	27	22.2	-0.21	100
1	12	22.1	-2.07	53.7		3	22.1	-5.12	
2	12	22.0	-0.92	19.5		2	21.9	5.30	
3	1	21.0	-20.4	82.4		4	21.6	1.83	
4	-	-	-	-		6	22.2	-1.39	

2.5.5 End-use disaggregation of meter data (endUseDisaggregation function)

The end-use disaggregation function inputs energy meter data, AHU- and zone-level HVAC controls network trend data, and Wi-Fi device count data and is based on the methodology described by Gunay *et al.* [11] and Darwazeh *et al.* [98]. The KPIs generated by this function are the EUIs for lighting and plug-loads, distribution (i.e., pumps and fans), and chiller for electricity energy use, EUIs for perimeter heating, the AHUs’ heating coils, and other appliances (i.e., domestic hot water) for heating energy use, and EUIs for the AHUs’ cooling coils for cooling energy use; visualizations depicting the weekly distribution of the major end-uses for electricity, cooling, and heating accompany the KPIs. In Gunay *et al.* [11], the GA was used to estimate the unknown parameters of separate disaggregation models for electricity use by lighting and plug-loads, distribution, and chillers. HVAC controls network data, such as fan state, and Wi-Fi data as proxy for

occupancy were used to provide contextual information to the operational state of each end-use, allowing for higher disaggregation resolution. In other words, HVAC controls network and Wi-Fi data were used to inform the disaggregation. Darwazeh *et al.* [98] extended this approach beyond electricity alone by monitoring heating and cooling energy use and using HVAC controls network data such as heating and reheat coil states as contextual information. The end-use disaggregation function develops a load disaggregation model for each of the end-uses for cooling, heating, and electricity, and uses the GA to estimate its unknown parameters and minimize the RMSE between measured and modeled energy use. Each parameter represents a multiplier associated with the fraction of energy use attributed to each end-use. The GA was performed with 10 generations and a population size of 10000. Note that since lighting and plug loads, distribution, and chiller are the assumed major electricity end-uses, this function is not suitable for buildings with other significant electricity end-uses such as elevators and data centers, unless such end-uses are separated from the energy meter data prior to using the function.

As Wi-Fi data were only available in the case study building from April 23rd, 2019 until the end of the year, these data along with energy meter data and AHU- and zone-level HVAC controls network trend data from the same period were used to demonstrate the capabilities of the end-use disaggregation function in the case study building. The generated visualization in Figure 2.9 depicts the weekly distribution of EUI, disaggregated per the major end-uses for electricity, heating, and cooling, and Table 2.4 shows the total EUI for each end-use. Note that due to the misaligned collection dates of data, specifically Wi-Fi data, the first 18 weeks of the plots are omitted. For electricity

consumption, the lighting, plug-in, and distribution loads stay relatively consistent throughout the year whereas chiller loads transition during shoulder seasons and peak in the cooling season. This is not unusual since chiller use is largely seasonal and lighting is not expected to vary drastically by time of year. Likewise, the relatively constant daily energy use for distribution is expected as the building requires active ventilation during occupied hours regardless of season or operating mode. For heating, domestic hot water energy consumption is consistent and is expected. However, the heavy perimeter heating energy use relative to AHU heating coils is a point of concern. The heating coils in both AHUs seem to engage minimally, even at the peak of the heating season. This anomaly was also observed in the AHU-level anomaly detection function where the fraction of active perimeter heaters was consistently high and AHU heating coil valves remained minimal even at peak heating season. Chiller use shows similar load patterns and magnitude during the peak cooling season, though one AHU sees greater use near the shoulder seasons than the other. This is expected since the building will only use one chiller until it alone is insufficient to meet the cooling demand.

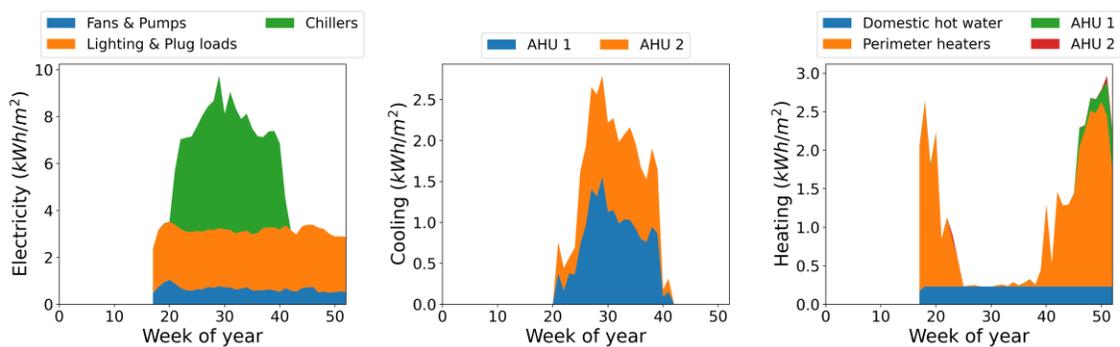


Figure 2.9: Weekly EUI profiles for (a) electricity, disaggregated by lighting and plug-loads, distribution, and chiller, (b) heating, disaggregated by hot water, perimeter heating, and AHU heating coils, and (c) cooling, disaggregated by AHU cooling coils.

Table 2.4: Total EUIs of major end-uses for electricity, heating, and cooling energy use (endUseDisaggregation function KPIs).

Utility	End-uses	Total EUI (kWh/m ²)*
Electricity	Chiller	101**
	Distribution	23.3
	Lighting & plug-in	99.6
Heating	Perimeter heating	34.3
	Heating coil (AHU 1, AHU 2)	0.754, 0.137
	Other	10.1
Cooling	Cooling coil (AHU 1, AHU 2)	18.1, 18.5

*From April 23rd, 2019 until Dec 31, 2019.

**Chiller electricity use includes use by an adjacent building.

2.5.6 Analysis of hot/cold complaints (complaintAnalytics function)

The complaint analytics function inputs CMMS data and extracts operator comments for thermal complaints for room air temperature. This function outputs four KPIs which indicate the daily frequency of hot complaints and cold complaints generated per 1000 m² during both the heating and cooling seasons. This is accomplished by employing the methods developed by Dutta *et al.* [74] whereby thermal complaints which contain the terms “hot” or “cool” are extracted from the CMMS data. As this function focuses on room air temperatures, the term “water” was also included in the search criteria to identify and remove complaints regarding water temperature. Three visualizations accompany the KPIs. The first depicts the hourly and monthly distribution of hot and cold complaints. The second is a decision tree model depicting the daily frequency of hot and cold complaints based on outdoor air temperature, indoor air temperature, and time of day. The third depicts the relationship between indoor and outdoor air temperature which trigger thermal

complaints. Note that the function is limited only to hot and cold complaints, and that building area must be manually inputted.

CMMS data from January 2nd, 2019 to December 23rd, 2019 were used to test the capabilities of the complaint analytics function in the case study building. However, the suitability of this function with the case study building, and its ability to generate any meaningful insight from its visualizations and KPIs, were hindered by the small sample size of hot/cold complaints. In the case study building, a small fraction of the building's occupants are full time employees while the majority are students with intermittent presence. The transient nature of occupancy may be a deterring factor for occupants to report thermal complaints [99], [100]. A building with a more static occupancy would be expected to generate a substantially larger number of complaints, and would allow the function to provide more conclusive insights.

Though the information provided by the function's visualizations and KPIs are inconclusive, the generated visualizations are still provided as an example of how extracted information is presented to the user. Figure 2.10 exemplifies the first set of visualizations generated by the function. These charts are provided to the user and are intended to break down the complaints by month, period of the day, and the type of complaint.

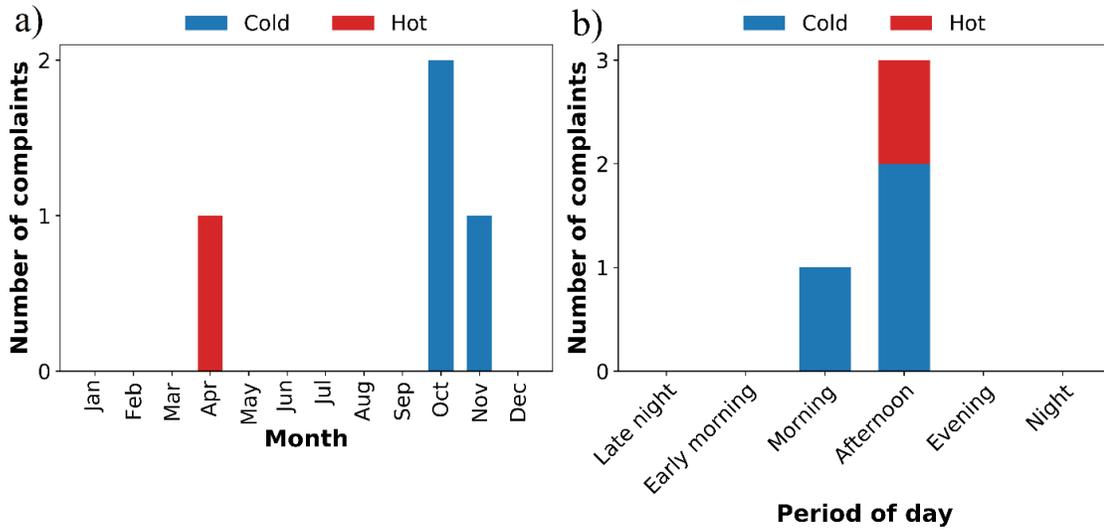


Figure 2.10: Example of first set of visualization generated by the complaint analytics function which provides a breakdown of hot and cold complaints by (a) month and (b) period of day. Due to the small sample size, no relationship between the complaints made and the month or time of day they were made was established.

The function also generates a regression tree which is intended to predict the expected frequency of hot/cold complaints based on the time of day, day of the week, and outdoor air temperature. Figure 2.11 is an example of the visualization.

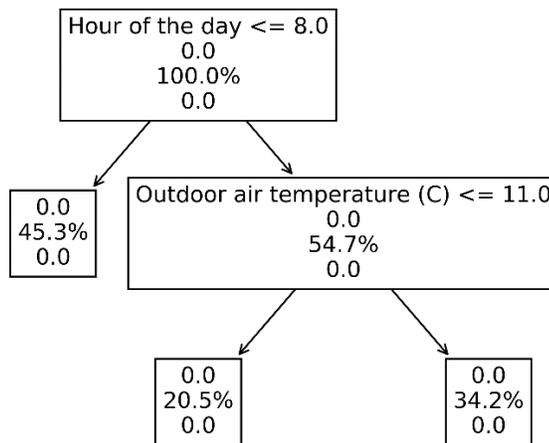


Figure 2.11: Example visualization of a regression decision tree model which is intended to predict the proportion of the cold/hot complaints based on time of day, day of the week, and outdoor air temperature. Note that nodes branching to the left represent the predicted proportion of complaints which satisfies the decision criteria listed in the preceding node. The decision criteria is listed at the top of the branch node and the predicted proportion is listed in between the 0.0s. Due to the small sample size, the decision tree model did not split the proportion of complaints by type of day (weekend or weekday) and the predicted proportions are inconclusive.

Figure 2.12 is an example of the third generated visualization which is intended to depict the complaints in relation to outdoor air temperature and indoor temperature.

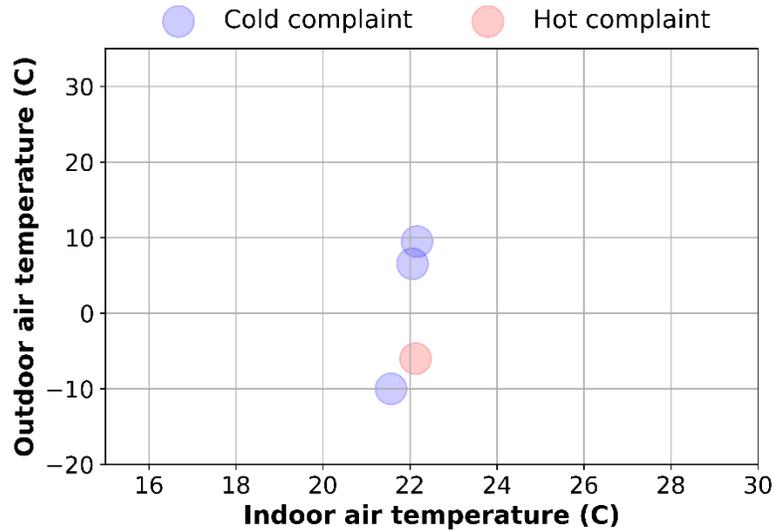


Figure 2.12: Example visualization which is intended to depict hot and cold complaints and their relationship between outdoor and indoor air temperature. Due to the small sample size, no relationship between the complaints made with respect to outdoor and indoor air temperatures was established.

2.5.7 Occupant count estimation (occupancy function)

Based on the methodology presented by Hobson *et al.* [83] and Gunay *et al.* [101], the occupancy function inputs Wi-Fi device count data to generate occupant count profiles of a building for weekdays and weekends separately. The function computes the earliest expected arrival time, latest expected departure time, and highest expected occupancy per floor, which are its KPIs. Visualizations depict the 25th, 50th, and 75th percentile occupant-count profile and the 75th percentile superimposed floor-level occupant-count profile for a typical weekend and weekday. The function uses Wi-Fi device count data as proxy for occupant count. Earliest arrival times are determined by the earliest timestep whereby the 75th percentile occupant count exceeded 10% of the maximum occupant count per floor. Likewise, latest departure times are determined by the latest timestep whereby the 75th percentile occupant count falls below 10% of the maximum occupant count per floor. Note

that the occupant-count proxy assumes 1.2 Wi-Fi devices per occupant; this relationship between occupants and devices was established in the case study building using ground truth data in a previous study [102]. This assumption may not hold true in different buildings and might even change over time as Wi-Fi device protocols evolve and the presence of Wi-Fi enabled devices increases. An occupant count cut-off of three or less occupants per floor was also implemented to minimize the inclusion of artefacts caused by devices like printers and computers that are left on overnight.

Wi-Fi device count data from April 23rd, 2019 to the end of the year were used to test the capabilities of the occupancy function in the case study building. Figure 2.13 shows the resulting plots which depict the hourly occupant count at the 75th, 50th, and 25th percentile for a typical weekday and weekend. Figure 2.14 disaggregates the 75th percentile occupant count per floor for weekdays and weekends. The earliest arrival time, latest departure time, and peak daily occupancy per floor KPIs are shown in Table 2.5.

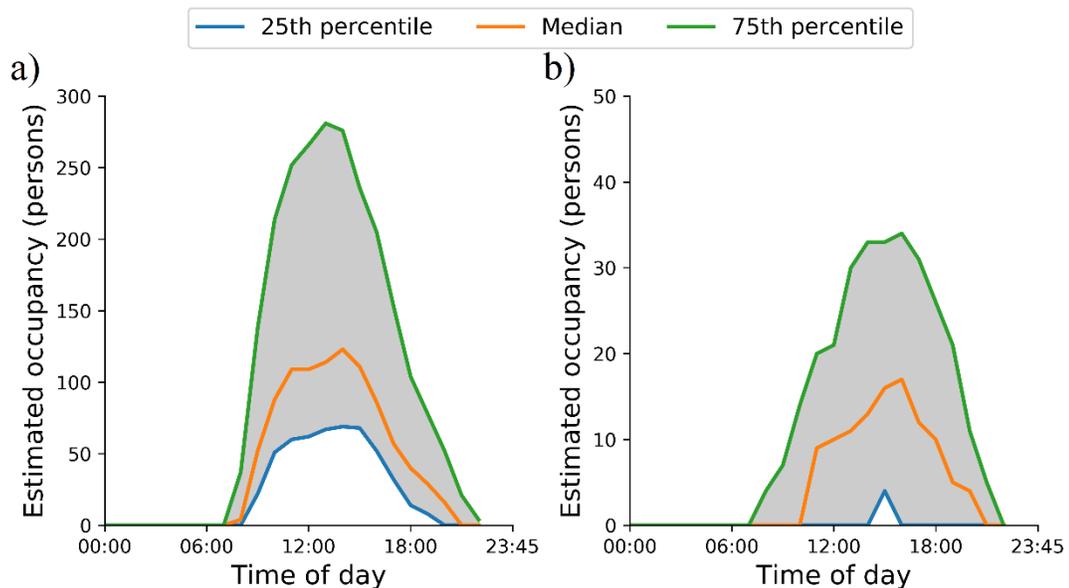


Figure 2.13: 25th, 50th, and 75th percentile occupant count profile for (a) weekdays and (b) weekends.

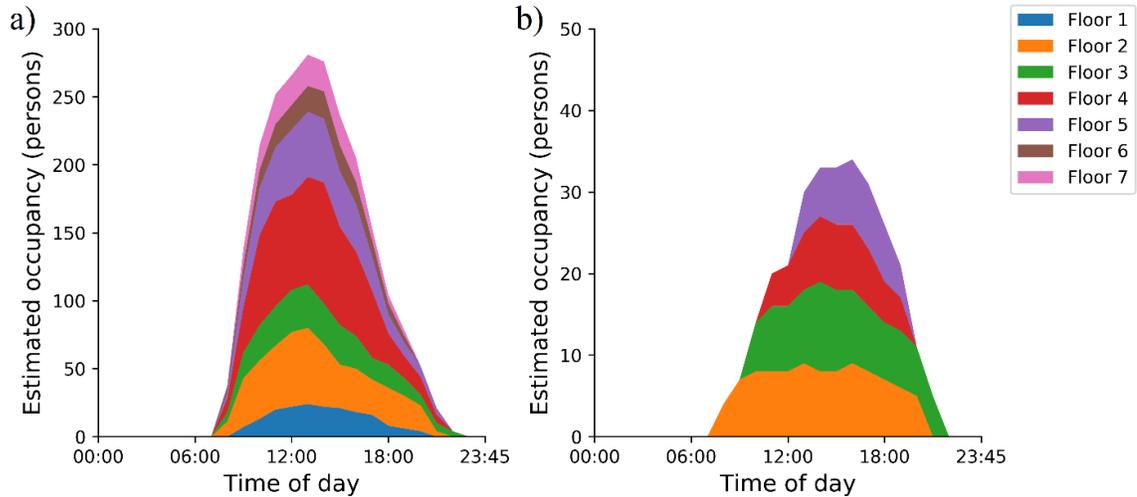


Figure 2.14: 75th percentile superimposed floor-level occupancy count profile for (a) weekdays and (b) weekends.

Table 2.5: Earliest arrival times, latest departure times, and peak occupancy per floor for weekdays and weekends (occupancy function KPIs), and design occupancy per floor.

Floor	Earliest arrival time		Latest departure time		Typical peak daily occupancy		Design occupancy per floor
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend	
1	8 am	-	7 pm	-	24	0	44
2	7 am	7 am	7 pm	7 pm	56	9	352
3	7 am	9 am	9 pm	8 pm	32	11	174
4	7 am	10 am	7 pm	6 pm	89	8	127
5	7 am	12 am	8 pm	6 pm	48	8	141
6	7 am	-	6 pm	-	20	0	98
7	8 am	-	6 pm	-	23	0	104

The total occupancy of the building during workhours peaks at 292 occupants at around noon for the typical 75th percentile workday. This is a drastic reduction from the 1040 occupants during occupied hours that was assumed during the case study building’s design. The plots show the variation in occupant-count throughout the day and can be used to derive occupancy-based ventilation controls, which can modulate ventilation rates to suit

the building's actual occupancy and reduce the overventilation. The generated occupancy profiles and KPIs formed the basis for an occupancy-based predictive control program which implemented different hourly outdoor air damper position profiles in the case study building's AHUs at the beginning of 2020 based on day-ahead occupant count predictions [103].

2.6 Discussion and limitations

The BEM toolkit's multifaceted analytical approach is intended to provide operators with a method to capture and monitor the propagated effects of operational deficiencies through various KPIs and visualizations. The amount of unique analytical methods under one toolkit offers a more comprehensive understanding of building operations, thus allowing for a more flexible and diverse approach in addressing deficiencies and improving operations. Underlying deficiencies are more likely to be captured since energy use anomalies and suboptimal operations undetected from one function can be detected from a number of the other functions. Furthermore, since many KPIs and visualizations are derived from the same data, underlying deficiencies can be captured from multiple unique interpretations of the data which can be consolidated to focus fault correction efforts to specific components or systems. In fact, these redundant indicators between functions are encouraged if they suggest a mutual deficiency since it strengthens the confidence of the toolkit's detection capabilities. For example, recall for the case study that the low heating SE KPI derived from the baseline energy function indicated excessive heating energy use outside of the AHUs' operating hours. Should analysis be limited to this KPI, a number of different root causes should be examined such as a lack of an AHU heating schedule, conflicts in the operational logic, or excessive heating demand during operating hours due

to overventilation. However, Figure 2.5 from the AHU anomaly detection function indicated excessive use of the perimeter heating devices; this issue was also seen in Figure 2.9(c) from the end-use disaggregation function. Presented with this additional insight, an end-user would be compelled to investigate the relationship between excessive heating energy use and the perimeter heating devices. The root cause was determined as a conflict between the perimeter heating devices, which operate based on zone temperature setpoints, and the zone temperature overheating threshold to engage the AHUs' economizer mode which led to excessive heating energy use; an overheating IT room propagated this issue to outside scheduled AHU operating hours. These issues have since been resolved. Though this example demonstrates the interrelated nature of the toolkit's functions, and the benefit of multiple interpretations of the data, it also demonstrates that the toolkit's ability to generate energy-saving insight is limited by the building operator's interpretation of the generated KPIs and visualizations. The scope of the toolkit's capability is limited to detection and monitoring of potential operational deficiencies whereas diagnosis and correction is left to the operator's discretion.

One of the inherent challenges of implementing a BEM toolkit to multiple buildings is the varying availability and granularity of data in different buildings. Underlying factors such as HVAC setups and physical building characteristics can also affect the toolkit's ability to generate reasonable and accurate KPIs, particularly for the anomaly detection functions that rely on detection thresholds which may exclusively suit a case study building. Consideration for varying accuracy and availability of installed sensing technologies in other buildings, as a result of different vendor practices or effective lifespans of sensors, should also be taken into account since these factors can also affect the reliability of the

outputted information. Though this toolkit was never intended as a universal solution but rather a framework from which derivations of BEM solutions may be fostered, its preliminary state, including its structure and methodologies, should at least demonstrate that a toolkit encompassing multiple domains of building operations data analytics can diversify opportunities to address operational deficiencies within an intended scope of building type and use; in this case study, medium to large commercial offices. Regarding availability of data, the functions operate independent of each other such that if a lack of certain data or data quality restricts one function, other functions that do not use that data can still be implemented and generate KPIs; this adds flexibility to the toolkit's overall implementation. It is expected that not all functions will be useful for every building. For example, the complaint analytics function was not well suited for the case study building since the building is primarily occupied by students who are intermittently present and less likely to complain. In this case, a low frequency of complaints should not suggest optimal indoor air setpoints but rather that the function cannot produce conclusive insight for the building. In addition, each function was also developed to input as little data as possible to generate KPIs, even if it meant excluding data streams that were available in the case study building and could have been used to improve reliability or accuracy of the KPIs.

The toolkit's seven selected functions were built upon established data-driven methods which are familiar to the building energy performance community, and is intended to provide a basis for future development of multi-source solutions. A preliminary version of the toolkit is publicly available online for others to learn from, adapt, and foster into more specialised versions of multi-source, data-driven toolkits. The major domains which were identified to attribute considerable improvements to building energy performance were

energy benchmarking, AFDD, and OCC. Each function, excluding the metadata inference function (*metadata*), can be categorized to address one of these three domains. Where multiple functions address a single domain, different data sources or combination of data sources are employed to allow flexibility of data availability in varying buildings.

2.7 Closing remarks

The framework for a multi-source, data-driven BEM toolkit as a synthesis of established data-driven approaches from the literature has been proposed with its inputs, algorithms used, and generated visualizations and KPIs. The toolkit is intended to provide building operators with a method to address operational deficiencies such as energy use anomalies and inappropriate schedules by interpreting metadata, detecting faults and upgrading sequences, and monitoring KPIs. The current capability of the toolkit was demonstrated using various disparate types of archived operational data from an academic building. Though energy use anomaly detection was discussed extensively, the toolkit's ability to produce energy benchmarking and occupancy-based analytics should not be overlooked as these insights are similarly instrumental for deriving energy-saving decisions. Should building operators capitalize on the toolkit's flexibility, the toolkit can be implemented in a wide range of different buildings.

A preliminary version of the toolkit is publicly available through GitHub (link in Appendix A) and is intended to facilitate community-driven efforts to refine the existing functions or serve as a basis upon which more specialized multi-source, data-driven solutions can be derived. A push to refine the functions should not be limited to improving reliability but should also work to improve robustness. Limiting data sources to what is absolutely required to generate the desired insight may not necessarily reduce the accuracy of the KPIs

but would allow the toolkit to operate in buildings that lack high resolution data. Derivations of multi-source toolkits may also incorporate additional, reduced, or altered functions to suit a particular set of buildings, or even establish explicit interdependencies between functions. Currently, the only interdependency between functions exists with the metadata inferencing function. Even then, it only serves to prepare the data for input in subsequent functions. A potential benefit of using the same set of data is that the generated KPIs and visualizations may concurrently suggest the same deficiencies. Such redundancies, as demonstrated by the baseline energy, AHU anomaly detection, and end-use disaggregation functions' role in detecting a conflict between the AHUs' economizer mode and the perimeter heating devices, would add a layer of confidence to the toolkit, especially in its anomaly detection capabilities, and can minimize the likelihood of false-positive detection. Revisions to the toolkit should attempt to integrate the KPIs and develop likelihoods of true faults in addition to just simple detection. If two or more functions generate seemingly corroborative insight, the likelihood of a specific fault or number of faults would increase. By contrast, should conflicting insights exist, the likelihood would decrease.

Future case studies will demonstrate the capabilities of the toolkit on other commercial and institutional buildings and the results will be disseminated to showcase its energy-saving potential and the need to further develop data-driven approaches to BEM. Furthermore, interviews with building operators will supplement understanding of how the toolkit's generated KPIs are actually used by operators, inform the development of performance metrics and visualization, and further improve the existing capabilities of the toolkit.

Chapter 3

This chapter has been accepted for publication as:

Markus, Andre A.; Hobson, Brodie W.; Gunay, H. Burak; and Bucking, Scott.

FRAMeWORK: A multi-source, web-based application to identify suboptimal energy use management. *Proc. of IBPSA-Canada's eSim 2022 Conf.*

3.0 FRAMEWORK: A web-based application to derive insights

3.1 Introduction

With building operations accounting for 80%-90% of a building's total life cycle energy [2], remediating energy use deficiencies resulting from suboptimal building operations and mechanical faults is paramount to addressing the widespread 'performance gap' and reducing carbon emissions [17], [104]. In the past two decades, attempts to address these deficiencies have manifested as various methods that identify energy use anomalies [32], [43], hard and soft faults resulting in excessive consumption [39], [105], and opportunities to reduce energy use resulting from overventilation [71] using only archived data streams, typically from BASs.

Several publications have proposed and demonstrated methods for AFDD, which pertains to the identification and diagnosis of hard and soft faults in a building's HVAC system [45]. Data-driven variants of this method, also known as black-box models [39], have been demonstrated using real-world data to identify inefficient operating states or achieve a certain level of energy savings from the remediation of identified faults [95], [106]. Likewise, data-driven energy anomaly detection and OCC approaches have been applied to directly remediate identified energy use anomalies and deficiencies in case study buildings [71], [107]. Multi-source data-driven approaches incorporating a combination of different data types have also been demonstrated to produce additional or refine existing energy-saving insights [84], [98].

Though case studies provide an effective and relatable benchmark for a method's capabilities, single-source data from one building may not be indicative of a method's robustness, especially considering the diverse characteristics and functions of the existing

building stock which may influence data quality and trends. Novel data-driven approaches have rarely been tested using data streams from more than one building and remain primarily proof-of-concept. This limitation may hide certain shortcomings of a method resulting from unforeseen variations such as building type, physical characteristics, and purpose. Expanding case study datasets to incorporate multiple buildings can expose limitations and allow further development and refinement of a method if widespread adoption by practitioners and researchers of such novel approaches is intended.

This chapter demonstrates the current capabilities of a novel multi-source, data-driven building energy management toolkit with four case study buildings. Available AHU- and thermal zone-level HVAC control network, energy meter, and motion detection data were extracted from each building and inputted into up to four of the toolkit's seven functions and the KPIs and visuals were presented and examined. The findings were used to further refine the toolkit's data pre-processing capabilities, processing robustness, and visual presentation. The toolkit is open-source and a web interface has been established to facilitate user input and result extraction, as well as to disseminate the toolkit's objectives.

3.2 Software platform architecture

The multi-source, data-driven building energy management toolkit is a novel software application that incorporates established data-driven inverse energy modelling, FDD, load disaggregation, and occupancy and occupant complaint analytic methods. The toolkit currently contains seven discrete functions which input one or a combination of AHU- and thermal zone-level HVAC controls network, energy meter, CMMS, Wi-Fi device count, and motion detection data, and output a variety of visualizations and KPIs intended to provide energy-saving insights for building energy professionals.

3.2.1 Backend

The toolkit is compiled in Python as a functions library and is stored in a public GitHub repository; the link is found in Appendix A. Each function, which can be invoked individually, employs a unique methodology to generate its own unique set of KPIs and visualizations. Upon selecting a function, a user is prompted to supply the required data for the function. The user-supplied data are stored in an intermediate folder (labelled “Unprocessed”) and the corresponding function automatically begins analysing the data if data are detected in this folder. Once analysis is successful, the function returns a report-style document containing the generated KPIs, visualizations, and descriptions of the KPIs and visualizations including how to interpret and utilize them to remediate building energy deficiencies; this is stored in the “Processed” folder for user retrieval. If an error occurs which prevents the function from completing its analysis, an error prompt is returned instead. Figure 3.1 illustrates the backend workflow of this process which periodically checks the “Unprocessed” folder for user-supplied data and processes the data through the appropriate function.

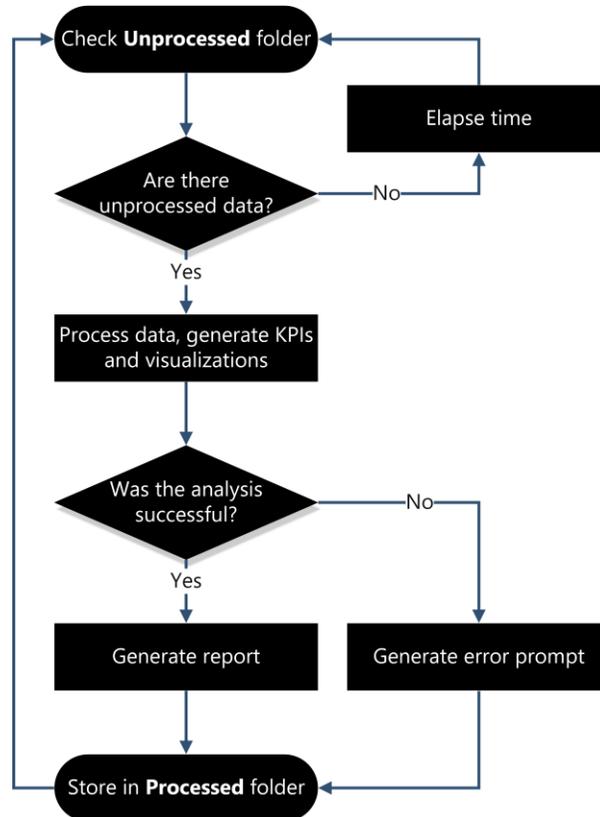


Figure 3.1: Backend workflow for user data storing and processing

Currently, the toolkit is comprised of seven discrete functions. These functions were demonstrated on a case study building in Ottawa, Canada, alongside their methodologies, reference literature, and significance of their KPIs [108]. The first function is the metadata inference function, which inputs a list of BAS-monitored object labels and respective address IDs and automatically identifies and associates the labels by type (i.e., temperature sensor, outdoor air damper position, etc.) and associating AHU or thermal zone. This function does not provide any energy-saving insights; it is intended to help users prepare the input data required for subsequent functions.

The second function is the baseline energy performance function which assesses a weekly AHU schedule’s ability to reduce after-hours energy consumption by comparing building-level energy use during and outside scheduled AHU operating hours. This function inputs hourly energy meter data, disaggregated by heating, cooling, and electricity use, and is

primarily intended to inform users who seek to recognize the actual effect, intended or otherwise, of a schedule's capacity to reduce overall energy consumption.

The third function is the AHU anomaly detection function which automatically identifies three common hard and three common soft faults associated with AHUs. These faults range from stuck outdoor air dampers and heating/cooling coil valves, to general energy use deficiencies resulting from inappropriate scheduling or supply air temperature (SAT) reset strategies. The fourth function is the zone anomaly detection function which identifies anomalous thermal zones based on their indoor air temperature and relative air flow control error. Both these functions require HVAC controls network data consisting of hourly data of various AHU- and thermal zone-level temperature sensors, flow rate sensors, and damper and valve positions. These functions serve to inform users of common issues with AHUs and VAVs which may result in excessive energy use or occupant discomfort if left unaddressed.

The fifth function is the end-use disaggregation function which disaggregates bulk energy consumption into common major end-uses; these end-uses are lighting and plug-loads, distribution (i.e., pumps and fans), chillers, perimeter heating devices, heating coils, cooling coils, and other appliances (i.e., domestic hot water). This function requires energy meter, HVAC controls network, and Wi-Fi device count data to operate, and can inform users of anomalous energy use in specific end-uses.

The sixth function is the occupant complaint analytics function which inputs CMMS work order logs and automatically extracts and analyses thermal complaints (i.e., hot and cold complaints) in the context of the prevailing outdoor conditions and equipment state. This function informs users of the conditions which most prominently give rise to complaints

and can be used to diagnose VAV and AHU operating conditions resulting in occupant discomfort.

The seventh function is the occupancy function which creates building- and floor-level occupant count profiles and computes arrival and departure times. This function can input either Wi-Fi device count or motion detection data, however, no visuals and reduced KPIs are generated using motion detection data. This function can be used to inform occupancy-centric control decisions (e.g., ventilation rates, scheduling).

Depending on the selected function, the amount of data supplied, and the processing hardware, the data analysis may take upwards of several minutes to complete, especially for functions which utilize a genetic algorithm to estimate model parameters.

3.2.2 Web interface

To facilitate data input and results retrieval, a web-based user interface was established to provide user interaction with the toolkit's functions. The web interface, which is hosted off-site using a virtual private server, contains descriptions of each function as well as instructions for data pre-formatting, inputting, and interpreting the generated outputs (i.e., KPIs, visualizations), along with examples of generated outputs. Figure 3.2 illustrates the intended workflow between users, the web interface, and the server.

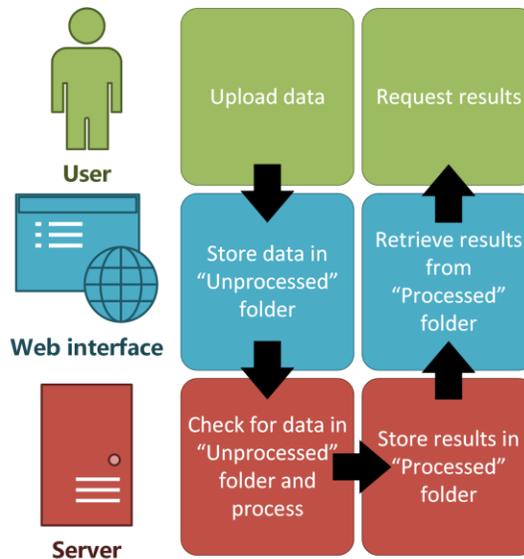


Figure 3.2: Workflow between users, web interface, and the server.

The web interface allows any user with access to the internet to directly upload data, invoke the functions, and retrieve the automatically generated report-style document containing KPIs, visuals, as well as descriptions of how to interpret the outputs; currently, seven discrete functions can be invoked through the web interface. Since the user is responsible for inputting the necessary data and formatting the data which is acceptable for each function, the web interface also serves to inform users of what data are required and provides examples for the input format. Figure 3.3 is an example of the upload webpage which allows users to upload and run the AHU anomaly detection function. Once the analysis has concluded, the user is able to retrieve the report through a randomly generated link which is automatically provided to the user upon successful data upload. Should the user attempt to download the report prior to a completed analysis, or if an error prevents the function to complete its analysis and generate a report, the user is informed accordingly. The web interface is compiled as several HTML and CSS files and are stored in the same GitHub repository as the function library; the link to the repository and web application is found in Appendix A.

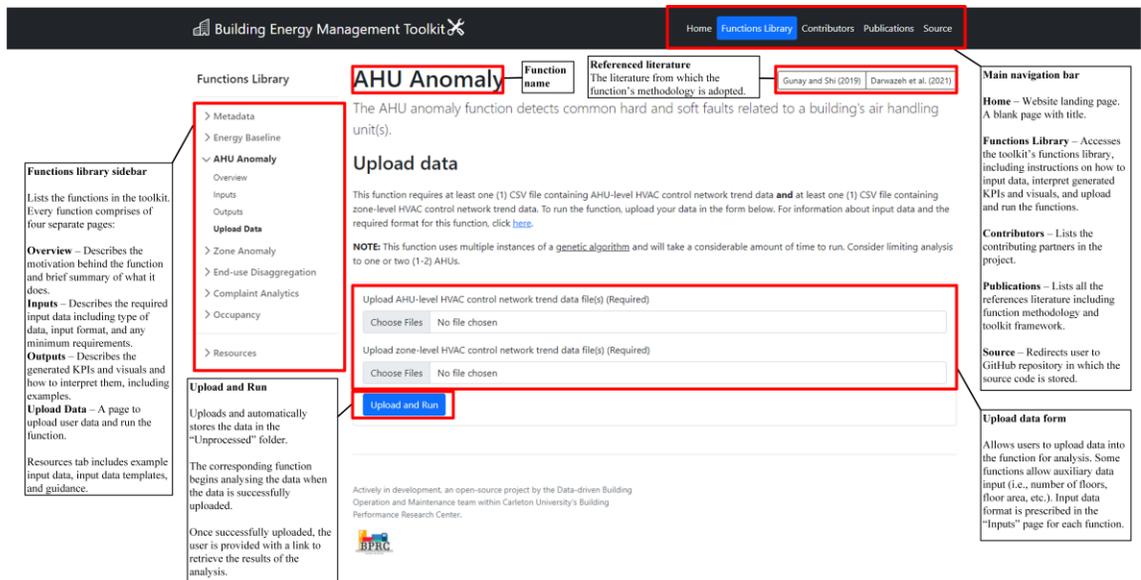


Figure 3.3: Annotated screenshot of upload page of the AHU anomaly detection function of the toolkit’s website.

3.3 Demonstration of the toolkit

To assess the toolkit’s current capabilities and robustness, four separate case studies were conducted. These case studies involved procuring building data of various types and quantity from four different buildings, each attributed to a different organization, formatting and inputting data through the toolkit’s frontend web interface, and retrieving the automatically generated reports. Data were solicited through the authors’ professional network and obtained through each organization’s dedicated energy or BAS monitoring dashboard, or provided through an organization representative. The availability of the collected, pre-processed data for each organization is presented in Table 3.1. The reference period for all data was from January 1, 2019 to December 31, 2019, with the exception of data from organization C with a reference period from February 14, 2019 to March 13, 2021. All buildings were located in Ottawa, Canada, except that of organization D’s, which was located in Borden, Canada.

Table 3.1: Availability of building data by organization

Organization	HVAC controls network		Energy Meter	Motion detection
	AHU	Zone		
A	✓	✓	✓	✓
B	✓	✓	✓	✓
C	✓	✓	✓	
D	✓	✓		

AHU-level HVAC controls network data consisted of hourly air temperature, return air temperature, outdoor air temperature, outdoor air damper position, heating coil valve position, cooling coil valve position, and supply fan state for each AHU. Zone-level HVAC controls network data contained hourly supply air-flow rate, supply air-flow rate setpoint, indoor air temperature, and fraction of active perimeter heating devices. Energy meter data contained hourly measured energy consumption for heating, cooling, and electricity separately. However, for organization A, the electricity consumption included chiller energy use from two nearby buildings which could not be separated, and thus could not be used.

Due to minimum data requirements imposed by certain functions, up to four functions per organization were used; these were the AHU anomaly detection, zone anomaly detection, energy performance benchmarking, and occupancy functions. Table 3.2 presents the number of AHUs and thermal zones analysed; these data were inputted into the AHU and zone anomaly detection functions. Note that the number of AHUs and zones analysed is not indicative of the total number of AHUs or zones attributed to the respective building. Certain AHUs and zones were omitted due to missing or insufficient data quantity.

Table 3.2: Number of AHUs and thermal zones analysed by organization

Organization	Number of AHUs analysed	Number of thermal zones analysed
A	4	98
B	2	16
C	12	32
D	5	73
Total	23	219

The collected data were formatted for input as prescribed by the toolkit’s functions and inputted through the web interface. For the baseline energy function, the available energy meter data were uploaded. However, the baseline energy function was known to produce inadmissible results if electricity use data from buildings with electricity-based heating is inputted; this dismissed the use of metered electricity data from organization B and C. The AHU anomaly detection function was inputted AHU- and zone-level HVAC controls network data from each organization while the zone anomaly detection only inputted the latter. The occupancy function was inputted motion detection data from each organization. Once data have been inputted, the invoked functions automatically generated a report-style document containing the generated visuals and KPIs; the reports were retrieved through an encrypted link that was provided upon successful data input. In total, 13 reports were generated.

3.4 Results

3.4.1 Baseline energy performance

The baseline energy function generated the ASHRAE Guideline 14 three-parameter univariate change point models for organizations A, B, and C, seen in Figure 3.4, for

heating (left column) and cooling (right column) energy use separately. In all cases, the heating energy use rate during (i.e., workhours) and outside (i.e., afterhours) scheduled AHU operating hours were similar, indicating an ineffective reduction of heating energy during afterhours. Organization C exhibited a worse-case scenario of identical workhours to afterhours energy use rates and this is reflected by the resulting 0% SE and highest AEUR of 65%; these are symptomatic of near negligible reductions to heating energy during afterhours. Albeit identical in rates, the afterhours model is slightly offset towards lower temperatures. This can be interpreted as the building having not implemented a schedule for the analysis period, or relying entirely on a simple temperature setback to achieve lower energy consumption. If no schedule existed, implementing one can more effectively reduce afterhours heating energy consumption and subsequently increase and decrease the SE and AEUR score, respectively, than a simple setback. If a schedule did exist, it was likely not working as intended as it produced minimal effect. The effects of internal heat gains through occupancy may result in lower-than-expected workhours rates. The off-schedule operating nature of hydronic perimeter heaters may also result in higher afterhours heating energy use rates.

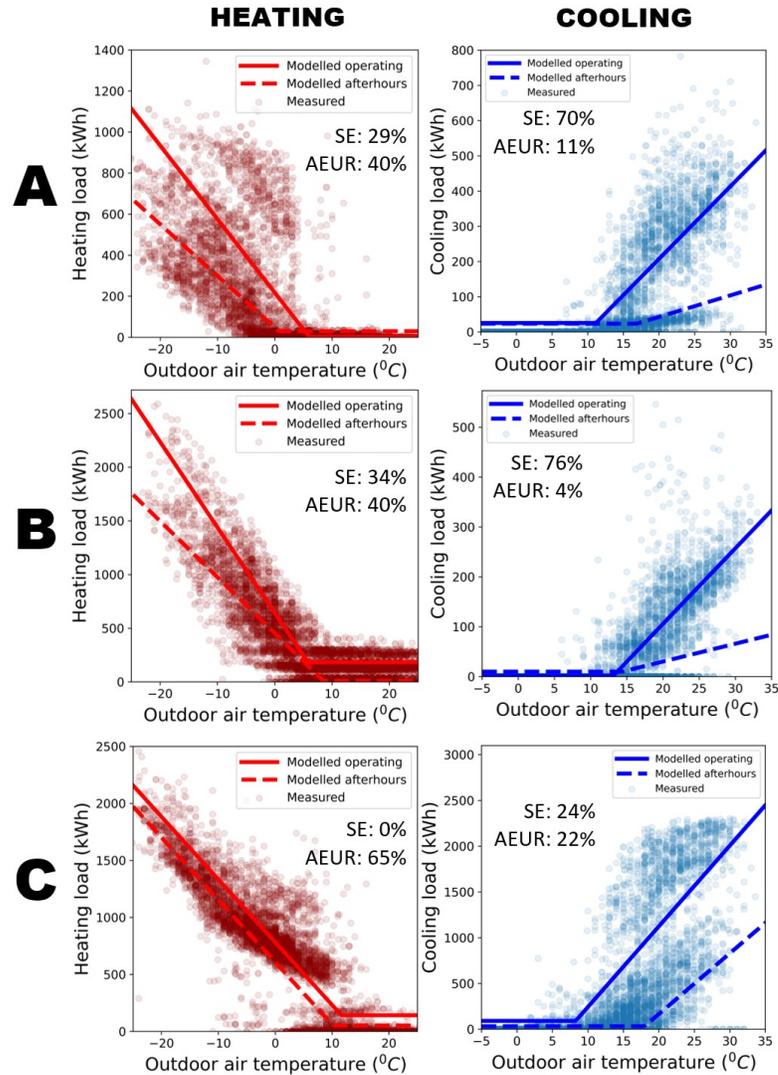


Figure 3.4: Heating (left column) and cooling (right column) energy use during (solid line) and outside (dashed line) AHU operating hours for organization A, B, and C.

As with heating, cooling energy use for Organization C exhibited near identical use rates, suggesting minimal reduction of cooling energy during afterhours. For cooling, a greater overall difference was observed between workhours and afterhours energy use rates compared to heating energy use. Considerations for internal heat gains and overall lower afterhours temperature induce the opposite effect as heating energy. Greater occupancy during workhours may command greater workhours cooling energy use and a lower overall temperature during afterhours, coupled with less solar heat gains, would further minimize cooling use during afterhours.

3.4.2 AHU anomaly detection

The AHU anomaly function produced a plot comparing SAT and return air temperatures for each analysed AHU, seen in Figure 3.5. In almost all AHUs in organization A and C, there was evidence of a SAT reset strategy present, some closely resembling the ideal region highlighted by the schematic. However, a reset strategy was not exhibited in any AHU for organizations B and D; SAT remained static for the entire range of outdoor air temperature and in some AHUs in organization D, SAT exhibited a contradicting scheme whereby the temperature would be higher in the cooling season. The return air temperature for the coolest and warmest rooms were also presented in the schematic and could be used to inform zone-based SAT reset strategies where SAT is modulated based on the coolest room in the heating season. However, the contradicting strategy posed by organization B and D are clearly anomalous and, as it presents no potential for energy reduction, is unlikely to be intentional.

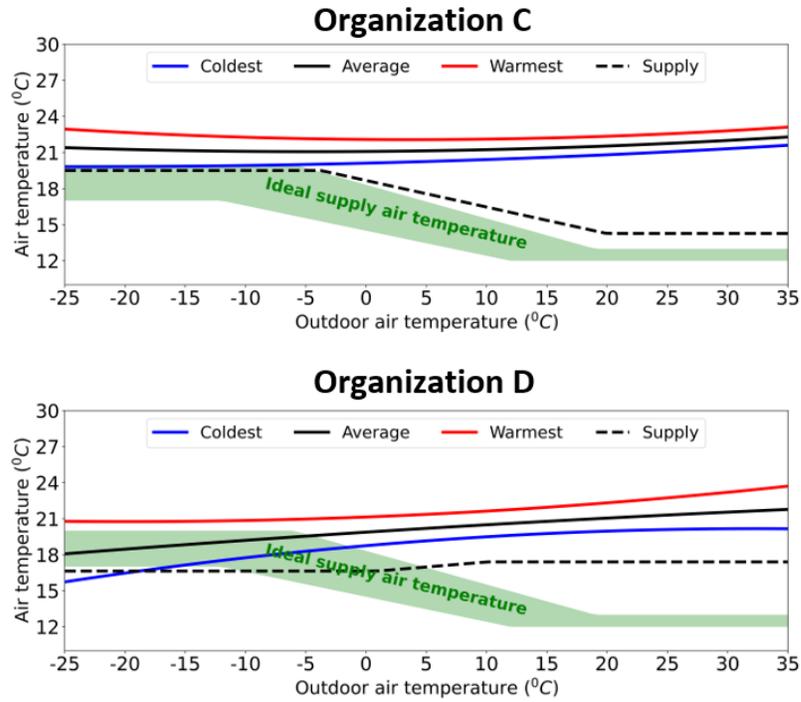


Figure 3.5: Comparison of split-range controller diagrams for AHUs in organization C and D. SAT follows a near ideal scheme for AHUs in organization C but exhibits a contradicting scheme in organization D.

Table 3.3 and Table 3.4 tally the number of hard and soft faults, respectively, identified by the AHU anomaly function by organization.

Table 3.3: Number of identified hard faults by organization

Organization	Number of identified hard faults		
	Stuck outdoor air damper	Stuck heating coil valve	Stuck cooling coil valve
A	1	0	0
B	1	1	0
C	2	1	0
D	1	0	0
Total	4	2	0

Table 3.4: Number of identified soft faults by organization

Organization	Number of soft identified faults		
	SAT reset logic	Schedule	Economizer
A	1	2	0
B	0	0	2
C	0	2	7
D	4	5	3
Total	5	9	12

Of the analysed AHUs, 30% exhibited at least one hard fault, 78% exhibited at least one soft fault, and 17% exhibited no faults. Of the six possible identified faults, the soft “Economizer” fault was the most prevalent and was flagged in the majority of AHUs in three organizations. This indicated inappropriate or non-existent economizer with cooling state settings and can be triggered by a number of factors. In most cases, the outdoor air damper did not exceed 70% open in its economizer with cooling state, despite ASHRAE Guideline 36 [4] recommending 100% (fully-open). For organization B, the maximum outdoor air damper position was as low as 20% open in the economizer with cooling state, which was exhibited in the generated split-range controller diagram in Figure 3.6.

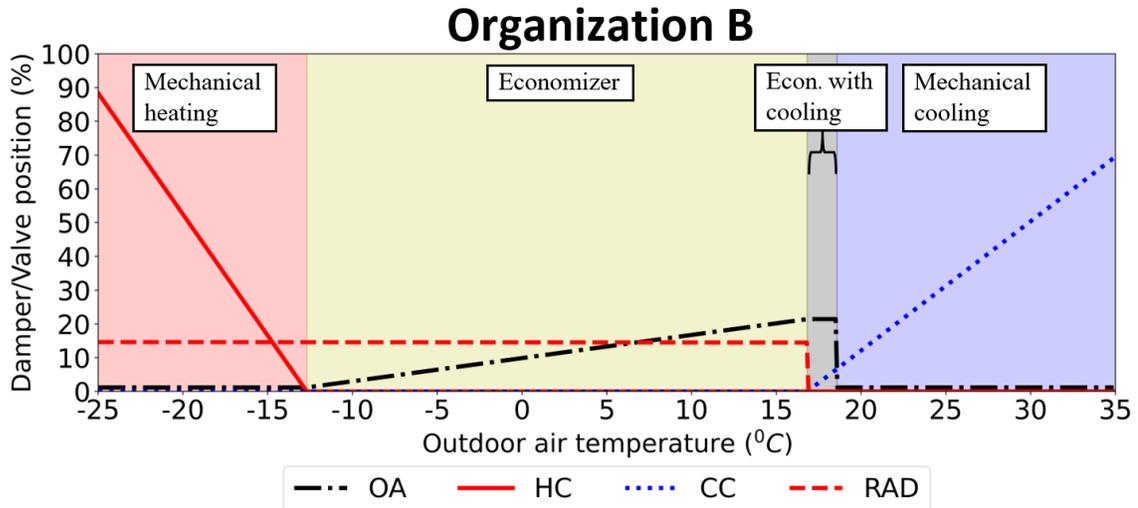


Figure 3.6: A split-range controller diagram for an AHU in organization B. Outdoor air damper position (OA), heating coil valve position (HC), cooling coil valve position (CC), and fraction of active perimeter heaters (RAD) are plotted along outdoor air temperature.

In some cases, the transition from the economizer with cooling to mechanical cooling state occurs prematurely (at outdoor air temperatures less than 15°C). In either case, the AHUs were unable to fully capitalize on cooler outdoor air to reduce energy that would otherwise be used for mechanical cooling. The second most commonly flagged fault was the “Schedule” fault which indicated excessive operation (i.e., greater than 100 hours per week) of the AHUs. This fault was most prevalent in organization D’s AHUs. The “SAT reset logic” fault was also frequently flagged in organization D’s AHUs, which indicated high perimeter heater use in the economizer state due to a lack of a SAT reset strategy. Though a reset strategy was also not observed in organization B’s AHUs, this fault was not flagged; this is likely due to the low intake of outdoor air which aided in minimizing perimeter heating use.

3.4.3 Zone anomaly detection

The zone anomaly function produced two zone cluster diagrams per organization, with the heating and cooling seasons plotted separately. Each cluster diagram was accompanied with a resultant zone health index and a table of zone names (by file name) with the

associated zone cluster. The zone health index is the ratio of zones within the acceptable range of indoor air temperature and airflow control error over the total number of analysed zones. Table 3.5 summarizes the zone health indices per season and organization.

Table 3.5: Summary of zone indices by organization

Organization	Zone health index (%)	
	Heating season	Cooling season
A	100	85.7
B	68.8	75.0
C	93.8	90.6
D	63.0	98.6

Only 15% of zones in the heating season and 10% of zones in the cooling season deviated from acceptable conditions. In most cases and for both seasons, zones which fell outside acceptable conditions gravitated towards cooler indoor air temperatures and negative airflow control error, which suggests a lesser than intended volumetric airflow rate. This trend was also exhibited in the resultant zone cluster diagram for the heating season seen in Figure 3.7 for organization D.

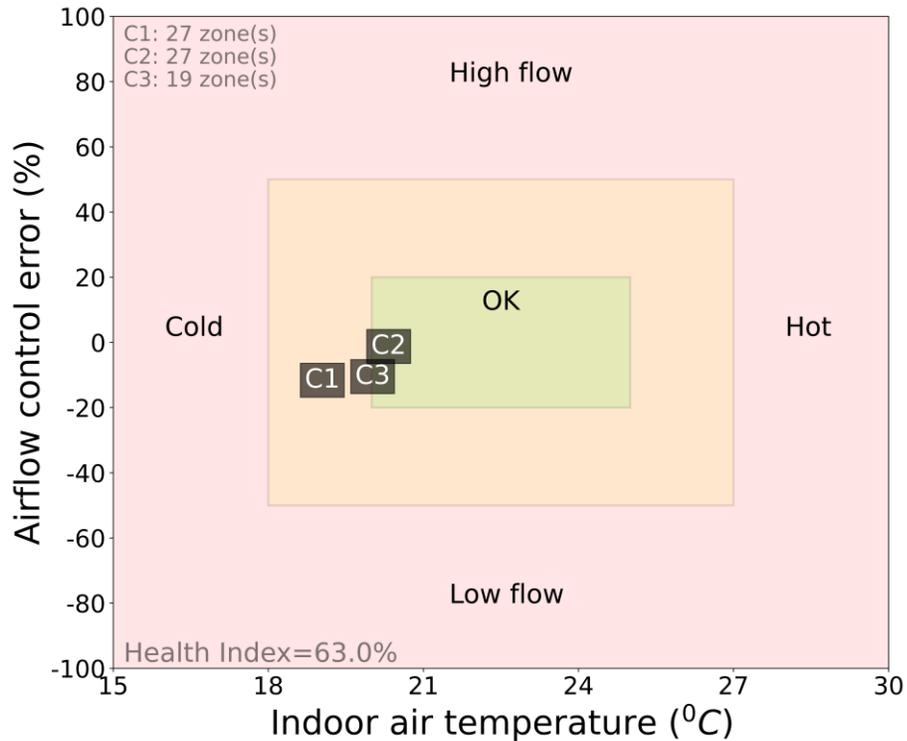


Figure 3.7: Heating season zone cluster diagram for organization D.

Low air-flow may be symptomatic of a stuck or uncalibrated VAV terminal damper, and an abnormally low indoor air temperature may be the result of a faulty reheat coil or perimeter heating devices, or excessive air flow rates. In the case of the latter, the minimum air-flow setpoint for these zones should be decreased. Although the function pinpoints zones which fall outside predetermined acceptable conditions, it does not attempt to identify any root causes such as leaky dampers or malfunctioning heating coils. Users would be advised to conduct further investigation on corresponding VAVs to determine if any faults did produce anomalous zone conditions.

Since the function samples average zone conditions over the analysis period (i.e., a three-month period for the heating and cooling season separately), anomalous zones should not be interpreted as short-term deviations from acceptable conditions, possibly resulting from occupants' impacts. Rather, anomalous zones should be interpreted as zones having

exhibited consistently anomalous conditions over the analysis period, which would be symptomatic of equipment malfunction.

3.4.4 Occupancy using motion detection

The occupancy function did not produce any visualizations since motion detection data, as opposed to Wi-Fi device count data, were inputted. However, the function computed the typical earliest arrival, latest arrival, latest departure, and longest break duration seen in Table 3.6.

Table 3.6: Earliest arrival and departure, latest departure, and longest break duration by organization

Organization	Earliest arrival time	Latest arrival time	Latest departure time	Longest break duration (h)
A	9:00	10:00	17:00	4
B	7:00	14:00	19:00	4

The computed times indicated a narrower period of arrival for organization A. Whereas the time elapsed between when occupants typically first and last arrive is one hour for organization A, this period is seven hours after the first arrival time for organization B. However, organization B exhibited an earlier first arrival time, later last departure time, and an overall longer period of occupancy of 12 hours rather than 8 hours for organization A. The computed KPIs can be used to inform building- and zone-level schedules to minimize excess airflow where there is vacancy. The earliest arrival time and latest departure times can suggest an optimal AHU operating schedule between 9 am and 5 pm for organization A and between 7 am and 7 pm for organization B. A simple logic may be implemented to shut off airflow to zones which have not had their motion detection sensor

triggered by the latest departure time, or if the motion detection sensor has not been triggered for a period excess of the longest break duration.

3.5 Discussion

Using the novel toolkit, various types of archived data from four buildings were analysed and the results were extracted. In the baseline energy performance reports, energy meter data from organizations A, B, and C were analysed. The afterhours and workhours heating energy use comparison were consistently more similar than cooling energy use. In other words, the SE was consistently higher for cooling energy use than heating. This is expected if one considers the nature of hydronic perimeter heating. In the AHUs' unoccupied mode, cooling energy can only be used if a subset of rooms exceeds a temperature setpoint. However, hydronic perimeter heaters will continue to provide heating, despite the AHUs' in the unoccupied mode. The effects of internal heat gains through occupancy, solar heat gains, and natural day-night temperature cycles should also be considered. These factors may encourage more similar afterhours and workhours energy use rates in the heating season and further polarize afterhours and workhours cooling energy use rates.

In the AHU anomaly detection reports, 23 AHUs from four organizations were analysed, of which 30% exhibited at least one hard fault, and 78% exhibited at least one soft fault. This can be interpreted as widespread suboptimal management of equipment, stemming from a lack of knowledge of best practices from operations staff or an oversight in the controller logic. A SAT reset strategy was not evident in any AHUs for organization B and D, and excessive AHU operation was flagged in all AHUs for organization D; this reinforces the similar workhours and afterhours energy use rates identified in the baseline

energy performance function, symptomatic of frequent night-cycling. The most prevalent of the soft faults was inappropriate economizer with cooling state settings.

In the zone anomaly reports, 219 thermal zones from four different organizations were analysed. Of which 15% of zones in the heating season and 10% of zones in the cooling season were outside acceptable conditions for indoor air temperature and airflow control error. In the heating season, more zones deviated from acceptable conditions, and variations in air temperature among zones were greater. Since the function resolves average zone conditions, it does not identify zones that exhibited instantaneous anomalies, but that are consistently outside acceptable conditions for the analysis period.

In the occupancy reports, earliest and latest arrival time, latest departure time, and duration of the longest break were calculated for organizations A and B. The times for both organizations could be used to inform occupant-based building- and zone-level schedules to reduce excess airflow where there is vacancy.

3.6 Closing remarks

The backend and frontend of a novel building energy management toolkit was presented which outlined the server-based handling of user data and web-based user interface of the toolkit. The toolkit was demonstrated using various building data from four separate organizations, which produced a total of 13 reports containing each function's generated visuals and KPIs. These reports are intended to help building energy professional address operational deficiencies and identify opportunities to further save energy. As such, interviews with members of the profession regarding their interpretation of the reports are planned and the feedback will be used to further improve the contents and presentation of the reports and website.

Key limitations of the toolkit were exposed by expanding the used dataset. Limitations on data quantity, which can be imposed by a building's available sensing technologies or data storage capacity, restricted use of the toolkit to a subset of the available functions. Furthermore, certain HVAC or metering configurations, like buildings which primarily rely on electricity-based heating or share a metering infrastructure with other buildings, may produce invalid results from certain functions, such as the baseline energy performance function. Future work may investigate the effects of further expanding datasets to buildings from different climate zones, better estimating the performance and financial impacts of addressing identified deficiencies, and taking steps to address or advise users of the toolkit's limitations.

Chapter 4

This chapter is under review as:

Markus, Andre A.; Hobson, Brodie W.; Gunay, H. Burak; and Bucking, Scott. Does a knowledge gap contribute to the performance gap? Interviews with building operators to identify how data-driven insights are interpreted. *Energy and Buildings*.

4.0 Interviews with building operators & facility managers

4.1 Introduction

The efficient operation of heating, cooling, and air conditioning (HVAC) systems in large institutional and commercial buildings often necessitates complex control procedures, regular maintenance of equipment, and energy use management, all of which are commonly performed by highly qualified operations personnel. Suboptimal operations, which can result in up to a 30% excess in actual building-level energy use [7], is considered a significant contributor in the wide-spread discrepancy between predicted and actual building energy use, also known as the ‘performance gap’. To ensure performance, optimal controls and scheduling strategies are prescribed and detailed as in ASHRAE Guideline 36-2018 [4], and ASHRAE Guideline 14-2014 [90] outlines reliable methods for assessing building-level energy performance. Furthermore, recent advancements in the data-driven building operations and maintenance analytics research field have fostered various novel methodologies to derive energy-saving insights using process-history data such as energy meter or HVAC controls network data [10], [40]. These data-driven approaches output various visualizations and KPIs, some inheriting ASHRAE guidelines’ figures and techniques, intended to inform optimized controls settings, identify and/or diagnose mechanical faults and anomalies, and improve building energy performance overall.

Though data-driven approaches serve to mitigate excess energy use through various methods, it remains unclear how well their outputs translate to energy-saving insights that can be utilized in practice. Building operations professionals (i.e., building operators and facility managers) are at the forefront of achieving building energy efficiency and use existing tools such as BAS and energy management systems (EMS) to inform their

decision-making processes [12]. However, data-driven approaches provide long-term insights derived through bulk data analysis which may augment their decision-making processes, potentially aiding these personnel in carrying out their duties. Accordingly, effective interpretation of the presented visualizations and KPIs is necessary to spur on energy-saving measures, though little is understood and documented on how such members of the building operations profession interpret these outputs.

To improve industry accessibility and the actionable potential of data-driven approaches to extracting energy-saving insights, an understanding of how building operators and facility managers interpret various generated outputs of established data-driven approaches in the literature is required.

4.2 Background and previous work

Despite scarce feedback from the building operations profession, data-driven approaches to extracting energy-saving insights continue to increase in quantity and complexity, further popularized by the growing accessibility and variation of archived building data and ML methods. These approaches aim to address the performance gap by identifying deficiencies in energy consumption, anomalies in physical components and operating states, and opportunities to non-intrusively improve energy efficiency, which may aid members of the profession in identifying and prioritizing operational changes.

4.2.1 Data-driven operation analytics

Though manifested in various forms and commonly utilizing ML and optimization techniques, data-driven approaches typically use archived data from numerous sensors and meters in the BAS to derive energy-saving insights. One prevalent approach is AFDD which serves to identify and diagnose hard and soft faults at specific hierarchies of a

building's HVAC system. This approach can further be categorized into white-box, grey-box and black-box model-based approaches. Shi and O'Brien [38] and Kim and Katipamula [39] reviewed several AFDD approaches, reporting a substantial increase in publications for black-box model-based approaches in the previous decade and utilizing, among other ML tools, ANN and regression analysis. These methods can augment a building operator's diagnosis of potential faults in equipment and provide additional information to ensure healthy operations. However, those in management positions who mostly deal with high-level energy management rather than day-to-day operations may not find such granular methods and insights particularly useful for their purposes.

Another common form of data-driven approaches serves to provide whole building-level energy consumption benchmarking and prediction. Li *et al.* [109] reviewed several methods for energy benchmarking incorporating one or several ML algorithms, noting its importance for monitoring and anomalous energy use detection. Amasyali and El-Gohary [104] further highlighted the importance of energy benchmarking approaches as an indicator for a suitable method for energy consumption prediction. For building operations professionals in management positions, particularly facility and energy managers, the insights generated by these methods can inform baseline energy use, help identify abnormal building-level energy consumption, and can serve to inform direction of resources and efforts.

OCC have provided methods for optimizing building controls for occupancy and occupant preferences. Though these methods have traditionally used a single data source to determine occupancy [64], recent studies have explored sensor-fusion [103]. These approaches can provide a more accurate and reliable method of occupant-count and

detection, further serving as a basis for optimizing control of the indoor environment such as adjusting equipment schedules or implementing demand-controlled ventilation (DCV) strategies.

4.2.2 Interviews with building operators and managers

Building operators are responsible for monitoring and optimizing energy use, maintaining equipment, and addressing issues that lead to occupant discomfort. There are various tools which collect and store building data that are used to inform their decision-making process. A BAS measures and collects multiple data points (e.g., temperature sensors, valve and damper positions, etc.), which can be used to monitor operating states and equipment and diagnose mechanical faults. Energy meters can be used to establish baseline figures and identify anomalous energy use, and CMMS provide a means of logging occupant complaints which can be used to diagnose anomalous conditions at the zone-level either manually or automatically using text analytics [73], [74].

Few studies have documented and analyzed the perspectives and challenges faced by building operations professionals pertaining the use and decision-making process associated with existing building operations tools. These studies collected and examined responses from interviews or surveys with multiple building energy professionals of varying roles, experience, and hierarchal standing. Goulden and Spence [110] conducted interviews with over 30 building energy and operations professionals of various hierarchal standing from four different organizations to examine the organizational impact in making energy-wise decisions, noting the expectations of facility managers to minimize energy and maintain occupant comfort. These expectations were not exclusive to a particular role in the management team, but have varying degrees of control and priority within the team's

hierarchy, thus warranting varying access to tools to administer each's responsibilities. Abuimara *et al.* [12] proposed goals to improve current building energy management practices by conducting interviews with building facility managers and operators regarding the use and challenges of building energy management tools. The study cited data inaccessibility, a lack of essential knowledge, and a lack of visuals which could effectively deliver insight consolidated from data to a human as current challenges in the practice. These interviews were conducted with 30 building energy and operations professionals in the Ottawa and Boston regions of varying titles and experience who were each asked the same 30 questions from a questionnaire. These studies sought to identify and categorize the challenges faced by building energy professionals and thus focused on the existing management infrastructure and cumulative experience from which professionals obtain decision-making insights.

Some studies have also identified the obstacles and perspectives of building energy and operations professionals regarding new or unfamiliar building energy management approaches, including novel protocols and management and information systems. Rock *et al.* [111] conducted separate, highly selective interviews with ten building facility managers between 7 to 25 years of experience regarding the challenges faced by undertaking more environmental-friendly operating practices in their buildings, finding a lack of financial incentive, lack of necessary knowledge, and undue deviation from stakeholder and tenant agreements as focal challenges. They sought to identify challenges presented by the adoption of novel practices in the profession by selectively interviewing managers who had experience in implementing these approaches. Similarly, Harris *et al.* [112] identified enabling and deterring factors for monitoring-based commissions using

energy management information systems, noting data preparation and quality and analyses reporting as the most common obstacles among users. It was also evident, by interviewing and surveying multiple organizations, that some of the potential factors can serve as a deterrent in one organization but be enabling in another if proactive measures are taken such as proper training and awareness to mitigate staff members' hesitancy to adopt new technology and methodologies.

Interviews pertaining to novel approaches should seek to identify the challenges presented by the factors which contribute to the approach's novelty in the context of existing protocols. Rather than considering purely hypothetical scenarios, participants should ideally have exposure to such novel approaches as well as experience with existing approaches, preferably in the context of one's own building (i.e., using data from one of multiple buildings in the participant's care). Accordingly, population sampling warrants a highly selective process based on the participant's exposure to novel approaches; Robinson *et al.* [113] recommends a sufficiently small sample size of no more than 16 participants such that the significance of individual responses is not lost within a large sampling size.

To the authors' knowledge, no previous research has focused on building operators or facility managers' interpretation of data-driven energy-saving insights to supplement their existing protocols. This is greatly in part due to the seldom adoption of such approaches in the profession, specifically towards building operators who are mainly tasked with day-to-day operations rather than long-term energy management. Understanding their challenges and decision-making process as a result of these derived insights is critical to bridging the gap between building operations data analytics research and practice, and taking a direct step towards improving building performance at the most fundamental scale.

4.3 Motivation and objectives

Despite recent advancements in data-driven approaches to operations and maintenance analytics, and their proven potential to improve building performance, these approaches have yet to reach mainstream utilization. One reason for this may be reluctance from building operations personnel stemming from a lack of transparency and the unfamiliarity associated with such novel approaches. For example, the outputs (i.e., visuals, KPIs, explanations) from such approaches may seem intuitive from a research perspective, but may be otherwise ineffective in a practical sense in informing an operator's existing duties unless additional information and context is provided; in this case, how much supplemental information can be provided to sufficiently inform the operator of the significance and practical uses of the outputs without overwhelming them? Therefore, gaining an understanding of how building operators and facility managers, members of the profession who are regularly involved in short- and long-term building operations and energy management, interpret data-driven approaches is imperative in establishing a synergistic relationship between building operations in academia and practice, improving building performance at the micro scale, and further developing practical data-driven approaches. To achieve this, members of the profession should be exposed to various outputs of data-driven approaches and inquired to derive their own conclusions and insights from data-driven visualizations and KPIs in context to what initiatives, either through controls adjustments or administratively, they would undertake in response to any energy use deficiencies (i.e., hard and soft AHU-level faults, excess ventilation, suboptimal supply air temperature settings, etc.) presented in the outputs; their responses would be qualitatively analyzed to determine how building operations personnel interpret and utilize data-driven

insights in practice and identify any barriers which may inhibit the intended use cases of these outputs.

4.4 Methodology

To assess the interpretations of building operators and facility managers of energy-saving insights derived from data-driven approaches, a unique methodology was adopted whereby the outputs presented to each group of participants (i.e., organization) were generated separately using data from their own buildings; these buildings were all located within ASHRAE climate zone 6A (Ottawa, Canada) and are exposed to a similar climate [114]. In other words, the generated fault alarms, visuals, and KPIs were bespoke to each organization, however, the format in which these outputs were presented (i.e., type of visual, type of KPIs, explanation text, order of presentation) was consistent throughout. Though this was done in part to encourage engagement in the interviews, it was decided that general interview-based methodologies which comprise simply of a list of non-specific questions cannot truly reveal the operator's overall ability to interpret and use data-driven insights in practice considering the participant's uniquely focused involvement and expertise of their own building. A qualitative approach was taken to analyze the participants' responses; however, the relevant anecdotal quotes were also taken to supplement key observations. The proceeding subsections are structured as follows: first, the questionnaire development is explored; then, the multi-source approach to extracting insights is detailed; finally, the participant recruitment and interview process is described.

4.4.1 Questionnaire development

The interviews consisted of 13 open-ended questions presented in Table 4.1; these questions were formulated to assess the participants' comprehension of the generated

outputs, including but not limited to identifying hard and soft faults, excess or inefficient energy consumption, and opportunities to improve energy performance through a supply air temperature (SAT) reset or optimized mode of operation settings. The questionnaire refrained from using questions such as, “how useful was the provided schematic,” or, “did you understand the significance of these KPIs,” since these questions are subjective and may imply differing connotations of understanding from individual to individual. Thus, questions primarily focused on what changes or remediating measures to building operations that participants would make, or rather not make, after reading the reports and why such initiatives would or would not be made. The questionnaire was reviewed and approved by the institution’s ethics review board.

Table 4.1: List of interview questions

Question No.	Question
1	What is your official organization title?
2	What is your experience/credentials that are most relevant to your position?
3	What problems in the building, if any, were you aware of prior to reading the reports?
4	If you were aware of problems prior to reading the report, did they show up in the reports? Specify.
5	What additional problems were you made aware of after reading the reports?
6	After reading the reports, what changes will you make to your building? Why?
7	If you do not plan on making any changes, why?
8	What known problems prior to reading the reports, if any, were missed in the reports?
9	How would you benefit from sharing these reports with your boss/colleagues/juniors?
10	What changes to the reports would help you make the right changes to the building?
11	Which report did you find most/least useful?
12	Based on the most useful report, why did you find it the most useful?
13	Based on the least useful report, why did you find it the least useful?

4.4.2 A multi-source approach to extracting energy-saving insights

To facilitate extracting insights from the various established data-driven approaches, a web-based software application employing established inverse modeling, FDD, and occupancy and occupant complaint analytics from the literature, guidelines, and standards was developed. The software application, which was based on an open-source library, contains seven discrete data-driven approaches (i.e., functions) which input a single or combination of data such as energy meter and HVAC controls network data, and automatically output a report-style document containing fault alarms, visuals, and KPIs; this is done separately for each function. Text explanation is also automatically outputted in the reports to inform the recipient on how to interpret the visuals and KPIs. Figure 4.1 presents the workflow for generating the reports through the web-based application.

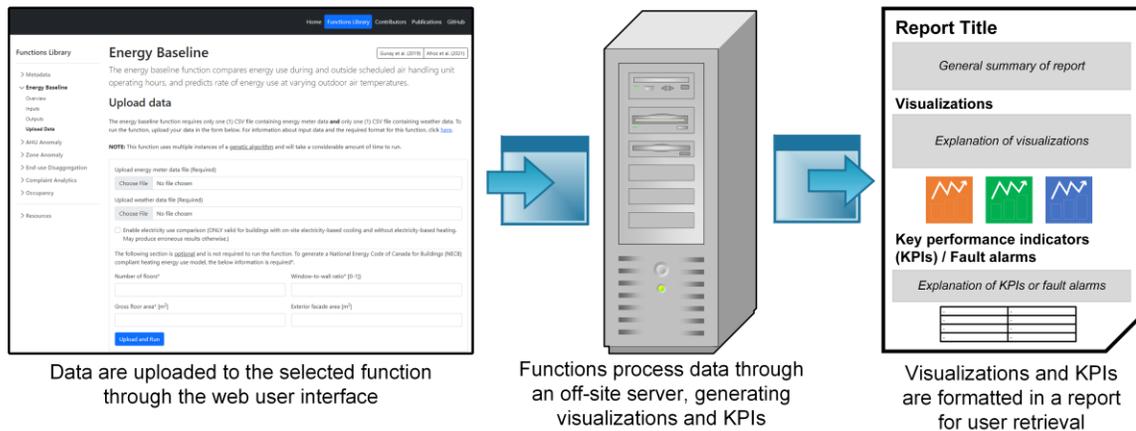


Figure 4.1: Workflow for generating reports using the web-based application.

The functions' outputs serve to inform the interviewee of operational deficiencies stemming from improper control strategies, hard and soft AHU- and zone-level faults, or suboptimal mode of operation settings. One of these functions, the baseline energy function, compares bulk energy use during (i.e., workhours) and outside (i.e., afterhours) scheduled AHU operating hours and is intended to help recipients assess a weekly AHU schedule's ability to reduce energy use outside the building's operating times. This function

utilizes inverse modeling to develop separate modified ASHRAE Guideline 14 [90] three-parameter change point models for workhours and afterhours energy use; Figure 4.2 is an example of this type of outputted visualization for heating energy use which was distributed to one organization's facilities management team. This figure was intended as a benchmarking visualization to inform the interviewee of the presence and actual effectiveness of a weekly AHU schedule to reduce energy use during a building's unoccupied hours. Text explanation informs the interviewee that, by comparing the slope of the afterhours and workhours slope, the interviewee can evaluate how well a weekly AHU schedule, if one exists, reduces energy use during afterhours; similar degrees of slope are symptomatic of similar rate of energy use during afterhours as workhours, which is considered suboptimal. The "ScheduleEffectiveness" (SE) KPI quantifies how dissimilar the slopes are relative to each other (0% if the models' slopes are identical and increases positively as the afterhours slope decreases relative to the workhours slope), and the "AfterhoursEnergyUseRatio" (AEUR) KPI presents the ratio of energy used during afterhours over the total energy use in the analysis period.

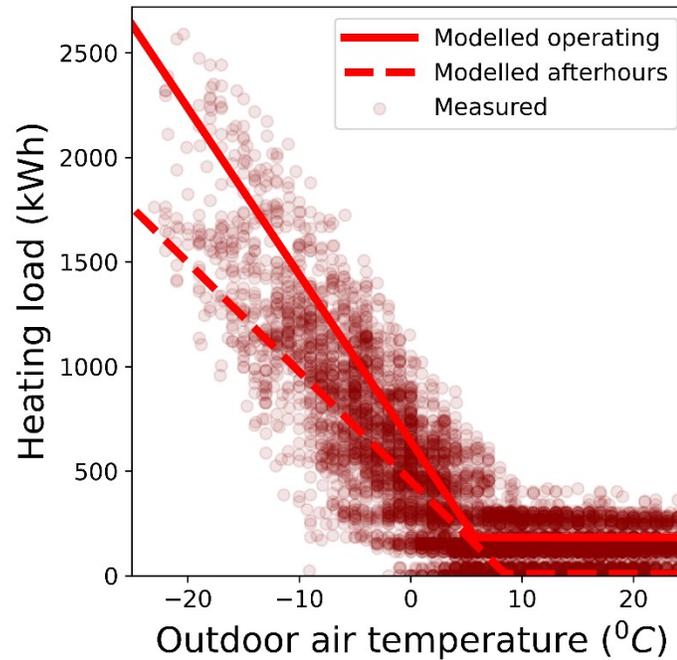


Figure 4.2: Example outputted visualization of modified AHSRAE Guideline 14 three-parameter change point models for workhours and afterhours heating energy use, as distributed to one organization's facilities management team. The "ScheduleEffectiveness" and "AfterhoursEnergyUseRatio" KPIs are 34% and 40%, respectively.

The AHU anomaly detection function serves to identify common hard and soft faults with the goal of informing the interviewee of underlying AHU-level hard faults (i.e., stuck outdoor air damper, stuck heating/cooling coil valve) or suboptimal controls settings (i.e., lack of weekly AHU schedule, lack of SAT reset strategy, inappropriate economizer mode settings). As well as presenting identified faults, this function also generates a diagram depicting the SAT as a function of outdoor air temperature alongside the coolest and warmest zone temperatures and an ideal range for an SAT reset strategy. Expository text informs interviewee that this visualization should be used to observe the presence and modulation of an SAT reset strategy, if one exists, and to inform optimal SAT reset controls settings; the interviewee is informed that an ideal SAT reset strategy should fall within the green highlighted area [97] and is provided with likely causes of a suboptimal strategy where the SAT falls outside the highlighted area. A split-range controller diagram per

ASHRAE Guideline 36 [4] compliments the previous diagram by illustrating outdoor air damper position, heating and cooling coil valve position, and fraction of active perimeter heating devices; this allows the recipient to observe any anomalous patterns of operation; Figure 4.3 is an example of this type of outputted visualization which was distributed to one organization’s facilities management team. Expository text encourages the interviewee to observe key characteristics of healthy AHU operations (i.e., heating coil valve should only be open in heating state, outdoor air damper should be near completely open in economizer with cooling state) and investigate AHUs which do not exhibit healthy operations. The text also provides diagnostic steps to the potential fault alarms.

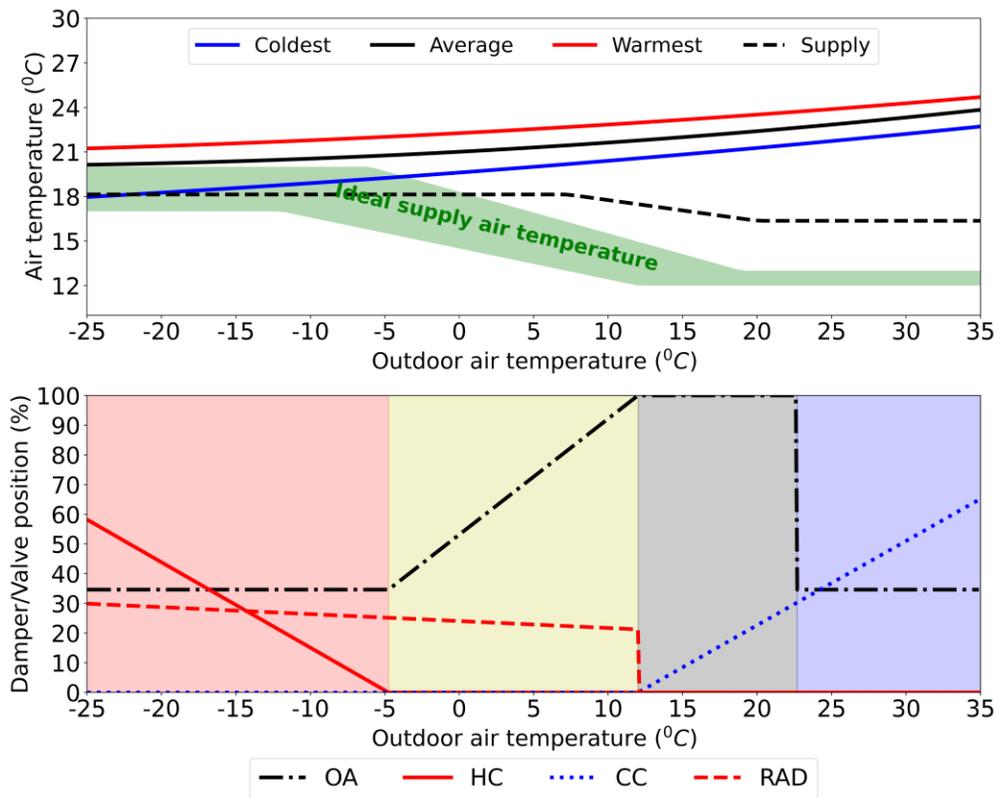


Figure 4.3: Example outputted visualization of a SAT diagram alongside the warmest and coldest zone temperatures and ideal range for supply air temperature reset strategy. A split-range controller diagram appears below the previous diagram, allowing the recipient to observe AHU outdoor air damper position (OA), heating coil valve (HC), cooling coil valve position (CC), and perimeter heater modulations (RAD) and identify any anomalous operations. These visualizations were distributed to one organization’s facilities management team.

The zone anomaly function generates a zone cluster diagram which plots clusters of thermal zones based on their average indoor air temperature and airflow control error separately for the heating (December through February) and cooling season (June through August). The purpose of this visualization is to inform the interviewee of potential VAV-level hard faults such as a stuck reheat coil valve or damper which may result in abnormal indoor air temperature or airflow. An example of this type of outputted visualization which was distributed to one organization's facilities management team is seen in Figure 4.4; the zones attributed to each cluster is also presented to the interviewee by file name. Text explanation informs the interviewee that this visualization should be used to identify and diagnose VAV terminals which serve thermal zones with anomalous indoor air temperatures and airflow. The interviewee is encouraged to ensure that the VAV units' terminal damper and airflow sensors are operating as intended for zones exhibiting high or low flow, and that perimeter heating devices or reheat coils are operating as intended for zones exhibiting abnormally hot or cold indoor air temperatures.

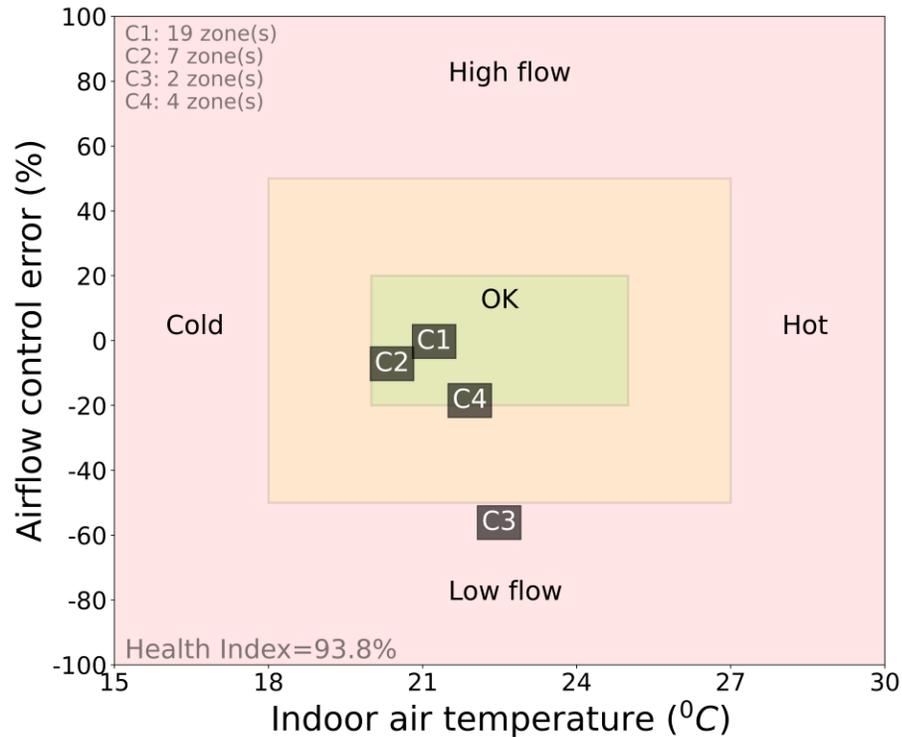


Figure 4.4: Example outputted visualization of zone cluster diagram, as distributed to one organization’s facilities management team.

The occupancy function calculates KPIs intended to help optimize the mode of operation settings, particularly by adjusting whole building- and zone-level occupied hours. These KPIs are the typical earliest arrival time, latest arrival time, latest departure time, and longest duration of vacancy which are derived from motion detection data; an example of these computed times which was distributed to one organization’s facilities management team is seen in Table 4.2. Text explanation informs the interviewee that the earliest arrival time and latest departure time should be used to adjust building-level schedule whereby a building’s “occupied” mode should start at the earliest arrival time and end at the latest departure time. Furthermore, the interviewee is also encouraged to use the latest arrival time and longest break duration for zone-level adjustments where if a particular zone is not yet occupied by the latest arrival time, or if there is a period of vacancy in excess of the longest break duration, its corresponding VAV terminal should be set to the “unoccupied”

mode. For this function, an exemplar pseudocode for real-time zone-level adjustment of “occupied” hours is also outputted as in Figure 4.5.

Table 4.2: Example outputted earliest and latest arrival time, latest departure time, and longest break duration, as distributed to one organization’s facilities management team.

Earliest arrival time	Latest arrival time	Latest departure time	Longest break duration (hours)
9:00	10:00	17:00	4

```

CURRENT_TIME = get_current_time()
# If schedule is set to occupied, set VAV_101 to 'occupied' mode.
if SYSTEM_SCHEDULE == 'OCCUPIED':
    VAV_101_CTRL_MODE = 'OCCUPIED'

    # ... and if motion is detected, indicate zone is STILL occupied and track last time occupied ...
    if motion_detection == True:
        zoneNotYetOccupied = False
        lastTimeOccupied = CURRENT_TIME

    # ... but if past 10 am (latest arrival time) and zone not yet occupied ...
    # ... set VAV_101 to 'unoccupied' mode.
    if hours > 10 and zoneNotYetOccupied == True:
        VAV_101_CTRL_MODE = 'UNOCCUPIED'

    # ... but if 4 hours (longest break duration) has elapsed since last time occupied and zone not yet occupied ...
    # ... set VAV_101 to 'unoccupied' mode.
    if lastTimeOccupied + 4 < CURRENT_TIME and zoneNotYetOccupied == True:
        VAV_101_CTRL_MODE = 'UNOCCUPIED'

# If schedule is set to unoccupied, set VAV_101 to 'unoccupied' state.
else:
    VAV_101_CTRL_MODE = 'UNOCCUPIED'

```

Figure 4.5: Outputted exemplar pseudocode for real-time adjusting zone-level VAV occupied hours utilizing latest arrival time (10 am) and longest break duration (4 hours), as distributed to all organizations’ facilities management team. If the example zone corresponding to VAV_101 is not yet occupied by the latest arrival time, or if no motion has been detected for longer than the longest break duration, VAV_101 should be set to the “unoccupied” mode.

The complaint analytics identifies hot and cold occupant complaints from CMMS work order logs in context to the prevailing operating and weather conditions which trigger those complaints; this function is intended to help operators adjust temperature setpoints for ideal occupant comfort.

4.4.3 Data collection and processing

Archived building data were collected from three buildings, each of a separate facilities management organization. These data consisted of energy meter, AHU- and zone-level

HVAC controls network, occupancy, and CMMS data. Collected energy meter data contained hourly measured energy consumption for heating, cooling, and in some building, electricity use. AHU-level HVAC controls network data contained hourly data for SAT, return air temperature, outdoor air temperature, outdoor air damper position, heating and cooling coil valve positions, and supply fan state. Zone-level HVAC controls network data contained hourly supply air flow rate, supply air flow rate setpoint, indoor air temperature, and state of perimeter heaters. Occupancy data consisted of event-based motion detection data, and CMMS data contained operator work order logs. Once obtained, the data were arranged into an acceptable format for input, uploaded through the web user interface, and processed off-site through a virtual private server. Depending on the availability of data per building, the data were inputted into up to five functions in the library; not all functions could be used due to minimum data requirements. These functions were the energy baseline, AHU anomaly detection, zone anomaly detection, occupancy, and complaint analytics function. Figure 4.6 presents the distribution of available data from each organization’s building to the invoked functions.

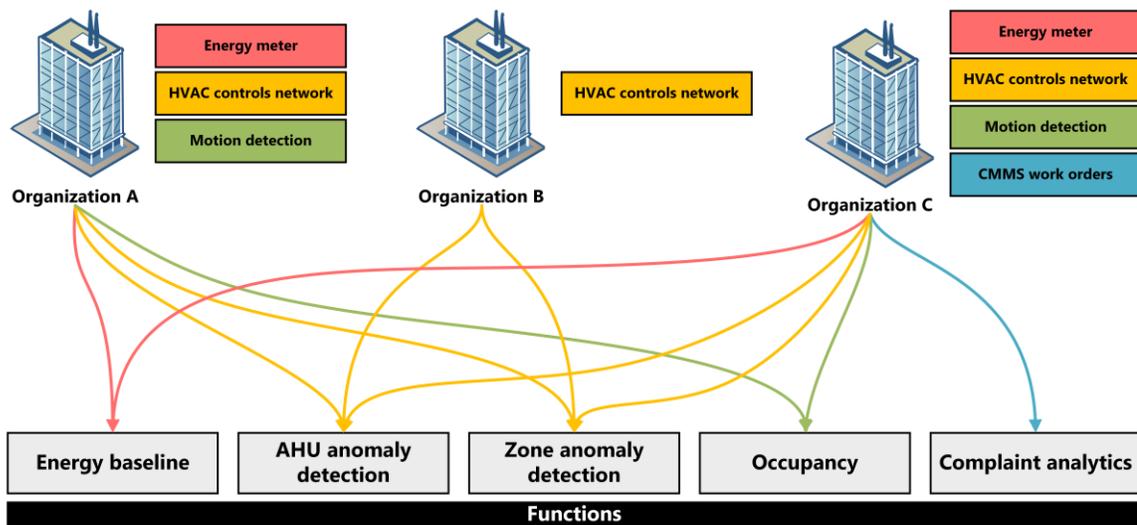


Figure 4.6: Distribution of available building data and organization to the functions

4.4.4 Report generation and distribution

Each function automatically outputted a report-style document containing the function’s generated visualizations and KPIs once the function’s analysis has concluded. In addition to the generated outputs, the document contained expository text to inform recipients of the significance of the KPIs and visuals, how to interpret them, possible symptoms of suboptimal operations, and suggestions to remediate deficiencies through improved controls strategies. These reports were distributed to the corresponding organization’s facilities management team including but not limited to facility managers, building operators, and operating managers prior to the interviews; this was to allow the participants to read and comprehend the contents of the report in preparation for the interview. Reports were only distributed to participants who had consented to the interviews. Table 4.3 presents the type of report and deriving function, and the corresponding organization to which the reports were distributed.

Table 4.3: Breakdown of reports distributed to organization's participating facilities management team

Org.	Type of report distributed (<i>function</i>)				
	Baseline energy (<i>baselineEnergy</i>)	AHU anomaly (<i>ahuAnomaly</i>)	Zone anomaly (<i>zoneAnomaly</i>)	Occupancy (<i>occupancy</i>)	Complaint analytics (<i>complaintAnalytics</i>)
A	✓	✓	✓	✓	
B		✓	✓		
C	✓	✓	✓	✓	✓

4.4.5 Participant recruitment and interview administration

Initial contact with the participating organizations was administered through the authors’ professional network. Once in contact with the appropriate facilities management team, a

recruitment email was distributed which solicited a video-conference interview pertaining to the generated reports. A total of 11 participants from the three different organizations agreed and consented to an interview; those who had consented were provided with the reports containing fault alarms, visuals, and KPIs generated using data from the participant's individual building portfolio to read in preparation for the interview. While the number of participants may seem sparse in the context of a traditional survey, it is important to consider that the solicitation process was highly selective, and was based on the availability of the provided data, the suitability of the data with the functions (i.e., sufficient quantity and quality of data to successfully generate outputs), and the participant's profession in building operations. The objective of this study is to qualitatively evaluate the recipient's interpretation of the generated outputs in context of the participant's managed building. Thus, "life history homogeneity" [113] subject to the participants' unique expertise in their own managed building, and in the building operations profession, was adopted to emphasize the contents of individual responses; the sample size thus adheres to the guidelines of Interpretative Phenomenological Analysis, a recognized research method for the use of interviews as data source. Participants' experience in their respective field ranged from 2 to 28 years and while the majority of participants were building operators, property/facility managers were also in attendance. After the interviews, participant responses were transcribed. The breakdown of the participants is presented in Table 4.4. In a few cases, participants did not consent to providing their experience.

Table 4.4: Breakdown of participants for semi-structured, conference-style interviews

Participant No.	Affiliated organization	Organization title	Years of experience
1	A	Property management and sustainability intern	2
2	A	Controls and Commissions Specialist	10
3	A	Director of operation and project	14
4	B	General manager	3
5	B	General manager	-
6	B	Building operator	10
7	B	Building operator	28
8	C	Manager of building operation	-
9	C	Building operations technician	-
10	C	Building operations technician	12
11	C	Building operations technician	-

4.5 Interview results

4.5.1 Interpretation of AHU faults and anomalies

Among interviewees, the most prevalent discussions were related to AHU-level faults and anomalies and the split-range controller and SAT diagram visualizations as per Figure 4.3. Six of the eleven interviewees, five building operators or operating managers and one facilities manager, noted that these fault alarms and visualizations were particularly useful for identifying which AHUs exhibited hard (i.e., stuck dampers and valves) and soft (i.e., inappropriate controls settings) faults, noting the intuitive table-style presentation of faults to “pick out which AHUs need fixing”. The five operators stated that, though they were generally already aware of various hard and soft AHU-level issues prior to reading the

reports and noted that some of these issues were realized in the reports, they were made particularly aware of new hard faults such as a stuck heating coil valve or stuck outdoor air damper and clearly specified that these components would be investigated and repaired if necessary. However, only two of the operators proposed solutions to identified soft faults such as programming an automatic reset strategy or revising economizer state settings. Through the split range controller and SAT diagram, the five operators/operating managers and one facilities manager noted that the actual SAT differed from the ideal region highlighted in green, however only these two building operators stated this issue would be addressed immediately by programming an automatic SAT reset strategy, while the rest believed “further investigation” was warranted before any remediating steps can be undertaken and did not specifically state what measures they would take. The facility manager noted that while they were aware of various issues at the AHU-level, they were unaware of what specific component-level issues were present and “are not involved in the details.”

4.5.2 Interpretation of zone-level faults and anomalies

Participants were informed of zones exhibiting anomalous indoor air temperatures and deficient airflow primarily through the zone cluster diagrams as per Figure 4.4. Four of the eleven interviewees, all of which building operators, expressed favor for this type of visualization for clearly identifying which zones fall outside the range of acceptable conditions, noting that the diagrams were easy to understand. One of these operators continued by highlighting the function’s ability to analyze zone-level HVAC data over an extended sample period which would “omit short term anomalies that trigger persistent flags”. Another operator was made aware that three of the 98 thermal zones in their building

had an actual volumetric airflow of 76% less than the setpoint value in the cooling season, but did not specify a solution to this issue. In fact, all interviewed operators noted that “further investigation” was required to definitively determine a correct course of action; this was case for all participants, despite expository text informing that such anomalous zones may be caused by faulty VAV terminal dampers or reheat coil valves. Furthermore, one operator preferred that the worst performing zones be identified and that trend data on these zones also be presented to aid in diagnostics. To facility managers, the responses mimicked that of AHU-level faults and anomalies; as they are not involved in the technical aspects of the building, they could not comment on any administrative or controls changes.

4.5.3 Interpretation of building-level energy use deficiencies

Building-level energy use deficiencies were seldom discussed in the interviews. Only two of the interviewees, both facility managers, were particularly interested in the ASHRAE Guideline 14 change point regression models as per Figure 4.2. One of these facility managers noted that the comparison of energy use during and outside scheduled operating hours would be useful as a “yearly summary report” of building performance. However, they did admit to having difficulty interpreting the SE KPI and, despite the text explanation, was unaware of what this represented. Another facilities manager also stated that this visualization was open to “broad” interpretations; they noted that by viewing the visual alone, they could not determine “what was optimal or not”, and that there was no clear indication of what was an ideal SE or AEUR score.

4.5.4 Interpretation of occupancy and occupant complaints-based KPIs

None of the interviewees found any practical implication of the presented earliest arrival, latest arrival, latest departure time, and longest break duration of building occupants. However, these metrics should have been used to optimize building- and zone-level schedules as explained in the expository text. During the interviews, two participants admitted that they did not understand the significance of these KPIs or the provided explanation of the KPIs. For the recipients of the KPIs related to occupant complaints, participants noted that very few complaints were identified and that, as a result of this, the consolidation of the complaints in context to operating and weather conditions provided no useful insights.

4.6 Discussion

The interpretations of building operators and facility managers to data-driven analytical tools were assessed through their intention to adopt novel or revise existing controls strategies and identify and remediate AHU- and zone-level hard faults in response to the outputs of these approaches, as well as their general reception of the outputs. Despite the majority identifying symptoms of hard faults in the reports, interviewees generally had difficulty specifying, or were hesitant in explaining, what initiatives they would undertake to remediate such issues, opting instead for “requires further investigation” without further clarification, specifically regarding issues concerning controls changes. There are probable explanations for this hesitancy and these reasons may not be mutually exclusive, however, the responses were largely varied and thus, a consensus to a universal barrier inhibiting effective interpretation and actionable energy-saving insights cannot be made.

A general unfamiliarity of data-driven analysis methodologies may draw question to the reliability of the outputs. These methodologies which incorporate highly specialized techniques such as genetic algorithms, inverse modeling, or clustering techniques to create visualizations and estimate fault detection parameters are comparatively esoteric next to more conventional hands-on and direct decision-making methods whereby faults and anomalies are typically first flagged by occupant complaints or through the manual screening of sensor data on an individual AHU basis. Understandably, to ascertain the outputs of such novel analytical tools without context is unconvincing, especially if interviewees are introduced to these methodologies for the first time. By extension, conventional tools used by building operators (i.e., BAS, EMS) are used to derive event-based insights, focusing on specific components over a relatively short period [115], [116], whereas data-driven tools identify long-term underlying deficiencies from bulk data which may omit short-term anomalies; this point was brought up by only one operator concerning the zone cluster diagram as in Figure 4.4. Operators may simply not be familiar with data-driven insights and therefore unsure how to utilize them. However, it was anticipated that the majority of the interviewees would be unfamiliar with the outputted KPIs and visualizations, as well as how to respond to the identified faults and inefficiencies; expository text was included to aid interpretation, although interviewees may simply have not read this text prior to the interviews and, subsequently, were unaware of how to interpret or act on the insights presented. The notion was raised several times that a building operator would be hesitant to set time out to absorb all of the information provided to understand the KPIs and visualizations, especially considering their sporadic and dynamic schedules. Reducing the amount of expository text may encourage interviewees to read the

text, but would risk the interviewee being insufficiently informed to appropriately interpret the outputs. By contrast, adding more may better inform interviewees of the practical applications of the outputs, but may overwhelm them with information. Intuitive visuals serve far more effectively as actionable insight than simple text, let alone paragraphs of explanations [117], [118]. The visuals should be able to convey aspects of optimal or suboptimal operations with little to no supplementary descriptions, though ideally, operators are already aware of the how these outputs are generated, what insights they can present, and how to effectively utilize them in practice.

Another reason may be a habitual nature not to make operational changes if not prompted with immediate or pressing circumstances otherwise [12]; this hesitancy may be compounded by the novelty of these approaches and lack of third-party validation of savings. This may explain why hard faults were more readily identified and prompted immediate remediation by some interviewees than soft faults concerning controls changes. The impact of hard faults, such as a faulty outdoor air damper or heating coil valve, may be easier to conceptualize and instill more immediate ramifications if not resolved than soft faults such as suboptimal scheduling of operating hours or inappropriate economizer state settings. It can be argued that the long-term, data-driven insights may seem intangible to building operators in the context of their day-to-day responsibilities. Furthermore, building operators can become preoccupied with short-term, event-based responsibilities such as managing occupant comfort and may not see the merit of long-term insights. If continual improvements to operations are not prescribed within their responsibilities, there would be little else to incentivise proactive measures to improve energy performance. In this case, a risk-adverse work environment may also contribute to a lack of interest in best practice,

that is optimal operating controls strategies. Recall that interviewees were provided with the reports several days prior to the interviews; their reluctance to read the reports may highlight their disinterest in optimizing operations. Best practice were the basis of the visualizations and KPIs of these data-driven tools. For example, the SAT diagrams should be used to observe and inform optimal SAT setpoint reset strategies; an optimal range was highlighted in green. In several AHUs under one facilities management team, the SAT did not modulate as expected of a proper reset strategy, instead remaining static throughout the cooling and heating season, which may indicate a conflict in the controls logic preventing an automatic reset strategy from properly actuating, or an outright absence of one. However, this issue prompted little concern among the interviewees, and one interviewee was initially unfamiliar with the concept of a reset strategy and its benefits. Accordingly, if the range for an ideal SAT reset strategy had not been highlighted, the recipients may have assumed that a fixed SAT throughout both seasons was fine whereas in reality, an outdoor temperature- or zone temperature-based SAT reset strategy can encourage economizer operation, minimize perimeter heating in the economizer state, and reduce fan energy use [4]. Similarly, the calculated earliest and latest arrival time, latest departure time, and longest vacancy period can be used to adjust scheduled occupied hours for the mode of operation at the zone- and building-level [101]. However, none of the interviewees suggested that such adjustments be considered after receiving these metrics; two interviewees admitted not knowing what value the KPIs offered in any practical sense. The reality is that the outputs of data-driven approaches are ineffective in producing actionable energy-saving insights if operators are not already knowledgeable of best operating practices, despite accompanying text explaining the significance of the visualizations and

KPIs. The two interviewees who were knowledgeable of such methods understood the visuals and their practical implications, either since they were familiar with them or supplement their current operating expertise. These interviewees, upon realizing that an automatic SAT reset had not been implemented in a few of their AHUs from the outputs, readily specified that one would be applied and automated immediately. They were more willing to implement controls changes in light of the outputs, despite a lack of any immediate ramifications otherwise, since the outputs compliment their existing knowledge. Still, some operators do shoulder a duty to improve overall energy efficiency [119] and doing so merits a combination of short-term interventions (e.g., fixing faulty components, adjusting temperature and airflow setpoints) and long-term interventions (e.g., implementing optimal schedules, implementing automated SAT reset strategies) to operations, though whether this obligation is contractual or moral may vary from individuals. The outputted reports contain the visualizations and KPIs should ideally stress the practical implications of the outputs in achieving long-term energy goals. Roussac and Huang [120] found that simply establishing “targets” can further encourage actionable insights among operators. The visuals and KPIs themselves should emphasize a realistic and tangible goal which operators would be encouraged to strive for such as highlighting an optimal range of operating controls parameters or a target KPI value for healthy operations. For example, the SAT diagram highlighted an ideal range for a SAT reset strategy; one facilities manager favored this type of visualizations since the highlighted area would help guide operators to establish ideal SAT setpoints. Similarly, four of the eleven interviewees expressed preference for the zone cluster diagram since its visual presentation succinctly identified which zones fell outside acceptable conditions. This

“target” method can be applied to the ASHRAE Guideline 14-2014 change point models for expected energy use based on the building’s physical characteristic. Furthermore, a metric can quantify the proportion of the observed SAT within the ideal SAT range, and similarly with energy use change-point models, where a realistic and attainable target value would be encouraged.

4.7 Closing remarks

To the authors’ knowledge, this is the first study of its kind to selectively inquire about the utility of multiple novel data-driven methodologies utilizing data from the participants’ individual building portfolio. Building operators and facility managers were provided with fault alarms, visualizations, and KPIs, generated using data collected from their own managed building and interviewed regarding what initiatives or remediating actions they would take as a result of any operating deficiencies exposed in the outputs. Though individual responses varied and a consensus to one industry-wide barrier to effective interpretation cannot be ascertained, these interviews revealed that, at least to the extent of the interviewed facilities management teams, an unfamiliarity with data-driven methodologies and tools, and the inherently risk-adverse nature of building operations leading to a disinterest in optimal controls strategies are partly responsible in deterring actionable insights from data-driven building operations analytical approaches.

Building operators and managers were understandably hesitant to rely on such unfamiliar tools to inform their current responsibilities. The nature of data-driven approaches is that long-term insights are generated through analysis of an extended sample period and these insights may seem distant to their existing decision-making methods to identify faults and deficiencies. However, improving energy efficiency should be a key role of building

operators and facility managers and any continual improvements to building energy performance constitutes short- and long-term interventions to controls and equipment maintenance.

At its core, data-driven insights serve to inform optimal control strategies. In most cases, a lack of interest of optimal control strategies for energy efficiency, stemming from a hesitation to invoke non-pressing controls changes, was evident either in the apparent misconception of a proper SAT reset strategy or how to optimize zone- or building-level operating schedules using the generated visualizations and KPIs. This knowledge gap ultimately prevents the outputs from supplementing and complimenting their own operating expertise; it might instead even cause confusion where existing operating procedures contradict seemingly novel and optimal approaches, leading to further hesitancy and distrust in data-driven methodologies.

The authors stress that the findings of this study are not intended to reveal the shortcomings or competency of building operators or facility managers, but to expose potential barriers in building operations practice which may deter widespread industry adoption of data-driven approaches. It should be noted that a number of participants did demonstrate a fair level of understanding of the data-driven approaches, the outputs, and their practical implications. Furthermore, the participant recruitment process was entirely voluntary and may inadvertently filter in participants who were more knowledgeable, innovative, or comfortable in interacting with such novel methods; volunteering interviewees may be better poised to answer the questions, leading to biases in the results. To these considerations, the authors advocate for the following points:

- **Ensure data literacy and coding proficiency** – Industry professionals should have prior exposure and knowledge on how to read code, make basic coding adjustments, and interpret data-driven visualizations and KPIs. From pneumatic to digital controls, and now more complex data-driven analytical techniques, the sophistication of a building operator’s duties has drastically evolved over the past several decades, and training programs should reflect this evolution. Operators and managers should be able to differentiate trends and observations of data-driven outputs symptomatic of optimal and suboptimal operations and be trained on how to respond to operating deficiencies stemming from suboptimal controls strategies and methods to apply KPIs in practice.
- **Encourage operations personnel to proactively make controls changes** – A risk-adverse work environment, leading to a hesitancy to make changes to existing operations controls, may contribute to a lack of interest in best practice sequence of operations. To this end, the responsibilities of operations personnel should include and emphasize optimizing building energy performance. This can be done by including and assessing performance metrics and incentivizing personnel to improve whole building- or component-level energy efficiency.
- **Ensure skills training on best practice sequences of operation** – Skills training should serve to better inform industry professionals of optimal controls strategies for improved energy efficiencies such as implementing a proper automatic SAT reset strategy and methods to reduce excess consumption through optimized mode of operation settings, and their subsequent effects to overall building performance. Training should include knowledge of relevant guidelines and standards such as

ASHRAE Guideline 14 and 36 – these may be endorsed by ASHRAE, International Facility Management Association, and Building Owners and Managers Association. In addition, ensure existing duties includes improving building energy performance and where possible, quantify and incentivize these improvements.

Future work should evaluate the participants' responses after implementing changes or initiatives in light of data-driven insights to see if their perspectives of data-driven approaches change as a result of real-world observed benefits. Furthermore, should the research methodology be imitated, interviewees should be incentivized to further encourage the interviewee's participation and critical decision-making capabilities to the presented outputs. Additionally, a greater understanding of how the risk of making changes for the sake of improving energy efficiency is placed among building operators or absorbed by a third-party energy service company would further add context to the participants' willingness to apply data-driven insights. Future work should also attempt to identify and understand the financial implications that factor for and against the integration of data-driven methods in industry.

5.0 Conclusions

5.1 Summary and contribution

This research explored the development, implementation, and industry reception of a novel multi-source, data-driven building energy management toolkit as a synthesis of established data-driven building operations analysis approaches in the literature. This toolkit was used to identify hard and soft AHU- and zone-level HVAC faults, identify energy use anomalies, augment sequences of operation settings (i.e., weekly AHU schedules). To this end, the resolution of such suboptimal operations may result in annual energy savings of up to 30% [1], [6], [7]. The summary of conclusions for each chapter is presented below.

5.1.1 A framework for a multi-source, data-driven building energy management toolkit

In this chapter, the structure of a multi-source, data-driven building energy management toolkit was proposed, incorporating several established data-driven building operations and maintenance analytical approaches. Seven discrete approaches were outlined, comprising the toolkit's open-source functions library, which encompassed FDD, inverse energy modeling, load disaggregation, and occupancy and occupant complaint analytics. The applications, inputs, outputs, and methodology of each function were described and its unique multi-faceted approach to identifying operational deficiencies was demonstrated on a seven-storey case study building. Energy meter, HVAC controls network, Wi-Fi device count, and work order logs from Jan 1, 2019 to Dec 31, 2019 were inputted into the toolkit's functions and its resultant visualizations and KPIs were analyzed. The key findings of this demonstration were as follows:

- Five faults pertaining to the AHUs' mode of operation and heating coil valves were identified and insights from the generated outputs (i.e., visualizations and KPIs) were used to pinpoint operational deficiencies stemming from inappropriate zone temperature overheating thresholds and perimeter heating devices.
- The inclusion of various data-driven approaches under one toolkit allowed for greater diversity of the derived insights and can provide greater flexibility in ensuing measures to improve overall performance.
- Underlying deficiencies can be captured through various interpretations of the data which can be consolidated to focus fault correction efforts to specific components and system levels. Redundant indicators between functions may suggest a mutual deficiency which would strengthen the confidence of the toolkit's detection capabilities.

5.1.2 FRAMeWORK: A web-based application to derive insights

This chapter explored the applicability of the toolkit's functionalities with four separate case study buildings, each located within ASHRAE climate zone 6A [114], and details the back-end handling and processing of user data and front-end user interaction of the toolkit's web-based application platform. The web-based platform was created to facilitate user data input of data and retrieval of the toolkit's outputs and further disseminate the capabilities of data-driven analytics and the toolkit's overlying objective to optimize building energy efficiency. Data from the four case study buildings were inputted into up to five functions in the toolkit using the web user interface, where select functions were omitted due to data limitations, and processed off-site through a virtual private server. The visualizations and

KPIs were extracted through the web user interface in the form of pre-formatted reports and analyzed. The key findings from the analysis were as follows:

- Heating energy use outside scheduled AHU operating hours (i.e., afterhours) were consistently more similar to heating energy use during scheduled operating hours (i.e., workhours), primarily caused by the always-on nature of hydronic perimeter heaters, and partly due to internal heat gains through occupants, solar heat gains, natural day-night temperature cycles, and the relative length of workhours compared to afterhours.
- Of the 23 AHUs analysed from four different buildings, 30% exhibited at least one hard fault and 78% exhibited at least one soft fault. This can be interpreted as widespread suboptimal controls management, stemming from a lack of knowledge or interest of best practice sequence of operation from operations personnel or an oversight in the controller logic.
- Of the 219 thermal zones analyzed from four different buildings, 15% in the heating season and 10% in the cooling season were outside acceptable conditions for indoor air temperature and airflow control error.
- The toolkit's functionalities were restricted by data quantity, imposed by limited availability of sensing technologies in a building or data storage capacity.
- Some functions (i.e., Baseline energy performance) can produce invalid results if data originates from certain HVAC or metering configurations, such as buildings which primarily rely on electricity-based heating or share a metering infrastructure with other buildings.

5.1.3 Interviews with building operators and facility managers

In this chapter, interviews with building operations personnel (i.e., building operators and facility managers) were conducted to understand how the intended recipients of the toolkit's outputs interpret data-driven visualizations and KPIs and translate them to energy-saving measures, and to identify barriers to effective interpretation. Three facilities management team were solicited, of which 11 participants consented for an interview. Data from the interviewees' individual building portfolio were obtained and analyzed using up to five of the toolkit's functions. Reports containing various data-driven visualizations and KPIs were distributed to the interviewees prior to the interview to read in preparation. Their interpretations of the toolkit's outputs were assessed by their intention to adopt novel or revise existing controls strategies and identify and remediate AHU- and zone-level hard faults in response to the outputs of these approaches, as well as their general reception of the outputs. The key findings from this analysis were as follows:

- The most prevalent discussions surround AHU-level faults and anomalies and the SAT and split-range controller diagram visualizations, and the least prevalent concerned occupancy and occupant-based complaint analytics.
- Most operators understood and were able to identify hard faults using the reports, but had difficulty interpreting soft faults such as lack of an SAT reset strategy or inappropriate economizer state settings; they were hesitant or had difficulty specifying what actions they would undertake to remediate these issues.
- A general unfamiliarity with data-driven approaches and subsequently how to interpret and utilize data-driven insights, and a disinterest in optimizing operations stemming from a risk-adverse work environment and a habitual nature not to make

operational changes, were identified as possible barrier that inhibit the toolkit's outputs from effectively inciting energy-saving measures to operating personnel.

5.2 Research contributions

5.2.1 A framework for a multi-source, data-driven building energy management toolkit

This chapter contributes to the diverse body of literature on data-driven building energy management strategies by demonstrating novel applications to identify and diagnose operating deficiencies and suboptimalities through the consolidation of various energy-saving insights from traditionally disparate data-driven approaches in the literature. Specifically, the presented inclusion of several established methodologies, simultaneous utilization of these methodologies on a single building, and unification of insights derived from the outputted KPIs and visualizations, which were facilitated by a preliminary multi-source building energy management software toolkit, provides a framework upon which researchers and practitioners can derive similarly purposed toolkits for different application cases. The findings of this chapter were published in journal article in *Energy and Building* as “A framework for a multi-source, data-driven building energy management toolkit” [108] to disseminate its capabilities and application potential, and presented virtually at the 2022 ASHRAE Winter Conference and AHR Expo at Las Vegas, Nevada. In addition, a public GitHub repository was established which contains the toolkit's source code in Python should researchers and practitioners seek technical inspiration; the link can be found in Appendix A.

5.2.2 FRAMEWORK: A web-based application to derive insights

Further to the presented novel framework and various data-driven approaches in the literature, a web-based application of a multi-source, data-driven building energy management toolkit was developed to facilitate its functionalities to a broader audience; case studies were conducted to demonstrate its capabilities. These case studies further highlighted the potential of multi-source, data-driven building operations analytics to identify various operational deficiencies and exposed the limitations of employing multi-sourced analytics to multiple buildings. The web application further disseminates user data input and results retrieval of the toolkit's seven discrete functions, as well as information regarding each function's motivation, capabilities, and how to respond to insights derived from its outputs. The results of this chapter are under review for publication at IBPSA-Canada's eSim 2022 Conference as "FRAMEWORK: A multi-source web application to identify operational deficiencies".

5.2.3 Interviews with building operators and facility managers

As demonstrated by Rock *et al.* [111] and Abuimara *et al.* [12], an understanding of the challenges that building operations personnel face in day-to-day affairs and novel approaches can augment research efforts in their respective fields and guide industry towards improved energy performance. To this end, this chapter addresses a gap in the literature by presenting the interpretations of building operations personnel exposed to data-driven insights; a knowledge gap was identified, abetted by a general unfamiliarity of data-driven analytical tools and an inherently risk-adverse environment, and initiatives were proposed with the goal of facilitating utilization of data-driven insights for the operations profession and advancing the prospective application of data-driven approaches

in the literature. The findings of this chapter are under review for publication in *Energy and Buildings* as “Does a knowledge gap contribute to the performance gap? Interviews with building operators to identify how data-driven insights are interpreted”.

5.3 Recommendations for future work

Though the development and preliminary applications of the toolkit have been researched, several research areas which warrant further study were identified; these were as follows:

- Throughout the research, data used to generate visualizations and KPIs primarily spanned a one-year period and at an hourly resolution. However, variations in sensing and metering capabilities and storage capacity in individual buildings may differ the available reference period and resolution. The effects of prolonging or shortening the reference period and the temporal resolution of data to the reliability of the information derived from the outputs should be further explored.
- Throughout the research, case studies buildings were limited to within ASHRAE climate zone 6A [114]. As such, many of the assessing parameters used to trigger hard and soft fault alarms in the AHU anomaly detection function may be valid exclusively for buildings experiencing similar environmental conditions. Future work should investigate the effects of further expanding datasets to buildings from different climate zones to verify if existing threshold hold true and the degree to which differing conditions warrant adjustment of these thresholds.
- The insights derived can be used to remediate underlying operational deficiencies and improve energy performance. A quantitative assessment of the financial and energy use impact of optimizing controls and fixing hard and soft faults should be conducted to further substantiate the toolkit’s practical implications. The insights

can also be used to improve indoor air quality and occupant satisfaction; a qualitative assessment should verify improvements to these if the appropriate measures incited by the outputs were taken.

- Due to limited access or availability of metadata labels, Wi-Fi device count, and operator work order data, testing and verification of the metadata inference, end-use disaggregation, occupancy using Wi-Fi device count, and complaint analytics functions were scarcely conducted. These functions can derive additional insights for operations personnel in informing their duties; their interpretations and thus, practical applicability, of these functions' outputs were also omitted. Future work should source these missing data types and extract energy-saving insight from these functions to improve their robustness and understand its intended recipient's comprehension of the outputs.
- The interviewed operators' hesitancy to make operational changes may be abetted by their unfamiliarity with data-driven analytical approaches and their ability to inform best practice. Future work should evaluate the responses of operators who have implemented changes as a result of data-driven insights, or are familiar with data-driven operation and maintenance approaches, to understand how practiced improvements affect their interpretations. Additionally, a greater understanding of how building operators assess risk of making operational changes or whether such risks are absorbed by a third-party ESCO would add context to the participants' willingness to apply data-driven insights. Further research should also attempt to identify the financial and technological barriers to widespread utilization of data-driven methods in industry.

- The practical applications of the toolkit rely on effective and informed understanding of the generated outputs and proper use of the functions (i.e., sourcing the required data types, formatting the data, invoking the correct function for the user's intended use case). The interpretive and manual nature of its utilization may deter users, specifically operations personnel, from using the toolkit, and opens the possibility of user misinterpretation of the outputs which may even produce adverse effects to energy performance. Future work should explore the benefits and challenges of integrating the capabilities of the toolkit within existing BASs and EMSs which can automatically implement corrective measures.

At its core, the toolkit serves as a template or example of a multi-source, data driven energy management toolkit. It is not intended as a universal solution that will accurately and reliably derive energy-saving insights in all buildings in all climates. For the scope of this research, the toolkit was developed with medium to large commercial and institutional buildings in mind. Residential buildings, building containing specialized facilities (e.g., laboratories), or buildings with nonstandard occupancy patterns would constitute a derivation of the toolkit to suit their unique operating conditions and protocols.

Appendix A

The GitHub repository containing the toolkit's functions library can be accessed here:

https://github.com/Carleton-DBOM-Research-Group/Building_energy_management_toolkit.git

The web-based application can be accessed here (Refer to the GitHub link if this does not work):

<http://building-energy-management-toolkit.com/home>

References

- [1] Natural Resources Canada, *Survey of commercial and institutional energy use - Buildings 2009*. Ottawa: Energy Publications, 2012.
- [2] T. Ramesh, R. Prakash, and K. K. Shukla, “Life cycle energy analysis of buildings: An overview,” *Energy and Buildings*, vol. 42, no. 10. Elsevier Ltd, pp. 1592–1600, 01-Oct-2010, doi: 10.1016/j.enbuild.2010.05.007.
- [3] U.S. Department of Energy, “2011 Buildings Energy Data Book,” *Off. Energy Effic. Renew. Energy*, 2012.
- [4] American Society of Heating Refrigerating and Air-Conditioning Engineers Inc, *ASHRAE Guideline 36-2018: High-Performance Sequences of Operation for HVAC Systems*. 2018.
- [5] K. Amasyali and N. M. El-Gohary, “A review of data-driven building energy consumption prediction studies,” *Renew. Sustain. Energy Rev.*, vol. 81, pp. 1192–1205, Jan. 2018, doi: 10.1016/J.RSER.2017.04.095.
- [6] S. Katipamula and M. R. Brambley, “Review article: Methods for fault detection, diagnostics, and prognostics for building systems—A review, part I,” *HVAC R Res.*, vol. 11, no. 1, pp. 3–25, 2005, doi: 10.1080/10789669.2005.10391123.
- [7] K. W. Roth, P. Llana, M. Feng, D. Westphalen, and M. Y. Feng, “The Energy Impact of Faults in U.S. Commercial Buildings,” 2004.
- [8] S. Katipamula and M. R. Brambley, “Review article: Methods for fault detection, diagnostics, and prognostics for building systems—a review, part II,” *HVAC R Res.*, vol. 11, no. 2, pp. 169–187, 2005, doi: 10.1080/10789669.2005.10391133.
- [9] E. Mills *et al.*, “The cost-effectiveness of commercial-buildings commissioning: A Meta-Analysis of Energy and Non-Energy Impacts in Existing Buildings and New Construction in the United States,” *LBNL – 56637*, 2004.
- [10] C. Fan, D. Yan, F. Xiao, A. Li, J. An, and X. Kang, “Advanced data analytics for enhancing building performances: From data-driven to big data-driven approaches Article History,” *Build. Simul.* 2020, pp. 1–22, Oct. 2020, doi: 10.1007/s12273-020-0723-1.
- [11] H. Burak Gunay, Z. Shi, I. Wilton, and J. Bursill, “Disaggregation of commercial

- building end-uses with automation system data,” *Energy Build.*, vol. 223, p. 110222, Sep. 2020, doi: 10.1016/j.enbuild.2020.110222.
- [12] T. Abuimara, B. W. Hobson, B. Gunay, W. O’Brien, and M. Kane, “Current state and future challenges in building management: Practitioner interviews and a literature review,” *J. Build. Eng.*, vol. 41, p. 102803, Sep. 2021, doi: 10.1016/J.JOBE.2021.102803.
- [13] P. De Wilde, “The gap between predicted and measured energy performance of buildings: A framework for investigation,” *Autom. Constr.*, vol. 41, pp. 40–49, May 2014, doi: 10.1016/j.autcon.2014.02.009.
- [14] B. Bordass and W. B. Associates, “Energy Performance of Non-Domestic Buildings: Closing the Credibility Gap,” in *8th International Conference on Improving Energy Efficiency in Commercial Buildings*, 2004.
- [15] C. van Dronkelaar, M. Dowson, C. Spataru, and D. Mumovic, “A Review of the Regulatory Energy Performance Gap and Its Underlying Causes in Non-domestic Buildings,” *Front. Mech. Eng.*, vol. 1, p. 17, Jan. 2016, doi: 10.3389/fmech.2015.00017.
- [16] C. van Dronkelaar, M. Dowson, C. Spataru, E. Burman, and D. Mumovic, “Quantifying the Underlying Causes of a Discrepancy Between Predicted and Measured Energy Use,” *Front. Mech. Eng.*, vol. 5, p. 20, May 2019, doi: 10.3389/fmech.2019.00020.
- [17] A. C. Menezes, A. Cripps, D. Bouchlaghem, and R. Buswell, “Predicted vs. actual energy performance of non-domestic buildings: Using post-occupancy evaluation data to reduce the performance gap,” *Appl. Energy*, vol. 97, pp. 355–364, Sep. 2012, doi: 10.1016/j.apenergy.2011.11.075.
- [18] H. P. Tuniki and P. Gultekin-Bicer, “A Comparative Case Study Approach: Identifying the Discrepancy between Energy Performance Results,” in *AEI 2017: Resilience of the Integrated Building - Proceedings of the Architectural Engineering National Conference 2017*, 2017, doi: 10.1061/9780784480502.065.
- [19] S. Katipamula *et al.*, “Small- and Medium-Sized Commercial Building Monitoring and Controls Needs: A Scoping Study,” Richland, WA (United States), Oct. 2012.
- [20] L. Pérez-Lombard, J. Ortiz, and C. Pout, “A review on buildings energy

- consumption information,” *Energy Build.*, vol. 40, no. 3, pp. 394–398, Jan. 2008, doi: 10.1016/j.enbuild.2007.03.007.
- [21] B. Dong, Z. O’Neill, and Z. Li, “A BIM-enabled information infrastructure for building energy Fault Detection and Diagnostics,” *Autom. Constr.*, vol. 44, pp. 197–211, Aug. 2014, doi: 10.1016/j.autcon.2014.04.007.
- [22] A. Ghahramani, S. A. Karvigh, and B. Becerik-Gerber, “HVAC system energy optimization using an adaptive hybrid metaheuristic,” *Energy Build.*, vol. 152, pp. 149–161, Oct. 2017, doi: 10.1016/j.enbuild.2017.07.053.
- [23] T. Abuimara, B. Gunay, A. Abdelalim, M. Ouf, and S. Gilani, “Modelling Occupants in Buildings: Stakeholders’ Workshop on Current Barriers, Challenges and Needs,” in *IBPSA-Canada eSim 2018*, 2018.
- [24] C. Miller *et al.*, “The Building Data Genome Project 2, energy meter data from the ASHRAE Great Energy Predictor III competition,” *Sci. Data*, vol. 7, no. 1, p. 368, 2020, doi: 10.1038/s41597-020-00712-x.
- [25] Y. Zhang, Z. O’Neill, B. Dong, and G. Augenbroe, “Comparisons of inverse modeling approaches for predicting building energy performance,” *Build. Environ.*, vol. 86, pp. 177–190, Apr. 2015, doi: 10.1016/j.buildenv.2014.12.023.
- [26] H. Burak Gunay, W. Shen, G. Newsham, and A. Ashouri, “Detection and interpretation of anomalies in building energy use through inverse modeling,” *Sci. Technol. Built Environ.*, vol. 25, no. 4, pp. 488–503, Apr. 2019, doi: 10.1080/23744731.2019.1565550.
- [27] Z. O’Neill and C. O’Neill, “Development of a probabilistic graphical model for predicting building energy performance,” *Appl. Energy*, vol. 164, pp. 650–658, Feb. 2016, doi: 10.1016/j.apenergy.2015.12.015.
- [28] M. L. Marceau and R. Zmeureanu, “Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings,” *Energy Convers. Manag.*, vol. 41, no. 13, pp. 1389–1403, Sep. 2000, doi: 10.1016/S0196-8904(99)00173-9.
- [29] E. Mocanu, P. H. Nguyen, and M. Gibescu, “Energy disaggregation for real-time building flexibility detection,” in *IEEE Power and Energy Society General Meeting*, 2016, vol. 2016-November, doi: 10.1109/PESGM.2016.7741966.

- [30] K. Basu, V. Debusschere, A. Douzal-Chouakria, and S. Bacha, “Time series distance-based methods for non-intrusive load monitoring in residential buildings,” *Energy Build.*, vol. 96, pp. 109–117, Jun. 2015, doi: 10.1016/j.enbuild.2015.03.021.
- [31] H. Akbari and S. J. Konopacki, “Application of an end-use disaggregation algorithm for obtaining building energy-use data,” *J. Sol. Energy Eng. Trans. ASME*, vol. 120, no. 3, pp. 205–210, Aug. 1998, doi: 10.1115/1.2888070.
- [32] Y. Ji, P. Xu, and Y. Ye, “HVAC terminal hourly end-use disaggregation in commercial buildings with Fourier series model,” *Energy Build.*, vol. 97, pp. 33–46, Jun. 2015, doi: 10.1016/j.enbuild.2015.03.048.
- [33] B. Doherty and K. Trenbath, “Device-level plug load disaggregation in a zero energy office building and opportunities for energy savings,” *Energy Build.*, vol. 204, p. 109480, Dec. 2019, doi: 10.1016/j.enbuild.2019.109480.
- [34] C. Miller, Z. Nagy, and A. Schlueter, “A review of unsupervised statistical learning and visual analytics techniques applied to performance analysis of non-residential buildings,” *Renewable and Sustainable Energy Reviews*, vol. 81. Elsevier Ltd, pp. 1365–1377, 01-Jan-2018, doi: 10.1016/j.rser.2017.05.124.
- [35] Y. Himeur, K. Ghanem, A. Alsalemi, F. Bensaali, and A. Amira, “Artificial intelligence based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives,” *Applied Energy*, vol. 287. Elsevier Ltd, p. 116601, 01-Apr-2021, doi: 10.1016/j.apenergy.2021.116601.
- [36] S. Touzani, B. Ravache, E. Crowe, and J. Granderson, “Statistical change detection of building energy consumption: Applications to savings estimation,” *Energy Build.*, vol. 185, pp. 123–136, Feb. 2019, doi: 10.1016/j.enbuild.2018.12.020.
- [37] Y. Li and Z. O’Neill, “A critical review of fault modeling of HVAC systems in buildings,” *Building Simulation*, vol. 11, no. 5. Tsinghua University Press, pp. 953–975, 01-Oct-2018, doi: 10.1007/s12273-018-0458-4.
- [38] Z. Shi and W. O’Brien, “Development and implementation of automated fault detection and diagnostics for building systems: A review,” *Automation in Construction*, vol. 104. Elsevier B.V., pp. 215–229, 01-Aug-2019, doi: 10.1016/j.autcon.2019.04.002.
- [39] W. Kim and S. Katipamula, “A review of fault detection and diagnostics methods

- for building systems,” *Sci. Technol. Built Environ.*, vol. 24, no. 1, pp. 3–21, Jan. 2018, doi: 10.1080/23744731.2017.1318008.
- [40] H. Burak Gunay, W. Shen, and G. Newsham, “Data analytics to improve building performance: A critical review,” *Automation in Construction*, vol. 97. Elsevier B.V., pp. 96–109, 01-Jan-2019, doi: 10.1016/j.autcon.2018.10.020.
- [41] A. Afram and F. Janabi-Sharifi, “Theory and applications of HVAC control systems - A review of model predictive control (MPC),” *Building and Environment*, vol. 72. Elsevier Ltd, pp. 343–355, 01-Feb-2014, doi: 10.1016/j.buildenv.2013.11.016.
- [42] S. Prívvara, J. Široký, L. Ferkl, and J. Cigler, “Model predictive control of a building heating system: The first experience,” *Energy Build.*, vol. 43, no. 2–3, pp. 564–572, Feb. 2011, doi: 10.1016/j.enbuild.2010.10.022.
- [43] J. Ma, J. Qin, T. Salsbury, and P. Xu, “Demand reduction in building energy systems based on economic model predictive control,” *Chem. Eng. Sci.*, vol. 67, no. 1, pp. 92–100, Jan. 2012, doi: 10.1016/j.ces.2011.07.052.
- [44] J. Široký, F. Oldewurtel, J. Cigler, and S. Prívvara, “Experimental analysis of model predictive control for an energy efficient building heating system,” *Appl. Energy*, vol. 88, no. 9, pp. 3079–3087, Sep. 2011, doi: 10.1016/j.apenergy.2011.03.009.
- [45] M. S. Mirnaghi and F. Haghghat, “Fault detection and diagnosis of large-scale HVAC systems in buildings using data-driven methods: A comprehensive review,” *Energy Build.*, vol. 229, p. 110492, Sep. 2020, doi: 10.1016/j.enbuild.2020.110492.
- [46] P. A. Mathew, L. N. Dunn, M. D. Sohn, A. Mercado, C. Custudio, and T. Walter, “Big-data for building energy performance: Lessons from assembling a very large national database of building energy use,” *Appl. Energy*, vol. 140, pp. 85–93, Feb. 2015, doi: 10.1016/j.apenergy.2014.11.042.
- [47] L. Chen, H. B. Gunay, Z. Shi, W. Shen, and X. Li, “A Metadata Inference Method for Building Automation Systems With Limited Semantic Information,” *IEEE Trans. Autom. Sci. Eng.*, pp. 1–13, May 2020, doi: 10.1109/tase.2020.2990566.
- [48] A. A. Bhattacharya, D. Hong, D. Culler, J. Ortiz, K. Whitehouse, and E. Wu, “Automated metadata construction to support portable building applications,” in *BuildSys 2015 - Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built*, 2015, pp. 3–12, doi:

10.1145/2821650.2821667.

- [49] J.-P. Calbimonte, Z. Yan, H. Jeung, O. Corcho, and K. Aberer, “Deriving semantic sensor metadata from raw measurements,” in *Proceedings of the 5th International Workshop on Semantic Sensor Networks at ISWC*, 2012, vol. 904, no. CONF, pp. 33–48.
- [50] J. Gao, J. Ploennigs, and M. Bergés, “A data-driven meta-data inference framework for building automation systems,” in *BuildSys 2015 - Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built*, 2015, pp. 23–32, doi: 10.1145/2821650.2821670.
- [51] Y. Zhao, T. Li, X. Zhang, and C. Zhang, “Artificial intelligence-based fault detection and diagnosis methods for building energy systems: Advantages, challenges and the future,” *Renewable and Sustainable Energy Reviews*, vol. 109. Elsevier Ltd, pp. 85–101, 01-Jul-2019, doi: 10.1016/j.rser.2019.04.021.
- [52] R. Yan, Z. Ma, G. Kokogiannakis, and Y. Zhao, “A sensor fault detection strategy for air handling units using cluster analysis,” *Autom. Constr.*, vol. 70, pp. 77–88, Oct. 2016, doi: 10.1016/j.autcon.2016.06.005.
- [53] B. Narayanaswamy, B. Balaji, R. Gupta, and Y. Agarwal, “Data Driven Investigation of Faults in HVAC Systems with Model, Cluster and Compare (MCC),” in *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, 2014, pp. 50–59, doi: 10.1145/2674061.2674067.
- [54] H. Wang, Y. Chen, C. W. H. Chan, J. Qin, and J. Wang, “Online model-based fault detection and diagnosis strategy for VAV air handling units,” *Energy Build.*, vol. 55, pp. 252–263, Dec. 2012, doi: 10.1016/j.enbuild.2012.08.016.
- [55] W. Wang, J. Chen, T. Hong, and N. Zhu, “Occupancy prediction through Markov based feedback recurrent neural network (M-FRNN) algorithm with WiFi probe technology,” *Build. Environ.*, vol. 138, pp. 160–170, Jun. 2018, doi: 10.1016/j.buildenv.2018.04.034.
- [56] W. Wang, J. Chen, and T. Hong, “Occupancy prediction through machine learning and data fusion of environmental sensing and Wi-Fi sensing in buildings,” *Autom. Constr.*, vol. 94, pp. 233–243, Oct. 2018, doi: 10.1016/j.autcon.2018.07.007.
- [57] Z. Wang, T. Hong, M. A. Piette, and M. Pritoni, “Inferring occupant counts from

- Wi-Fi data in buildings through machine learning,” *Build. Environ.*, vol. 158, pp. 281–294, Jul. 2019, doi: 10.1016/j.buildenv.2019.05.015.
- [58] Y. Zhao, W. Zeiler, G. Boxem, and T. Labeodan, “Virtual occupancy sensors for real-time occupancy information in buildings,” *Build. Environ.*, vol. 93, no. P2, pp. 9–20, Nov. 2015, doi: 10.1016/j.buildenv.2015.06.019.
- [59] E. Longo, A. E. C. Redondi, and M. Cesana, “Accurate occupancy estimation with WiFi and bluetooth/BLE packet capture,” *Comput. Networks*, vol. 163, p. 106876, Nov. 2019, doi: 10.1016/j.comnet.2019.106876.
- [60] A. Williams, B. Atkinson, K. Garbesi, E. Page, and F. Rubinstein, “Lighting controls in commercial buildings,” *LEUKOS - J. Illum. Eng. Soc. North Am.*, vol. 8, no. 3, pp. 161–180, Jan. 2012, doi: 10.1582/LEUKOS.2012.08.03.001.
- [61] Z. Nagy, F. Y. Yong, M. Frei, and A. Schlueter, “Occupant centered lighting control for comfort and energy efficient building operation,” *Energy Build.*, vol. 94, pp. 100–108, May 2015, doi: 10.1016/j.enbuild.2015.02.053.
- [62] G. Lowry, “Energy saving claims for lighting controls in commercial buildings,” *Energy and Buildings*, vol. 133. Elsevier Ltd, pp. 489–497, 01-Dec-2016, doi: 10.1016/j.enbuild.2016.10.003.
- [63] S. Schiavon and A. K. Melikov, “Energy saving and improved comfort by increased air movement,” *Energy Build.*, vol. 40, no. 10, pp. 1954–1960, Jan. 2008, doi: 10.1016/j.enbuild.2008.05.001.
- [64] S. Naylor, M. Gillott, and T. Lau, “A review of occupant-centric building control strategies to reduce building energy use,” *Renewable and Sustainable Energy Reviews*, vol. 96. Elsevier Ltd, pp. 1–10, 01-Nov-2018, doi: 10.1016/j.rser.2018.07.019.
- [65] A. Aswani, N. Master, J. Taneja, D. Culler, and C. Tomlin, “Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control,” *Proc. IEEE*, vol. 100, no. 1, pp. 240–253, 2012, doi: 10.1109/JPROC.2011.2161242.
- [66] S. Goyal, H. A. Ingle, and P. Barooah, “Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance,” *Appl. Energy*, vol. 106, pp. 209–221, Jun. 2013, doi: 10.1016/j.apenergy.2013.01.039.

- [67] J. Y. Park *et al.*, “A critical review of field implementations of occupant-centric building controls,” *Building and Environment*, vol. 165. Elsevier Ltd, p. 106351, 01-Nov-2019, doi: 10.1016/j.buildenv.2019.106351.
- [68] H. Zou, Y. Zhou, H. Jiang, S. C. Chien, L. Xie, and C. J. Spanos, “WinLight: A WiFi-based occupancy-driven lighting control system for smart building,” *Energy Build.*, vol. 158, pp. 924–938, Jan. 2018, doi: 10.1016/j.enbuild.2017.09.001.
- [69] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal, “Sentinel: Occupancy based HVAC actuation using existing wifi infrastructure within commercial buildings,” in *SenSys 2013 - Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, 2013, pp. 1–14, doi: 10.1145/2517351.2517370.
- [70] N. Alishahi, M. Nik-Bakht, and M. M. Ouf, “A framework to identify key occupancy indicators for optimizing building operation using WiFi connection count data,” *Build. Environ.*, vol. 200, p. 107936, Aug. 2021, doi: 10.1016/j.buildenv.2021.107936.
- [71] B. W. Hobson, D. Lowcay, H. B. Gunay, A. Ashouri, and G. R. Newsham, “Opportunistic occupancy-count estimation using sensor fusion: A case study,” *Build. Environ.*, vol. 159, p. 106154, Jul. 2019, doi: 10.1016/j.buildenv.2019.05.032.
- [72] R. A. Martin, C. C. Federspiel, D. M. Auslander, and A. Dean, “Supervisory Control for Energy Savings and Thermal Comfort in Commercial Building HVAC Systems,” in *Proceedings of the AAAS 2002 Spring Symposium*, 2002.
- [73] S. Dutta, H. Gunay, and S. Bucking, “Using Text Analytics on Operator Logbooks for Performance Benchmarking: A Case Study,” *ASHRAE Trans.*, vol. 125, p. 408+, Oct. 2019.
- [74] S. Dutta, H. Burak Gunay, and S. Bucking, “A method for extracting performance metrics using work-order data,” *Sci. Technol. Built Environ.*, vol. 26, no. 3, pp. 414–425, Mar. 2020, doi: 10.1080/23744731.2019.1693208.
- [75] H. B. Gunay, W. Shen, and C. Yang, “Text-mining building maintenance work orders for component fault frequency,” *Build. Res. Inf.*, vol. 47, no. 5, pp. 518–533, Jul. 2019, doi: 10.1080/09613218.2018.1459004.
- [76] S. Assaf and I. Srour, “Using a data driven neural network approach to forecast

- building occupant complaints,” *Build. Environ.*, vol. 200, p. 107972, Aug. 2021, doi: 10.1016/j.buildenv.2021.107972.
- [77] T. Hong *et al.*, “Commercial Building Energy Saver: An energy retrofit analysis toolkit,” *Appl. Energy*, vol. 159, pp. 298–309, Dec. 2015, doi: 10.1016/j.apenergy.2015.09.002.
- [78] A. Costa, M. M. Keane, J. I. Torrens, and E. Corry, “Building operation and energy performance: Monitoring, analysis and optimisation toolkit,” *Appl. Energy*, vol. 101, pp. 310–316, Jan. 2013, doi: 10.1016/j.apenergy.2011.10.037.
- [79] T. Zhang, A. Al Zishan, and O. Ardakanian, “ODToolkit: A Toolkit for Building Occupancy Detection,” 2019, doi: 10.1145/3307772.3328280.
- [80] H. Li, C. Szum, S. Lisauskas, A. Bekhit, C. Nesler, and S. C. Snyder, “Targeting Building Energy Efficiency Opportunities: An Open-source Analytical & Benchmarking Tool,” in *ASHRAE Transactions*, 2019.
- [81] A. Miller, K. Carbonnier, and M. Frankel, “Sample Municipal Portfolio Benchmarking Report: A Portfolio Analysis of Municipal Buildings in a Sample City Mark Frankel,” Portland, Feb. 2017.
- [82] H. N. Rafsanjani and C. Ahn, “Linking Building Energy-Load Variations with Occupants’ Energy-Use Behaviors in Commercial Buildings: Non-Intrusive Occupant Load Monitoring (NIOLM),” in *Procedia Engineering*, 2016, vol. 145, pp. 532–539, doi: 10.1016/j.proeng.2016.04.041.
- [83] B. W. Hobson, H. B. Gunay, A. Ashouri, and G. R. Newsham, “Clustering and motif identification for occupancy-centric control of an air handling unit,” *Energy Build.*, vol. 223, p. 110179, Sep. 2020, doi: 10.1016/j.enbuild.2020.110179.
- [84] H. B. Gunay, W. Shen, G. Newsham, and A. Ashouri, “Modelling and analysis of unsolicited temperature setpoint change requests in office buildings,” *Build. Environ.*, vol. 133, pp. 203–212, Apr. 2018, doi: 10.1016/j.buildenv.2018.02.025.
- [85] W. J. N. Turner, A. Staino, and B. Basu, “Residential HVAC fault detection using a system identification approach,” *Energy Build.*, vol. 151, pp. 1–17, 2017, doi: 10.1016/j.enbuild.2017.06.008.
- [86] C. Fan, F. Xiao, C. Yan, C. Liu, Z. Li, and J. Wang, “A novel methodology to explain and evaluate data-driven building energy performance models based on

- interpretable machine learning,” *Appl. Energy*, vol. 235, pp. 1551–1560, Feb. 2019, doi: 10.1016/j.apenergy.2018.11.081.
- [87] Z. Afroz, B. W. Hobson, H. B. Gunay, W. O’Brien, and M. Kane, “How Occupants Affect Decision-Making Building Operators’,” *ASHRAE J.*, vol. 62, no. November, pp. 22–30, 2020.
- [88] A. Ghahramani, K. Zhang, K. Dutta, Z. Yang, and B. Becerik-Gerber, “Energy savings from temperature setpoints and deadband: Quantifying the influence of building and system properties on savings,” *Appl. Energy*, vol. 165, pp. 930–942, Mar. 2016, doi: 10.1016/j.apenergy.2015.12.115.
- [89] Z. Afroz, H. Burak Gunay, W. O’Brien, G. Newsham, and I. Wilton, “An inquiry into the capabilities of baseline building energy modelling approaches to estimate energy savings,” *Energy Build.*, vol. 244, p. 111054, Aug. 2021, doi: 10.1016/j.enbuild.2021.111054.
- [90] American Society of Heating Refrigerating and Air-Conditioning Engineers Inc, *ASHRAE Guideline 14-2014: Measurement of Energy, Demand, and Water Savings*. 2014.
- [91] *National Energy Code of Canada for Buildings*. National Research Council, 2017.
- [92] American Society of Heating Refrigerating and Air-Conditioning Engineers Inc, *ASHRAE Standard 62.1-2007: Ventilation for Acceptable Indoor Air Quality*. 2007.
- [93] J. Koh, B. Balaji, D. Sengupta, J. McAuley, R. Gupta, and Y. Agarwal, “Scrabble: Transferrable semi-automated semantic metadata normalization using intermediate representation,” in *BuildSys 2018 - Proceedings of the 5th Conference on Systems for Built Environments*, 2018, vol. 18, pp. 11–20, doi: 10.1145/3276774.3276795.
- [94] H. B. Gunay and Z. Shi, “Cluster analysis-based anomaly detection in building automation systems,” *Energy Build.*, vol. 228, p. 110445, Dec. 2020, doi: 10.1016/j.enbuild.2020.110445.
- [95] D. Darwazeh, B. Gunay, and J. Duquette, “Development of Inverse Greybox Model-Based Virtual Meters for Air Handling Units,” *IEEE Trans. Autom. Sci. Eng.*, pp. 1–14, Jul. 2020, doi: 10.1109/tase.2020.3005888.
- [96] L. Wang, P. Mathew, and X. Pang, “Uncertainties in energy consumption introduced by building operations and weather for a medium-size office building,” *Energy*

- Build.*, vol. 53, pp. 152–158, Oct. 2012, doi: 10.1016/j.enbuild.2012.06.017.
- [97] H. B. Gunay, G. Newsham, A. Ashouri, and I. Wilton, “Deriving sequences of operation for air handling units through building performance optimization,” *J. Build. Perform. Simul.*, vol. 13, no. 5, pp. 501–515, Sep. 2020, doi: 10.1080/19401493.2020.1793221.
- [98] D. Darwazeh, J. Duquette, and B. Gunay, “Virtual metering of heat supplied by hydronic perimeter heaters in variable air volume zones,” in *Proceedings of the 5th International Workshop on Non-Intrusive Load Monitoring*, 2020, pp. 49–53, doi: 10.1145/3427771.3429389.
- [99] C. C. Federspiel, “Predicting the frequency and cost of hot and cold complaints in buildings,” *HVAC R Res.*, vol. 6, no. 4, pp. 289–305, 2000, doi: 10.1080/10789669.2000.10391418.
- [100] C. C. Federspiel, R. A. Martin, and H. Yan, “Recalibration of the complaint prediction model,” *HVAC R Res.*, vol. 10, no. 2, pp. 179–200, 2004, doi: 10.1080/10789669.2004.10391098.
- [101] H. B. Gunay, W. O’Brien, I. Beausoleil-Morrison, W. Shen, G. Newsham, and I. Macdonald, “The effect of zone level occupancy characteristics on adaptive controls,” in *15th International Conference of the International Building Performance Simulation Association, San Francisco, CA, USA*, 2017, doi: <https://dx.doi.org/10.26868/25222708.2017.252>.
- [102] A. Ashouri, G. R. Newsham, Z. Shi, and H. B. Gunay, “Day-ahead Prediction of Building Occupancy using WiFi Signals,” in *IEEE International Conference on Automation Science and Engineering*, 2019, vol. 2019-August, pp. 1237–1242, doi: 10.1109/COASE.2019.8843224.
- [103] B. W. Hobson, H. B. Gunay, A. Ashouri, and G. R. Newsham, “Occupancy-based predictive control of an outdoor air intake damper: A case study,” *Proc. IBPSA-Canada eSim 2021 Conf.*, 2021.
- [104] K. Amasyali and N. M. El-Gohary, “A review of data-driven building energy consumption prediction studies,” *Renewable and Sustainable Energy Reviews*, vol. 81. Elsevier Ltd, pp. 1192–1205, 01-Jan-2018, doi: 10.1016/j.rser.2017.04.095.
- [105] B. Narayanaswamy, B. Balaji, R. Gupta, and Y. Agarwal, “Data driven investigation

- of faults in HVAC systems with Model, Cluster and Compare (MCC),” in *BuildSys 2014 - Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, 2014, pp. 50–59, doi: 10.1145/2674061.2674067.
- [106] K. Yan, C. Zhong, Z. Ji, and J. Huang, “Semi-supervised learning for early detection and diagnosis of various air handling unit faults,” *Energy Build.*, vol. 181, pp. 75–83, Dec. 2018, doi: 10.1016/J.ENBUILD.2018.10.016.
- [107] Z. Chen, M. K. Masood, and Y. C. Soh, “A fusion framework for occupancy estimation in office buildings based on environmental sensor data,” *Energy Build.*, vol. 133, pp. 790–798, Dec. 2016, doi: 10.1016/J.ENBUILD.2016.10.030.
- [108] A. A. Markus, B. W. Hobson, H. B. Gunay, and S. Bucking, “A framework for a multi-source, data-driven building energy management toolkit,” *Energy Build.*, vol. 250, p. 111255, Nov. 2021, doi: 10.1016/J.ENBUILD.2021.111255.
- [109] Z. Li, Y. Han, and P. Xu, “Methods for benchmarking building energy consumption against its past or intended performance: An overview,” *Appl. Energy*, vol. 124, pp. 325–334, Jul. 2014, doi: 10.1016/J.APENERGY.2014.03.020.
- [110] M. Goulden and A. Spence, “Caught in the middle: The role of the Facilities Manager in organisational energy use,” *Energy Policy*, vol. 85, pp. 280–287, Oct. 2015, doi: 10.1016/J.ENPOL.2015.06.014.
- [111] S. Rock, M. R. Hosseini, B. Nikmehr, I. Martek, S. Abrishami, and S. Durdyev, “Barriers to ‘green operation’ of commercial office buildings: Perspectives of Australian facilities managers,” *Facilities*, vol. 37, no. 13–14, pp. 1048–1065, Sep. 2019, doi: 10.1108/F-08-2018-0101/FULL/PDF.
- [112] N. Harris, T. Shealy, H. Kramer, J. Granderson, and G. Reichard, “A framework for monitoring-based commissioning: Identifying variables that act as barriers and enablers to the process,” *Energy Build.*, vol. 168, pp. 331–346, Jun. 2018, doi: 10.1016/J.ENBUILD.2018.03.033.
- [113] O. C. Robinson, “Sampling in Interview-Based Qualitative Research: A Theoretical and Practical Guide,” <https://doi.org/10.1080/14780887.2013.801543>, vol. 11, no. 1, pp. 25–41, Jan. 2013, doi: 10.1080/14780887.2013.801543.
- [114] American Society of Heating Refrigerating and Air-Conditioning Engineers Inc, *ASHRAE Standard 169-2013: Climatic Data for Building Design Standards*. 2013.

- [115] B. Balaji, N. Weibel, and Y. Agarwal, “Managing Commercial HVAC Systems: What do Building Operators Really Need?,” *arXiv Prepr. arXiv1612.06025*, Dec. 2016, doi: <https://doi.org/10.48550/arXiv.1612.06025>.
- [116] P. Domingues, P. Carreira, R. Vieira, and W. Kastner, “Building automation systems: Concepts and technology review,” *Comput. Stand. Interfaces*, vol. 45, pp. 1–12, Mar. 2016, doi: 10.1016/J.CSI.2015.11.005.
- [117] C. Duarte, B. Acker, R. Grosshans, M. Manic, K. Van Den Wymelenberg, and C. Rieger, “Prioritizing and Visualizing Energy Management and Control System Data to Provide Actionable Information for Building Operators,” in *Western Energy Policy Research Conference*, 2011.
- [118] X. Yang and S. Ergan, “Evaluation of visualization techniques for use by facility operators during monitoring tasks,” *Autom. Constr.*, vol. 44, pp. 103–118, Aug. 2014, doi: 10.1016/J.AUTCON.2014.03.023.
- [119] M. Aune, T. Berker, and R. Bye, “The missing link which was already there: Building operators and energy management in non-residential buildings,” *Facilities*, vol. 27, no. 1–2, pp. 44–55, 2009, doi: 10.1108/02632770910923081/FULL/PDF.
- [120] A. C. Roussac and H. Huang, “For building operators, what difference does a target make?,” *Energy Effic.*, vol. 13, no. 3, pp. 459–471, Mar. 2020, doi: 10.1007/S12053-019-09788-W/FIGURES/8.