

Automated Anomaly Detection and Diagnosis for Real Time Carbon Minimization

Hao Huang, and Craig Roussac, Buildings Alive

ABSTRACT

Energy models are increasingly being used to support automated detection of system anomalies in commercial buildings. Most approaches analyze energy demand profiles at the system and equipment levels to identify significant energy performance issues and optimization opportunities. However, this approach has the potential to mislead operators seeking to minimize greenhouse gas emissions (GHG) rather than raw energy usage, because energy models do not address the time-varying GHG emissions intensity of electricity networks. To the best of the authors' knowledge, no published studies apply real-time GHG predictive modeling for anomaly detection in buildings.

In this study, we propose a novel anomaly detection framework which combines energy modeling with real-time GHG emissions intensity analysis. First, we built a machine learning model to predict daily electricity consumption, based on exogenous inputs such as weather factors, calendar effects and occupancy level. The model was then used to identify anomalies by considering the deviations between actual and predicted energy profiles. A GHG intensity model was run in parallel to calculate the real-time GHG emissions intensity of electricity supplies by considering grid generation mix. The GHG impacts of the anomalies were then assessed based on the time-varying GHG emissions intensity data. The output was a list of the top equipment anomalies to be investigated, and their energy and carbon impacts. The experimental results from two office buildings located on different electricity networks show that when considering real-time GHG emissions, it is possible to prioritize the anomalies that present the largest GHG emission reduction opportunity. The result also shows that even if two sets of anomalies have similar energy impacts, their percentage difference in GHG impacts can be as high as 95%. The proposed method can help building operators and technicians to manage anomalies/faults and reduce GHG emissions by considering their buildings' demand flexibility.

Introduction

According to the International Energy Agency, for the global economy to reach net zero greenhouse gas (GHG) emissions by 2050, annual clean energy investment worldwide will need to more than triple by 2030 to around USD \$4 trillion. This ambition will require huge increases in electricity system flexibility by 2030—such as demand response, behavioral changes, energy storage and 'zero-carbon-ready' buildings, among many other things—to stabilize loads and maintain reliable supplies (IEA, 2021).

In the United States, buildings account for approximately 75% of electricity use and 40% of all primary energy use and associated greenhouse gas (GHG) emissions (DoE 2015), so they have the potential to make a significant contribution to balancing supply variability. Grid-integrated and *zero-carbon-ready buildings* (a term used extensively in IEA, 2021) can provide a flexible resource for electricity systems, using connectivity and automation to manage network demand and the operation of energy storage capacity. This flexibility can, in turn, underpin the economic case for investments in renewable electricity generation infrastructure.

Anomaly detection has proven to be an effective tool for detecting faults swiftly and reducing the impact of them in commercial buildings, particularly in heating, ventilation and air-conditioning (HVAC) systems. Anomaly detection, also referred to as outlier detection, deals with finding patterns in the signal that are abnormal, unexpected, or interesting. By identifying anomalies, one can trigger alerts to either repair a piece of broken equipment, faulty sensors, or to suggest more optimized control settings. In the building energy domain, the use of historical data to detect and diagnose faults is becoming increasingly popular. A common practice to detect anomalous building operation is to implement predictive energy models using real time interval meter data through Energy Information Systems (EIS) (Kramer et al. 2020).

Automated fault detection and diagnosis (AFDD) methods can be classified into quantitative, qualitative, and process history-based methods (Katipamula and Brambley 2005). When process history-based AFDD algorithms are considered, regression models are built using field data. This approach sometimes leverages some degree of engineering knowledge to enable more detailed fault diagnosis (Granderson et. al. 2020). An advantage of using historical data to detect sensor faults is that modelling is typically less complicated than quantitative and qualitative approaches. It has been shown that FDD users achieve median savings of 9% after two years of implementation (Kramer et al. 2020)

There are two types of anomaly detection in the building energy domain: equipment fault diagnosis and anomaly detection of energy consumption data (Zhou et al. 2021). The advantage of using energy consumption data to perform anomaly detection is that it can capture major energy-related issues without the need to collect Building Management System (BMS) data from the field. Seem (2007) used a robust statistical method to determine if a building's energy consumption was significantly different to previous energy consumption. The method clusters consumption data into different groups based on day of week features and then picks up the outliers using boxplots. Stephen *et al.* (2016) presented a hybrid AFDD approach that combines building energy models and a machine-learning algorithm to detect and diagnose faults using whole-building data only. Their approach requires fewer sensor data and configurations. Some researchers have also tried to combine the energy consumption data with field data so that one can apply a top-down approach to detect major energy anomalies and find their related causes (Huang, Roach and Roussac 2020).

Conventional anomaly detection approaches focus on identifying materially significant system anomalies that impact building energy use. This was appropriate when most electricity supplies were generated from fossil fuels. However, with decarbonization of electricity grids, negative price and low / unstable demand events are becoming increasingly common around the world and they usually correspond with times of maximum renewable energy generation and, therefore, lower GHG emissions electricity. All of this requires insight into real-time price and GHG emissions factors; without it building operators are not able to actively manage their impacts. This paper addresses these challenges. To the best of the authors' knowledge, although the applications of anomaly detection have been broadly discussed in the past, to-date no study has been conducted to apply real-time GHG predictive modeling to perform anomaly detection in buildings.

In this study, we propose a new anomaly detection framework that considers the real-time GHG emissions intensity of electricity grids. We apply a semi-supervised learning anomaly detection approach described in our previous work (Huang, Roach, and Roussac 2020), coupled with a real-time GHG intensity model running in parallel to calculate real-time GHG emissions intensity of electricity supplies. We then assess both the energy and carbon impacts of identified anomalies based on the grid's time-varying GHG emissions intensity. Experimental results from two office buildings under different grid systems show that when considering the GHG emissions, the priority levels of the identified anomalies change according to the time-varying characteristics of the grid: anomalies occurring at the high carbon intensity period may have higher carbon impacts, even though their energy impacts may be

small. The anomalies that occurred during low carbon intensity periods seem to have less significance. Through experiment we found that even if two groups of anomalies have similar energy impact, their percentage difference in carbon impacts can be as high as 95%. The developed method can help building operators and technicians to manage anomalies/faults for lower cost and environmental impact when considering demand flexibilities in their buildings.

Methodology

Data Description

The building used as the primary subject of this study is a multi-story, multi-tenanted office building located in the downtown of Adelaide, South Australia. The energy data were collected at 15-minute intervals from utility network meters supplying the base-building consumption including heating, ventilation, and air-conditioning (HVAC), common area lighting, elevators, and other loads – but not tenancy consumption which included office light and power. The measured and forecast weather data were obtained from the Australian Bureau of Meteorology (BOM).

Figure 1 plots the daily average temperature against daily energy consumption from March 2019 to March 2022. From the figures, we can see that there is a notable difference in data patterns between the period impacted by COVID-19 lockdowns and the pre-pandemic period, for both weekdays and weekends. Therefore, when selecting the baseline model for the training period, it was considered important to not mix the two periods in the training set otherwise the modelling performance would be degraded. Since our purpose is to use the model to detect anomalies in real-time, we always use the most recent data for training to capture the most recent status of the building operation.




Figure 1. Daily interval consumption data compared to temperature for weekdays (left) and all working hours (right).

Energy Prediction Model

In our previous work, we demonstrated a peak demand prediction model which can accurately predict a building's electricity demand for any 15-minute interval up to 120 hours (5 days) in advance (Roussac and Huang 2018). A similar modelling approach is applied in this work. Because the model is used for detecting anomalies instead of demand response, the forecast weather inputs are replaced with observed inputs. For predicting the interval daily energy consumption, we trained a random regression (RF) model based on a series of independent variables, such as weather, occupancy status, and calendar effects. Since many factors are available but not all of them are related to the outputs (energy usage),

we ran a feature selection approach to select the most significant variables and rank the importance of each feature as provided by the algorithm. The final feature list is listed in Table 1. It was found that some variables such as wind speed have little impact on the outputs, so they were excluded in the final model.

As the second step, we ran a hyper-parameter tuning process to maximize the performance of a regression model. The random forest (RF) builds N independent regression trees, splitting the data at each node. The features considered to perform each splitting are a random subset of the total number of features. The results obtained for the N trees were then averaged to provide a model output. At the growth of N, the performance improves. However, the convergence rate was shown to saturate quite quickly and thus we fixed the number of trees to N=100. The optimal maximum number of features to be used at each splitting has an optimal value, chosen to be 10.

Table 1. Description of input variables for 15 min interval model

Feature name	Type	Value
Occupancy status	Categorical	1% to 100%
Outdoor temperature	Numerical	°C
Moisture content	Numerical	%
Temperature readings from previous hours	Numerical	°C
Time of day indicator	Categorical	[1,2....96]
Monday/Post-holiday flag	binary	[1,0]
Day of week indicator	Categorical	[1,2,3,4,5,6,7]

We trained a random forest model using one year’s worth of historical data, from Nov 2020 to Nov 2021. As new data are fed into the system, the model automatically retrain to include the most recent status. This allows us to keep tracking changes happening at the building. Figure 2 displays the actual electricity consumption (blue) overlayed with the outputs of the developed random forest regressor model (green) during the training window. It shows that the random forest regression model can accurately predict the interval building metering data.




Figure 2. Actual electricity consumption (blue) versus a random forest regressor model for a typical building. The dashed dark line is the ambient temperature.

The above random forest regression model returns the “most likely” energy usage, as the algorithm is designed to minimize overall errors of a costs function. However, for buildings that are continually being fine-tuned, the energy performance may improve leading to the possibility that models trained off historical data may over-predict demand. Similarly, a model could under-predict a

building's energy performance if a building's performance is trending negatively due to the existence of system faults. To solve this problem, we made some improvements to the model so that it always returns the lowest possible prediction (target usage) instead of “the most likely prediction”. The target performance represents the best performance a building has achieved for a given period under similar weather and occupancy conditions. A detailed description of a statistic approach to generate the target is presented in (Roussac and Huang 2020). The “target usage” scenarios will be used for anomaly detection in the rest of this paper.

Building Level Anomaly Detection

We applied the top-down anomaly detection approach which was introduced in our previous study (Huang, Roach, and Roussac 2020) to detect building energy anomalies. For a given day that we are interested in, the actual metering readings are compared with the model prediction, and if the deviation is above certain thresholds, the time stamp will be flagged as anomalies.

For simplicity, we use a fixed threshold to detect anomalies in this study. An interval point was identified as anomalous if its absolute percentage error against prediction was above 15% or the difference in power demand is above 10 kW. The reason we apply both percentage and absolute kW errors to determine thresholds is because buildings have quite different usage patterns at different times of a day. For example, during after-hours operation the absolute error between actual and prediction is typically quite small because the building is mostly unoccupied during that time. But there could still be anomalies we are interested in such as 24/7 pump operation, and improper purge control. The percentage error threshold allows us to capture the anomalies regardless of the amplitudes of the energy usage happened at that time. We used 15% because that is the typical errors bound of the random forest model. The absolute percentage error is calculated using the following equation.

$$P = \frac{|Q_{actual,t} - Q_{target,t}|}{Q_{actual,t}} \quad (1)$$

where Q_{actual} and $Q_{prediction}$ denote actual and target usage at each time stamp, respectively, and P represents the deviation between actual and target usage. Figure 3 illustrates the anomaly detection results of an example building for two weeks. One can see that the actual (blue line) matches the prediction (green line) for most of the time. Whenever the actual deviates from the prediction and reaches the pre-defined threshold, the time intervals were then marked as anomalies in orange.




Figure 3. Detecting energy anomaly using RF model. The forecasting is shown in green, the actual values of the time series in blue, and the anomalies detected are highlighted in orange.

Equipment Level Anomaly Detection

Once the anomalies are identified at the building level, such as in Figure 5, the next question to ask is “what has caused the anomalies to happen?”. When physical submeters are installed, this question can be answered easily. However fine-grained physical meters are not always available, and their readings do not always cover the component levels (e.g., a supply air fan). Here we use BMS data and technical specifications of the equipment to estimate the energy usage of individual HVAC components by creating ‘soft meters’. For detailed methods related to soft meter modelling please refer to (Huang, Roach, and Roussac 2020). With the soft meter models prepared, we can then make a linkage between the anomalies at the building level and subsystem level anomalies. We rank the soft meters based on the difference between actual soft-meter usage and predicted softer-meter usage. The “top anomalies” represent the equipment that contributes the most to the building level anomalies. We use three different impact values (high, medium, and low) to indicate different levels of anomalies associated with different components. For example, if the percentage deviation of a time stamp is above 50% and kW deviation is above a pre-define kW threshold, then the point is defined as a “high impact” anomaly. If the percentage deviation of a time stamp is between 15% and 50% and kW deviation is below the threshold, then the point is defined as a “medium impact” anomaly. The point will be flagged as a low impact anomaly if the percentage deviation of a time stamp is below 15%. The pre-defined kW threshold is set based on the building’s actual demand profile.

Figure 4 illustrates an example of the top anomalies list and their energy impacts. One can see that in this case, the most anomalous equipment includes the return air fan and supply air fan of AHU-10-08 and chilled water pumps because they have the highest deviations against the model expectation (here we refer to the prediction as “best ever”).




Figure 4. Top anomalies list. The green line represents predicted soft meter usage, and the blue line represents the actual soft meter usage.

Carbon Intensity Calculation

Electricity is generated from a mixture of fuels. Different fuel types have different GHG emission factors (measured in units of CO₂-equivalent/MWh) which are published by authoritative sources (see for example EIA 2022). There are two widely recognized approaches to calculating the GHG emissions intensity of electricity supplies: average and marginal. The average emissions intensity corresponds to the overall, e.g., region-wide electricity production including net imports feeding into a grid at a given point in time. The marginal intensity factor reflects the GHG intensity of the marginal generator, i.e., the last and least economical generating unit contributing to the supply. There has been some debate about whether one should use marginal emission factors or average emission factors to

evaluate the carbon savings when performing demand response (Finenko and Cheah 2016; Evans and Sidat 2018). In this study we focus on average carbon intensity because the load in commercial buildings is inconsequential when compared to the overall demand on generators, and thus making no discernible difference to the generation fuel mix (Lowry 2018).

We calculate the average grid carbon intensity using a weighted sum of the emission factors for different fuels according to their contribution to the generation mix. If there is power exchange between two grid systems, the GHG emissions intensity of the imported electricity is also considered. The calculation of GHG emissions intensity requires the following three data sets:

- The total generation (MWh) output from each relevant generating system.
- The GHG emissions per unit of electricity (t CO₂-e /MWh) generated by each relevant generating system.
- Imported generation (MWh) and its GHG emissions intensity, calculated as per the above two points.

The above data sets are publicly available and can be downloaded from National Electricity Market (NEM) web (AEMO 2022) the carbon intensity can then be calculated using the following equations:

$$CDEII = \frac{\sum_i CDE_i + E_{import} * EF_{import}}{\sum_i E_i + E_{import}} \quad (2)$$

$$CDE_i = E_i \times EF_i \quad (3)$$

where CDE denotes Carbon Dioxide Equivalent emissions (t CO₂-e) from a generating system, CDEII denotes grid carbon intensity (kg CO₂-e /kWh), E denotes generation output (MWh) from a generating system, E_{import} is imported generation (MWh), and EF_{import} = Carbon intensity of the imported electricity. Figure 5 shows the 5-min interval generation mix in South Australia (SA) from November 11 to November 28, 2021. It shows that wind energy accounts for a large proportion of electricity production in the state. Using fuel mixed data, the time varying carbon intensity is calculated using equations 1 and 2, and the result is illustrated in Figure 6. As we can see, the carbon intensity data shows uncertain patterns, which is mainly due to the uncertainties associated with wind energy generation.




Figure 5. Fuel mix data in South Australia.



Figure 6. Carbon intensity of South Australia grid.

Consider time-varying carbon intensity in anomaly detection

Based on the real time carbon intensity calculated in Figure 6, we can then obtain the “real emissions profiles” for both actual and predicted usage of the building, using the following equation:

$$e_i = \text{CDEI}_i \times Q_i \quad (4)$$

where Q denotes the energy usage at interval i , and e presents the emissions of the building at time interval i . Figure 7 shows actual emissions versus predicted emissions for the same period.

Interestingly, we can see the time series plot of emissions has a totally different trend compared to its consumption profile (Figure 3). The gap between actual emissions and predicted emissions becomes less visible, because the deviations become much smaller when multiplying the energy deviation with the carbon intensity data at each time stamp.



Figure 7. Actual and predicted emission profiles of a SA building.

Results & Discussions

The proposed algorithm was deployed with the anomaly detection algorithm continuously processes weather data, interval metering data, BMS data and fuel mix data feeding into a cloud server. Whenever new data comes in, it triggers the anomaly detection algorithm to perform analysis. The top four detected anomalous faulty components and their corresponding root cause analyses are reported in real-time. We use two office buildings located at different cities of Australia, one in South Australia (SA) and the other in New South Wales (NSW), to illustrate our approach.

SA example

The first test building is a 15-storey office building located in the downtown area of Adelaide, SA. Figure 8 shows the anomaly detection result on 03/16/2022. The graph shows there are three distinguishable anomaly periods that occurred on this day: midnight, early morning, and evening. The red dashed line shows the carbon intensity level was low in the morning and then reached a peak in the late afternoon (~18:00). Figure 9(a) shows that the most anomalous equipment returned from the algorithm is chiller-3. The graph clearly shows that chiller-3 was running out-of-hours and was cycling on/off a couple of times in the early morning, which matches the anomaly symptoms reflected by the building level energy data. Further investigation shows that the cooling valve position of AHU-09 was maintained above 20% throughout day (figure 9b), which explains why chiller-3 was running out-of-hours. Additionally, we noticed that two AHUs running at a higher speed than usual (figures 9b and 9c), which are other contributors to the higher energy usage on this day.

Table 2 compares the carbon impacts and energy impacts of the three anomaly periods. Interestingly, we found the anomalies that occurred from 8:00 pm to 11:45 pm have the highest carbon impacts among the three anomaly clusters, even though their energy impacts were only ranked as the second highest. The anomaly cluster 2 seems to have very minor carbon impacts because the grid electricity had a lower carbon intensity at the time when the energy increase occurred. The carbon impacts associated with anomaly cluster 3 is 74 lbs (95%) higher than cluster 2, even though their energy impacts are quite similar. So, in this case, we can conclude that fixing the after-hours chiller operation is more rewarding in terms of environmental benefits, as compared to fixing the morning spikes or increased AHUs consumption on the experimental day.




Figure 8. Building level anomaly detection result on 16/03/2022.



Figure 9a. Chiller-3 consumption




Figure 9b. AHU-9 cooling valve position




Figure 9c. AHU-10 usage




Figure 9d. AHU-16 usage

Figure 9. Equipment level anomaly detection results at an Adelaide building.

Table 2. Carbon and energy impacts of different anomalies at an Adelaide building

Anomaly cluster	Occurred time	Related equipment	Energy impact	Carbon intensity	Carbon impact
1	00:00 am to 3:00 am	Chiller-3, AHU-9	69 kWh	0.28 lbs/kWh	19 lbs
2	07:30 am to 10:15 am	Chiller-3, AHU-10, AHU-16	158 kWh	0.25 lbs/kWh	40 lbs
3	8:00 pm to 11:45 pm	Chiller-3, AHU-9	152 kWh	0.75 lbs/kWh	113 lbs

NSW example

In the second example, we apply the same approach to an office located in Sydney, NSW, Australia. Different from SA, coal dominates in NSW which makes up more than 60% of generation (AEMO 2021). The main source of renewable energy is solar energy in NSW. Therefore, the NSW grid has a higher average carbon intensity level and a more regular carbon intensity pattern as compared to SA.

Figure 10 compares the actual electricity profile with the predicted electricity profile at the building level. The algorithm has identified that the most anomalous period of the day was between 7:30 am and 9:00 am. We can observe spiky usage from the actual meter readings during this time, while the prediction suggests that the morning profile should be much smoother. The morning spike aligns with the spikes that occurred at the condenser water pump and cooling tower consumption profiles, as shown in Figures 11 (b) and (d). Figures 11 (a) and 11 (c) illustrate that the cycling energy data pattern is closely correlated with chiller and chilled water pump operations. We can see the

increased after-hours usage that happened at 8:00 pm was also related to the chiller and chilled water pump operation.




Figure 10. Building level anomaly detection result on 23/03/2022.




Figure 11a. Chiller consumption.




Figure 11b. Cooling tower consumption.




Figure 11c. Chilled water pump consumption




Figure 11d. Condenser water pump consumption

Figure 11. Equipment level anomaly detection results at a Sydney building

Table 3 compares the carbon and energy impacts of the two anomaly periods. In this case, we can see both carbon and energy impacts of anomaly cluster 1 is high. This is not surprising because the most significant energy increase occurred during this time, despite the impact of the small decrease in GHG emissions intensity during the day. However, it appears that the spike which occurred at 8:00 am requires the most attentions from the site team because the carbon intensity level was still high at that time of day prior to the solar generation starting to take effect. Therefore, in this case, fixing the morning spikes associated with the cooling towers and condenser water pumps may result in higher carbon savings as compared to other anomalies.

Table 3. Carbon and energy impacts of different anomalies at a Sydney building

Anomaly cluster	Occurred time	Related equipment	Energy impact	Carbon intensity	Carbon impact
1	07:30 am to 1:45 pm	Cooling tower, Condenser water pump, Chiller, chilled water pump	338 kWh	1.4 lbs/ kWh	472 lbs.
2	8:00 pm to 11:45 pm	Chiller, chilled water pump	34 kWh	1.6 lbs/ kWh	55 lbs

Conclusions

In this study, we have proposed a novel anomaly detection framework that combines energy consumption prediction with the calculation of real-time grid GHG emissions intensity. The experimental results from two office buildings located on two significantly different electricity grids show that this anomaly detection method can identify operational anomalies that have the highest GHG emissions impacts. The anomalies that occurred at the high carbon intensity period tended to have more significant carbon impacts, even when their energy impacts may have been small. The anomalies that occurred during low carbon intensity periods were less important and could be ranked as low priority issues. It was also found that even if two groups of time-based anomalies have similar energy impact, their percentage difference in carbon impacts can be as high as 95%. We also found that on a grid with relatively low renewable energy penetration, the change in the anomaly ranking was relatively small as compared to a grid with greater renewable energy penetration. Our method allows building operators to prioritize their anomalies/faults by considering the carbon impact of their buildings' demand. However, because the grid carbon intensity may fluctuate from one day to another, the same anomaly may be regarded as having high carbon impacts on one day but fewer impacts on another. Therefore, the future research will investigate the frequency of the occurrence of the anomalies across several months to obtain more useful insights.

References

- AEMO (Australian Energy Market Operator). 2021. *Quarterly Energy Dynamics Q1 2021*. aemo.com.au/-/media/files/major-publications/qed/2021/q1-report.pdf?la=en
- DoE (Department of Energy). 2015. *Quadrennial Technology Review: An Assessment of Energy Technologies and Research Opportunities*. U.S. Department of Energy, Washington www.energy.gov/quadrennial-technology-review-2015
- EIA (Energy Information Administration). 2022. *State Electricity Profiles*. [online] www.eia.gov/electricity/state/ [Accessed March 26, 2022].
- Evans, B., and S. Sabbir. 2018. "The use of temporal factors for improved CO2 emissions accounting in buildings". *Building Services Engineering Research and Technology*, 39(2): 196-210. doi.org/10.1177/0143624417753297

- Finenko, A., and L. Cheah. 2016. "Temporal CO₂ emissions associated with electricity generation: Case study of Singapore". *Energy Policy*, 93: 70-79. doi.org/10.1016/j.enpol.2016.02.039
- IEA (International Energy Agency). 2021. *Net Zero by 2050*. International Energy Agency, Paris www.iea.org/reports/net-zero-by-2050.
- Granderson, J., R. Singla, E. Mayhorn, P. Ehrlich, D. Vrabie., and S. Frank. 2017. "Characterization and Survey of Automated Fault Detection and Diagnostic Tools". Lawrence Berkeley National Laboratory. LBNL-2001075
- Katipamula, S, and M.R. Brambley. 2005. "Review Article: Methods for Fault Detection, Diagnostics, and Prognostics for Building Systems—A Review, Part I." *HVAC&R Research* 11 (1): 3–25.
- Kramer, H., G.J. Lin, C. Curtin, E. Crowe, and J. Granderson. 2020. "Proving the Business case for building analytics". *Lawrence Berkeley National Laboratory*, doi.org/10.20357/B7G022
- Huang, H, C. Roach, and A.C. Roussac. 2020. "From Local Optima to Overall Optimum: A Field Study of Model-Based Anomaly Detection in Commercial Buildings". In *Proceedings of the 2018 ACEEE Summer Study on Energy Efficiency in Buildings*, Pacific Grove, California: ACEEE
- Keane, D., N. Harmsen, and S. Tomevska. 2021. "Solar panels switched off by energy authorities to stabilise South Australian electricity grid". Australian Broadcasting Corporation (ABC). March 17. www.abc.net.au/news/2021-03-17/solar-panels-switched-off-in-sa-to-stabilise-grid/13256572
- McConnell, D., S. Holmes à Court, S. Tan, and N. Cubrilovic. 2022. OpenNEM project. opennem.org.au/
- Parkinson, G. 2021. "'World first': South Australia achieves 100pct solar, and lowest prices in Australia". *RenewEconomy*, January 29. reneweconomy.com.au/world-first-south-australia-achieves-100pct-solar-and-lowest-prices-in-australia/
- Roussac, A.C. and H. Huang. 2018. "Forewarned Is Forearmed: Reducing Peak Demand in Commercial Buildings with Behavior Science". In *Proceedings of the 2018 ACEEE Summer Study on Energy Efficiency in Buildings*, Pacific Grove, California: ACEEE
- Roussac, A.C. and H. Huang. 2020. "For building operators, what difference does a target make?" *Energy Efficiency*, 13(3): 459-471. doi.org/10.1007/s12053-019-09788-w
- Seem, J. 2007. "Using intelligent data analysis to detect abnormal energy consumption in buildings. " *Energy and Buildings*, 39(1), pp.52-58. doi.org/10.1016/j.enbuild.2006.03.033
- Zhou, X., T. Yang., L. Liang., XH. Zi., JW. Yan., and D. Pan. 2021. "Anomaly detection method of daily energy consumption patterns for central air conditioning systems". *Journal of Building Engineering*, 38: 102179. doi.org/10.1016/j.jobbe.2021.102179