FRAMeWORK: A multi-source web-based application for identifying suboptimal energy use management

Andre A. Markus[1], Brodie W. Hobson[1], H. Burak Gunay[1], Scott Bucking[1]
[1]Carleton University, Ottawa, Canada

Abstract

Uninformed building operations and mechanical faults and anomalies often lead to poor building energy performance. Though analytical software tools incorporating several data-driven methods that identify energy deficiencies and energy-saving opportunities have been proposed as an accessible method for building operators to improve energy performance, these tools are primarily prototypes and have scarcely seen application in more than one building. In this paper, the software platform of a novel multi-source, data-driven energy management toolkit incorporating inverse energy modelling, fault detection and diagnosis, load disaggregation, and occupancy and occupant complaint analytic methods, was presented and demonstrated on four case study buildings; the generated key performance indicators and visuals were presented and examined.

Introduction

With building operations accounting for 80%-90% of a building's total life cycle energy (Ramesh, Prakash, and Shukla 2010), remediating energy use deficiencies resulting from suboptimal building operations and mechanical faults is paramount to addressing the widespread 'performance gap' and reducing carbon emissions (Amasyali and El-Gohary 2018; Menezes et al. 2012). In the past two decades, attempts to address these deficiencies have manifested as various methods that identify energy use anomalies (Ma et al. 2012; Ji, Xu, and Ye 2015), hard and soft faults resulting in excessive energy use (Kim and Katipamula 2018; Narayanaswamy et al. 2014), and opportunities to reduce energy use resulting from overventilation (Hobson et al. 2019) using only archived data streams, typically from the building automation system (BAS).

Several publications have proposed and demonstrated methods for automatic fault detection and diagnostic (AFDD), which pertains to the identification and diagnosis of hard and soft faults in a building's heating, ventilation, and air conditioning (HVAC) system (Mirnaghi and Haghighat 2020). Data-driven variants of this method, also known as black-box models (Kim and Katipamula 2018), 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 (Darwazeh, Duquette, and Gunay 2020). Likewise, data-driven energy anomaly detection and occupant-centric controls approaches have been applied to directly remediate identified energy use anomalies and deficiencies in case study buildings (Chen, Masood, and Soh 2016; Hobson et al. 2019). 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 (Gunay et al. 2018; Darwazeh, Duquette, and Gunay 2020).

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 paper demonstrates the current capabilities of a novel multi-source, data-driven building energy management toolkit with four case study buildings. Available air handling unit (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 generated key performance indicators (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.

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, fault detection and diagnosis, 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, computerized maintenance and management system (CMMS), Wi-Fi device count, and motion detection data, and output a variety of visualizations and key performance indicators (KPIs) intended to provide energy-saving insights for building energy professionals.

Backend

The toolkit is compiled in Python as a functions library and is stored in a public GitHub repository (Markus 2021). 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 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.

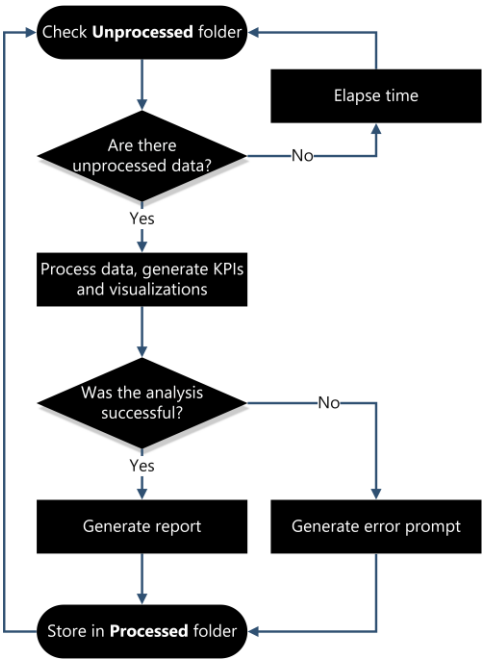

Figure 1: Backend workflow for user data storing and processing

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 (Markus et al. 2021). 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 presents and compares the rate of 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 assess the actual impact, intended or otherwise, of an AHU operating schedule in reducing building-level energy use, particularly outside a building's scheduled operating hours.

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 variable air volume (VAV) terminal units 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 use 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 augment ideal indoor air temperature setpoints and ventilation schedules.

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 augment occupancy-centric control decisions (e.g., ventilation rates, scheduling) to reduce excess ventilation/conditioning during partial occupant capacity or vacancy.

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.

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 using a virtual private server (VPS), contains descriptions of each function as well as instructions for data pre-formatting, inputting, and interpretating the generated outputs (i.e., KPIs, visualizations), along with examples of generated outputs. Figure 2 illustrates the intended workflow between users, the web interface, and the server.

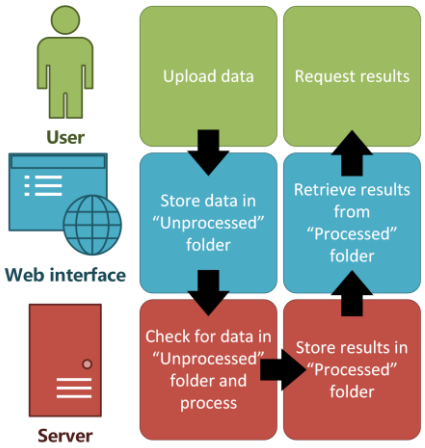

Figure 2: Workflow between users, web interface, and the server.

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 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 (Markus 2021).

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 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 1: Availability of building data by organization

Org.	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. Meter data contained hourly measured energy use for heating, cooling, and electricity separately. However, electricity meter data from organization A included chiller energy use from two nearby buildings which could not be disaggregated; these data could not be used since meter data must pertain exclusively to the analysed building.

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 2 presents the number of AHUs

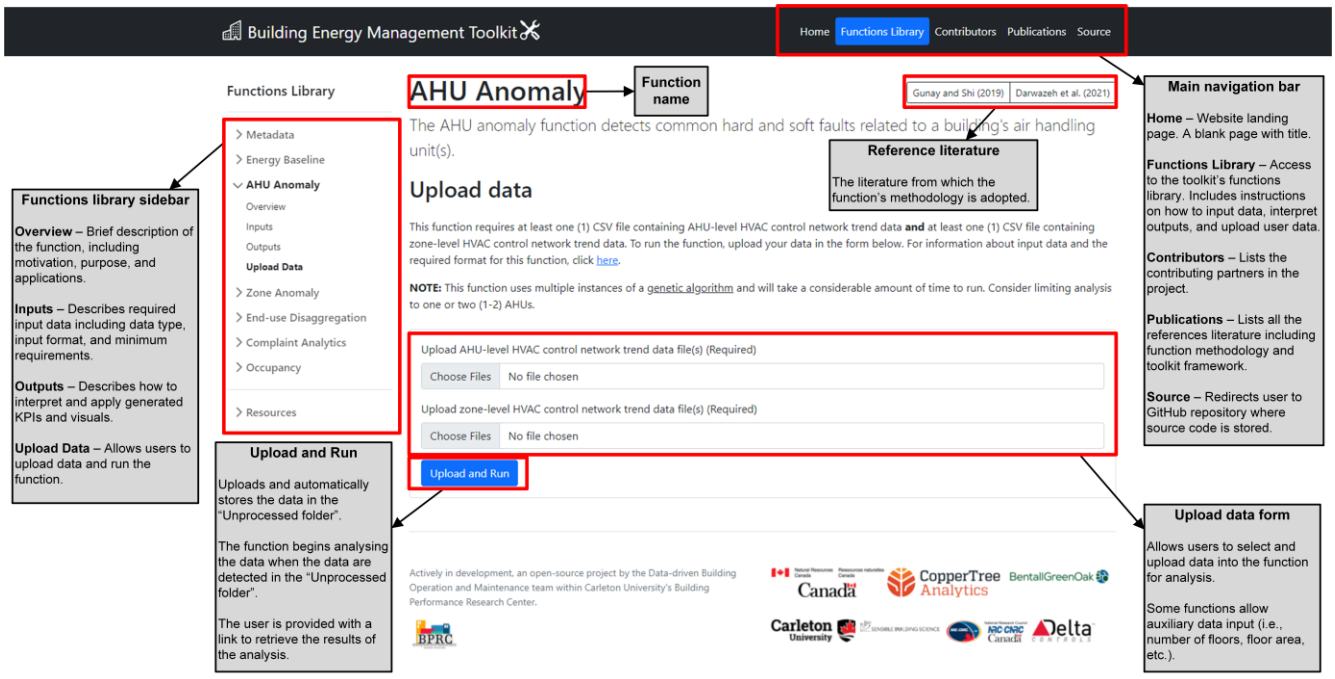

Figure 3: Annotated screenshot of upload page of the AHU anomaly detection function of the toolkit's website

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 2: Number of AHUs and thermal zones analysed by organization

Org.	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 functions and inputted through the web interface. Available meter data were uploaded to the energy baseline function, excluding electricity data from organization B and C since these buildings primarily operated on electricity-based heating; this function inherently assumed that electricity data include only plug-in/lighting, fan/pump, and on-site chiller loads. 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 available motion detection data.

Once data were 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.

Results

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 4, for heating and cooling energy use separately. Meter data for organization D were unavailable and, along with electricity data for the aforementioned limitations, were omitted from this analysis. In all cases, the rate of heating energy use 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 rate of energy use and this is reflected by the resulting 0% schedule effectiveness (SE) and highest afterhours energy use (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 use. If no schedule existed, implementing one can more effectively reduce afterhours heating energy use 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 a higher rate of afterhours heating energy use.

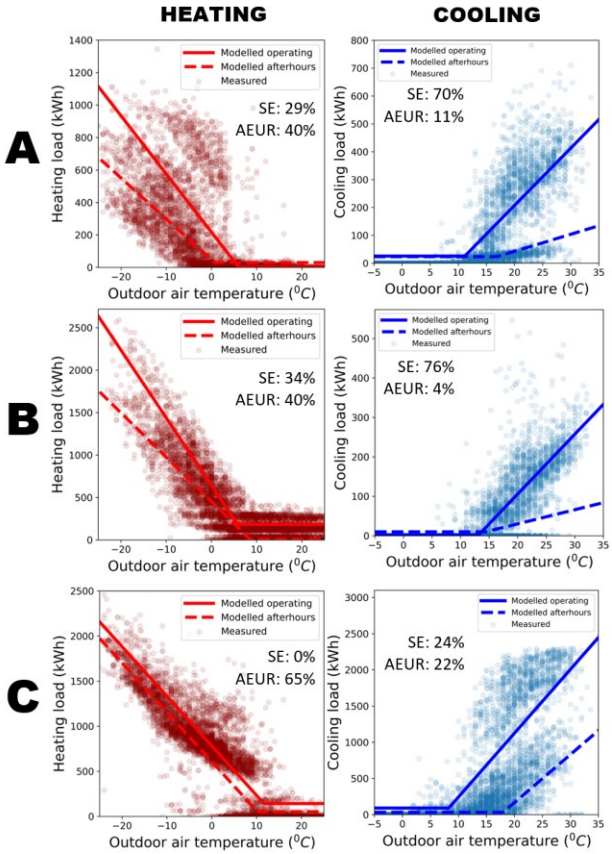

Figure 4: Heating (left column) and cooling (right column) energy use during (solid line) and outside (dashed line) AHU operating hours for organization A, B, and C.

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.

AHU anomaly detection

The AHU anomaly function produced a plot comparing SAT and return air temperatures for each analysed AHU, seen in Figure 5. In almost all AHUs in organization A and C, there was evidence of a SAT reset strategy present, some closely resembling ideal SAT settings highlighted by the schematic; this ideal region was determined by (Gunay et al. 2020). 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 efficiency, is unlikely to be intentional.

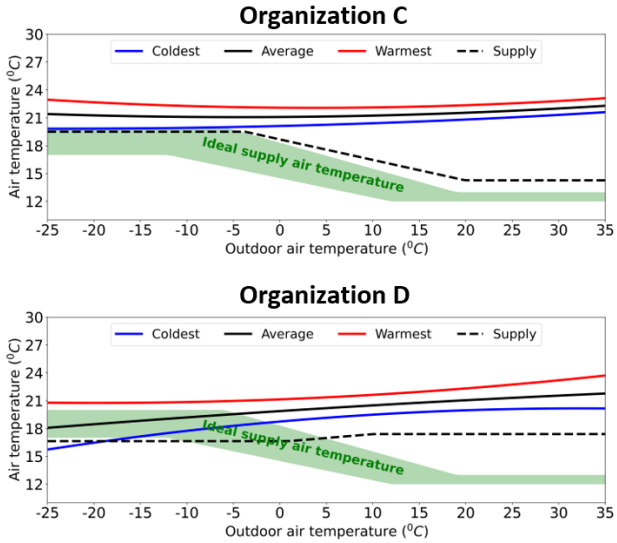

Figure 5: Comparison of split-range controller diagrams for AHUs in organization C and D. SAT follows a near ideal scheme (Gunay et al. 2020) for AHUs in organization C but exhibits a contradicting scheme in organization D.

Table 3 and Table 4 tally the number of hard and soft faults, respectively, identified by the AHU anomaly function by organization.

Table 3: Number of identified hard faults by organization

Org.	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 4: Number of identified soft faults by organization

Org.	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 (ASHRAE 2018) 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 6.

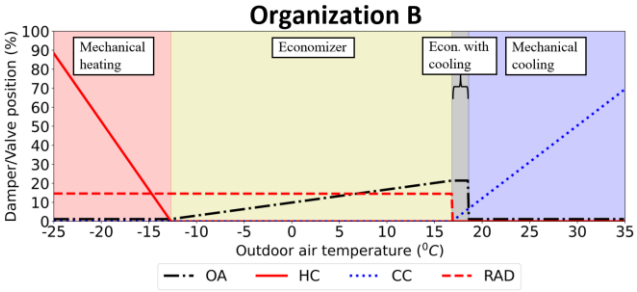

Figure 6: Split-range controller diagram for an AHU in organization B. Outdoor air damper position (OA), heating coil valve position (HC), cooling coil valve position (CC), and fraction of active perimeter heaters (RAD) are plotted along outdoor air temperature.

In other 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 cool outdoor air to offset energy from 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 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.

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 inputted data 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 5 summarizes the zone health indices per season and organization.

Table 5: Summary of zone indices by organization

Org.	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 7 for organization D.

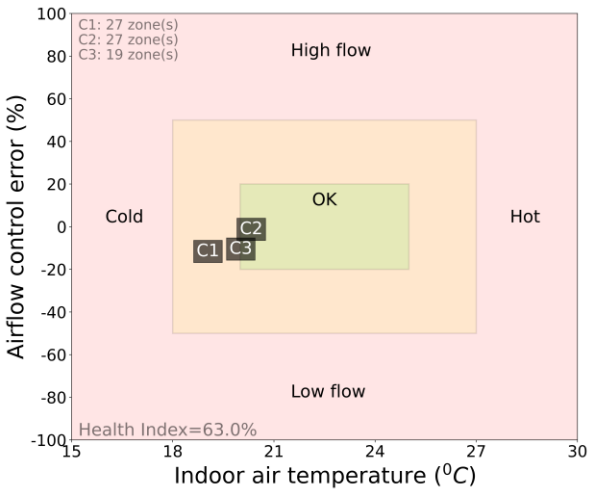

Figure 7: Heating season zone cluster diagram for organization D.

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.

Occupancy using motion detection

The occupancy function did not produce any visualizations since motion detection data, in place of Wi-Fi device count data, were inputted. However, the function computed typical building-level earliest arrival, latest arrival, latest departure, and longest break duration seen in Table 6. Motion detection data for organization C and D were unavailable and were omitted from this analysis.

Table 6: Earliest arrival and departure, latest departure, and longest break duration by organization

Org.	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.

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 schedule effectiveness 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.

Conclusion

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 multi-sourced building operations analytics were exposed by expanding the inputted 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 be incompatible with certain functions and produce inconclusive results, as

with 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.

Acknowledgement

This research is supported by the research contract GIBE-2187 with Natural Resources Canada, CopperTree Analytics, Delta Controls, Bentall GreenOak, and Sensible Building Science.

References

Amasyali, Kadir, and Nora M. El-Gohary. 2018. "A Review of Data-Driven Building Energy Consumption Prediction Studies." *Renewable and Sustainable Energy Reviews*. Elsevier Ltd. https://doi.org/10.1016/j.rser.2017.04.095.

American Society of Heating Refrigerating and Air-Conditioning Engineers Inc. 2018. *ASHRAE Guideline 36-2018: High-Performance Sequences of Operation for HVAC Systems*. *ASHRAE*.

Chen, Zhenghua, Mustafa K. Masood, and Yeng Chai Soh. 2016. "A Fusion Framework for Occupancy Estimation in Office Buildings Based on Environmental Sensor Data." *Energy and Buildings* 133 (December): 790–98. https://doi.org/10.1016/J.ENBUILD.2016.10.030.

Darwazeh, Darwish, Jean Duquette, and Burak Gunay. 2020. "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*, 49–53. New York, NY, USA: Association for Computing Machinery (ACM). https://doi.org/10.1145/3427771.3429389.

Gunay, H. Burak, Guy Newsham, Araz Ashouri, and Ian Wilton. 2020. "Deriving Sequences of Operation for Air Handling Units through Building Performance Optimization." *Journal of Building Performance Simulation* 13 (5): 501–15. https://doi.org/10.1080/19401493.2020.1793221.

Gunay, H. Burak, Weiming Shen, Guy Newsham, and Araz Ashouri. 2018. "Modelling and Analysis of Unsolicited Temperature Setpoint Change Requests in Office Buildings." *Building and Environment* 133 (April): 203–12. https://doi.org/10.1016/j.buildenv.2018.02.025.

Hobson, Brodie W., Daniel Lowcay, H. Burak Gunay, Araz Ashouri, and Guy R. Newsham. 2019. "Opportunistic Occupancy-Count Estimation Using Sensor Fusion: A Case Study." *Building and Environment* 159 (July): 106154. https://doi.org/10.1016/j.buildenv.2019.05.032.

Ji, Ying, Peng Xu, and Yunyang Ye. 2015. "HVAC Terminal Hourly End-Use Disaggregation in Commercial Buildings with Fourier Series Model." *Energy and Buildings* 97 (June): 33–46. https://doi.org/10.1016/j.enbuild.2015.03.048.

Kim, Woohyun, and Srinivas Katipamula. 2018. "A Review of Fault Detection and Diagnostics Methods for Building Systems." *Science and Technology for the Built Environment* 24 (1): 3–21. https://doi.org/10.1080/23744731.2017.1318008.

Ma, Jingran, Joe Qin, Timothy Salsbury, and Peng Xu. 2012. "Demand Reduction in Building Energy Systems Based on Economic Model Predictive Control." *Chemical Engineering Science* 67 (1): 92–100. https://doi.org/10.1016/j.ces.2011.07.052.

Markus, Andre. 2021. "Building Energy Management Toolkit." GitHub Repository. 2021. https://github.com/Carleton-DBOM-Research-Group/Building_energy_management_toolkit.git.

Markus, Andre A., Brodie W. Hobson, H. Burak Gunay, and Scott Bucking. 2021. "A Framework for a Multi-Source, Data-Driven Building Energy Management Toolkit." *Energy and Buildings* 250 (November): 111255. https://doi.org/10.1016/J.ENBUILD.2021.111255.

Menezes, Anna Carolina, Andrew Cripps, Dino Bouchlaghem, and Richard Buswell. 2012. "Predicted vs. Actual Energy Performance of Non-Domestic Buildings: Using Post-Occupancy Evaluation Data to Reduce the Performance Gap." *Applied Energy* 97 (September): 355–64. https://doi.org/10.1016/j.apenergy.2011.11.075.

Mirnaghi, Maryam Sadat, and Fariborz Haghighat. 2020. "Fault Detection and Diagnosis of Large-Scale HVAC Systems in Buildings Using Data-Driven Methods: A Comprehensive Review." *Energy and Buildings* 229 (September): 110492. https://doi.org/10.1016/j.enbuild.2020.110492.

Narayanaswamy, Balakrishnan, Bharathan Balaji, Rajesh Gupta, and Yuvraj Agarwal. 2014. "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*, 50–59. Association for Computing Machinery, Inc. https://doi.org/10.1145/2674061.2674067.

Ramesh, T., Ravi Prakash, and K. K. Shukla. 2010. "Life Cycle Energy Analysis of Buildings: An Overview." *Energy and Buildings*. Elsevier Ltd. https://doi.org/10.1016/j.enbuild.2010.05.007.