Fault detection in residential buildings using sub-hourly meter data

Gilbert Larochelle Martin[1], Benoit Delcroix[1], Simon Sansregret[1] Ahmed Daoud[1]
[1]Laboratoire des technologies de l'énergie (LTE), Hydro-Québec, Shawinigan, Québec, Canada

Abstract

The current deployment of smart meters and home energy management systems provides utilities' residential customers with more customer-centric energy data than ever before. This data can be used to provide residential customers with meaningful insights into their energy consumption and peak power demand. In this paper, machine learning techniques and smart meter data are leveraged to identify personalized periods of over or under consumption which could be indicative of abnormal energy consumption or detrimental occupant behaviour. The lower and upper centiles are modelled using a feed-forward neural network trained using a quantile loss function. The model is then applied to a set of sub-metered residential houses to assess its performance. This model could be used to provide residential building occupants with personalized energy feedback aimed at improving their energy-related behaviour and lower their peak demand and energy consumption.

Introduction

The energy consumption of the residential building sector represents 18% of all energy consumed in Quebec, 66% of which is electricity (Whitmore & Pineau, 2021). Furthermore, the electrical demand registered by Hydro-Québec in January 2022 has reached an historical peak of 40 510 MW. While energy efficiency and demand response programs and technologies have yielded important gains over the years, low-hanging fruits in both fields are getting scarcer. There is still a clear need for new ways of reducing electricity consumption and peak demand province-wide.

In Quebec, 66% of residential homes are heated with electric baseboards heaters (NRCAN, 2011). At the same time, 62 % of the main thermostats in homes are not programmed or not programmable and only 53% used wintertime nightly setbacks with a programmable thermostat (NRCAN, 2019). As connected devices, such as connected thermostats and appliances, become more and more available, it is easier to program and schedule their use. Still, more advanced uses than simple programming may be within reach of residential consumers in the coming years that may yield further energy efficiency and demand response potential.

The commercially available smart home systems are maturing and offer the energy sector stakeholders a more detailed window into consumer energy-related behaviour than ever before. Hence, in the coming years, there will be an opportunity to leverage the consumer trend data reported by those devices beyond the laboratory setting into real world operations. For example, the measured data may be used to promote energy appropriate behaviours and new energy-related equipment selection. While advanced consumer feedback is not new and has been proven effective in many studies (Ehrhardt-Martinez et al., 2010), there are more advanced features that can consume, interpret and communicate information to the consumer based on this data. One of the looming opportunities is to provide an advanced and personalized feedback to each customer in the form of a Fault Detection and Diagnosis (FDD) tool. This tool can help identify abnormal energy consumption periods and guide the consumer to improve their energy-related behaviours.

FDD has been applied extensively in Commercial and Institutional (CI) buildings at the whole building and systems level. Katipamula and Brambley (2005) described the fundamental FDD methods applied to building systems and proposed a selection method to identify the most appropriate type of algorithm. In a more recent review, Kim and Katipamula (2018) identified the recent published information grouped by FDD method and application. They only identified 24 papers aimed at whole building FDD. Frank et al. (2018) assessed the technical, market and cost barriers of FDD tools in small-scale commercial buildings. They also proposed cost targets for future FDD tools. Machine learning methods applied to FDD have been reviewed by Zhao et al. (2019) and they might see more applications as more sensors, such as the ones reviewed in Dong et al. (2019), become available.

At the residential level, the scientific literature is scarce and only a few studies could be identified, mainly at air conditioning systems. Rogers et al. (2019) highlighted the potential for the application of FDD methods in the residential sector to reduce maintenance costs, reduce energy consumption, improve equipment commissioning and reduce peak demand. They also concluded that the

initial studies using FDD methods do not attempt to identify faulty behaviour in specific homes and that the field would grow considerably in the next decade with large-scale analysis becoming available. Indeed, applications of big data tools and machine learning in the residential building stock has recently been attempted by Khani et al. (2021) using custom indicators and k-means clustering and Guo et al. (2021) using multivariate statistics and the Mahalanobis distance.

Continuing this trend, this paper proposes a FDD methodology based on an artificial feed forward neural network using a quantile loss function to identify possible « *faults* » or abnormal hourly consumption in residential customers' data at the whole building level. The methodology leverages measured data and machine learning to provide advanced feedback to the residential consumers.

Methods

The methodology employed to develop the FDD black-box algorithm can be presented in two parts: (1) data exploration and (2) model development.

Data exploration

Data exploration is an essential step to identify the shortcomings of basic regression modelling. The data used in this paper comes from a previous study aimed at assessing energy savings, winter demand response impacts and general customer experience with Line-Voltage Communicating Thermostats (Fournier et al., 2018). The data used from this study was recorded over the span of a year beginning May 2017 for 30 sub-metered houses located in the vicinity of Trois-Rivières, Québec. The data includes house descriptions, hourly electric baseboard heaters thermostats measured room temperature, setpoint and energy consumption supplemented with 1 minute interval data of the main electric loads (spa, dryer, water heater, etc.). Table 1 presents an overview of the characteristics of the sub metered houses. In this paper, only the whole house energy consumption data is used. There is still a lot more value to the dataset that will be discussed in the following sections.

Table 1. Household, building and practical characteristics of the submetering study. (Adapted from (Fournier et al., 2018))

Characteristics	
Number of participants	30
Average annual consumption (kWh)	26 280
Normalized average annual consumption (kWh/m²)	123.71
Household size (average)	3.3
House built before 1970	33.3%
Floor area less than 1500 ft2	53%
Wall mounted heat pump presence	27%
Constant setpoint (no setback)	47%
Household occupancy follows a regular pattern/schedule	97%

One of the simplest way to model the energy consumption of any building is the « *PRInceton Scorekeeping Method (PRISM)* » (Fels, 1986). This method is commonly used in multiple aspect of building energy analysis. Excluding the base consumption from the piece-wise regression, the simplest PRISM model reduces to Equation (1).

$$y = ax + b \quad (1)$$

The energy consumption (y) is modelled as a linear function of the degree days (x) using a regression coefficient (a) and an intercept (b). A simplified version of this PRISM model using a fixed base temperature was applied to the available data using multiple time aggregation levels. Figure 1 and Figure 2 present respectively the boxplots of coefficients (a) and the intercepts (b) of PRISM models for different levels of time aggregation.

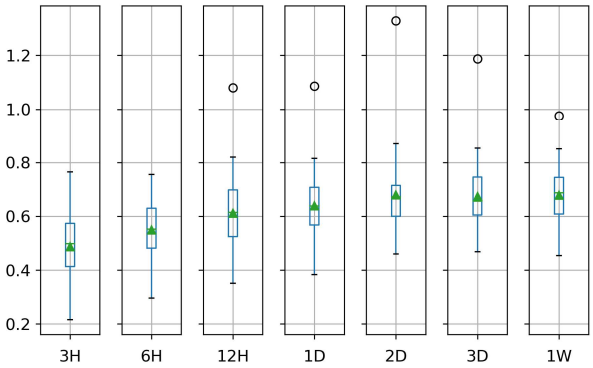

Figure 1. PRISM models coefficients (kWh/DD)

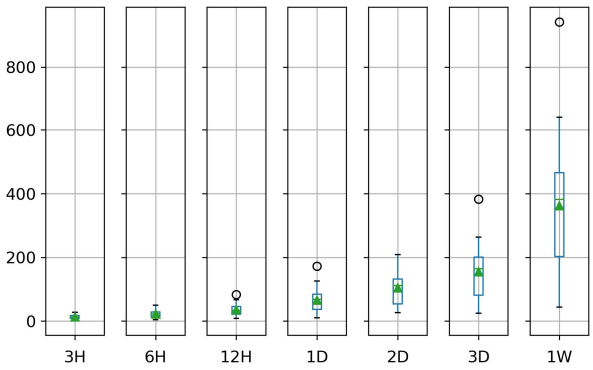

Figure 2. PRISM models intercepts (kWh)

Figure 2 shows an expected growth in the intercept value results as the intercept is expected to grow in relation to a higher time aggregation level. Figure 3 presents the boxplot of the regression scores (R^2) of the linear models given the different time aggregation levels where the green line is the median, the green triangle is the mean and the boundaries of the whiskers are within 1.5 times the interquartile range. From inspection of Figure 3, the performance of the model is clearly deteriorating as the time aggregation is lowered. Model scores below 0.7-0.8 are usually heuristically considered inappropriate for analysis in this context.

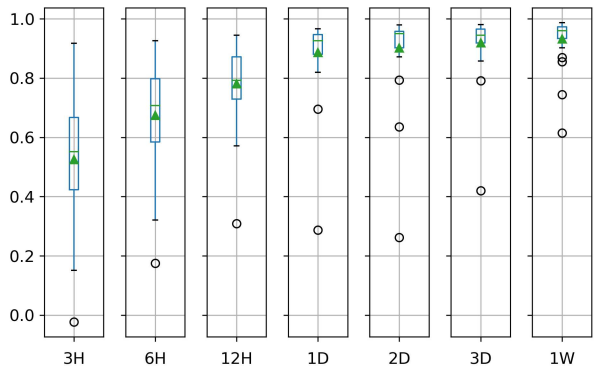

Figure 3. PRISM models scores

Regarding the development of a residential FDD tool, it is clear from the presented results that a more advanced energy consumption prediction method is needed at the hourly level. In this paper, a machine learning black box algorithm model development was performed using a simple feed-forward neural network.

Model Development

To address the identified lack of performance from the simple model outlined in the previous section, a black-box model approach leveraging a machine learning algorithm was developed using a simple feedforward neural network. The developed model leverages sub-hourly data (15 min) aggregated to an hourly frequency from each individual house. The data is used to train a feed-forward artificial neural network, which can be schematically represented in Figure 4.

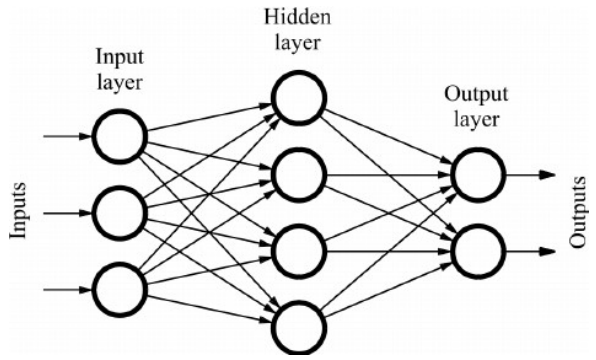

Figure 4. Representation of a feed-forward artificial neural network (ANN).

The input layer of the ANN uses only a few available features :

- Heating Degree-Hours (HDH);
- Hour of the day;
- Day of the week.

The base used for computing the HDH is 13ºC, which was arbitrarily chosen based on heuristics for a residential setting. This is compatible with the average residential balance temperature range of 7ºC to 16ºC (Mitchell & Braun, 2012). The neural network uses two hidden layers of 32 neurons using ReLU activation functions as shown in Figure 5.

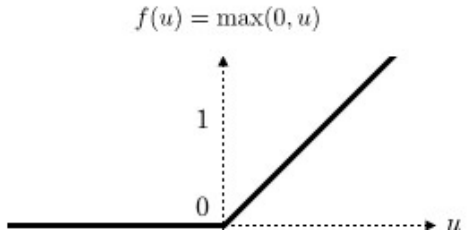

Figure 5. Illustration of the ReLU activation function.

Hence, only the time-stamped energy consumption and local weather data are needed to train the model. The quantile loss function used during training of the neural network is described in equation (1).

$$\mathcal{L}_\alpha(y, \hat{y}) = \sum_{i=1}^{n} 1_{y_i<\hat{y}_i} (1-\alpha) |y_i - \hat{y}_i| + 1_{y_i \geq \hat{y}_i} \alpha |y_i - \hat{y}_i| \quad (1)$$

Hence, the loss function given a specific quantile ($\mathcal{L}_\alpha$) is computed using the specified quantile (α), the measured data (y_i) and the estimated data ($\hat{y}_i$). In this paper, the 1st and 99th percentiles are arbitrarily used as the threshold for fault detection. These criteria could be adjusted or supplemented with other threshold levels for different feedback sensitivity aimed at the consumer. One drawback of using neural networks is the danger of overfitting the training data and concurrently yielding poor performance on unseen datasets. To avoid overfitting the training data, the dropout technique is used on the hidden layers of the network in conjunction with early stopping. The dropout technique is illustrated in Figure 6.

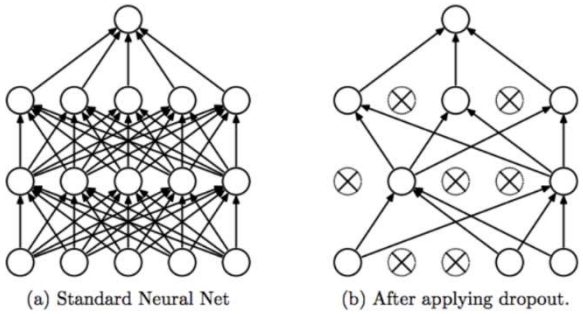

Figure 6. Illustration of a (a) standard neural network and a (b) neural network with a probability of dropout on the input and hidden layers.

A scoring rule is needed to evaluate the quantile coverage of the model. In this paper, the model is evaluated using the interval score (Gneiting & Raftery, 2007). This metric is described in Equation 2 using the desired quantile (α), data vector (x), the upper bound (u) and lower bound (l). The function represented by $\mathbb{1}$ is equal to one if the expression in the brackets is true.

$$IS_\alpha(l, u: x) = (u-l) + \frac{2}{\alpha}(l-x)\mathbb{1}\{x<l\} + \frac{2}{\alpha}(x-u)\mathbb{1}\{x>u\} \quad (2)$$

The performance of the model is then evaluated using stratified k-fold cross-validation (Pedregosa et al., 2011) with the previous metric.

Results

Figures 7 to 11 show selected results of the proposed modelling approach using the whole house electrical energy consumption. Overconsumption (shown in red) and underconsumption (shown in blue) are identified for an hourly data span of a year. The original data was aggregated to an hour from a 15-minute interval.

In Figure 12, the interval score computed from equation (2) are presented for both the training and testing set. The aim is to obtain the lowest possible interval score. While different predictions can be compared against one another, there is no "threshold" value for the interval score aimed at this setting. This is a major practical problem, that will be addressed in the near future with more data and consumer interaction available.

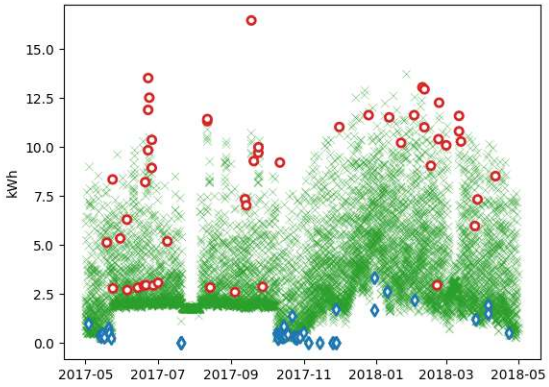

Figure 7. Identified over (red) and under (blue) electrical consumption for House ID 68

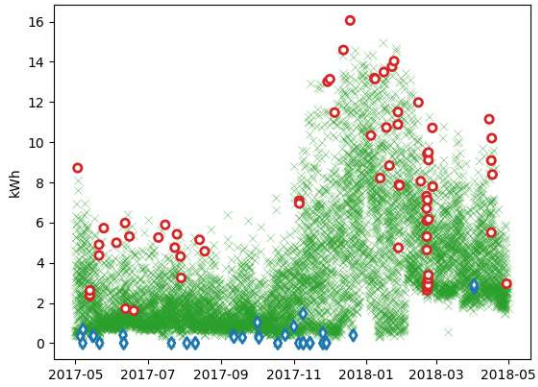

Figure 8. Identified over (red) and under (blue) electrical consumption for House ID 79

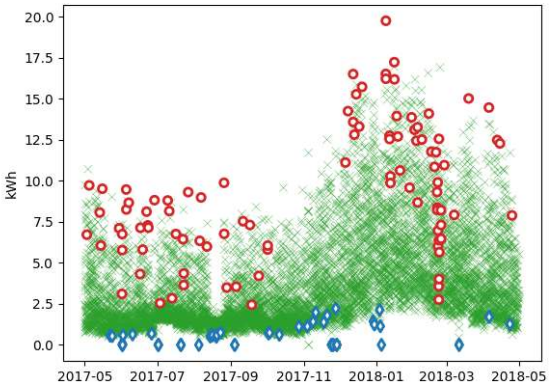

Figure 10. Identified over (red) and under (blue) electrical consumption for House ID 85

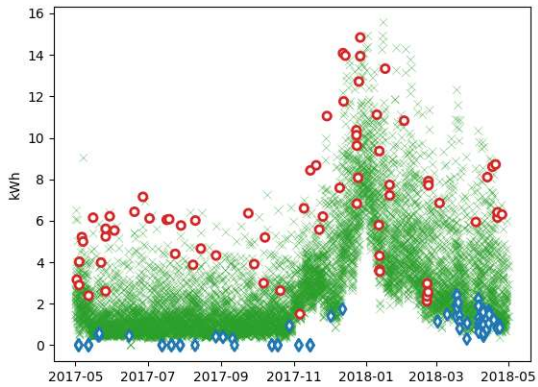

Figure 9. Identified over (red) and under (blue) electrical consumption for House ID 80

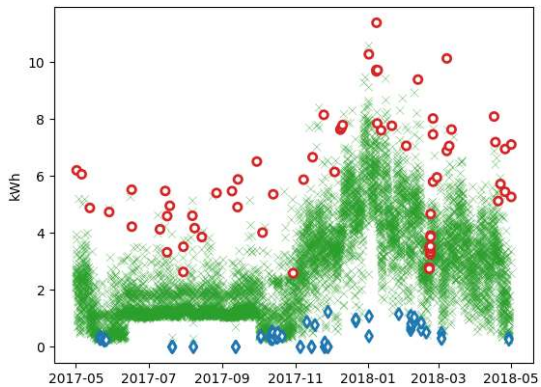

Figure 11. Identified over (red) and under (blue) electrical consumption for House ID 88

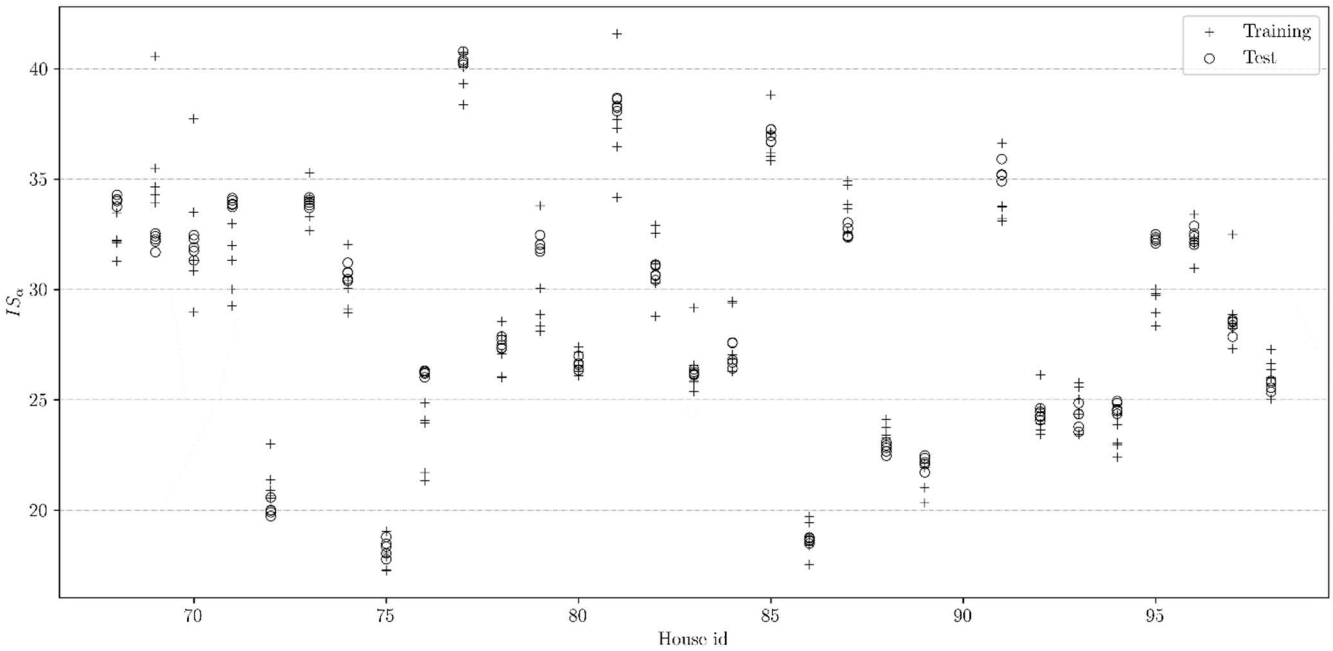

Figure 12. Interval score (IS_α) for the houses of the data set for folds of the training and test data.

The boxplots in Figure 13 present the dispersion of the interval score standard deviation for the training and testing datasets. In accordance with Figure 12, the performance metric of the model seems more stable for the testing data set than for the training dataset.

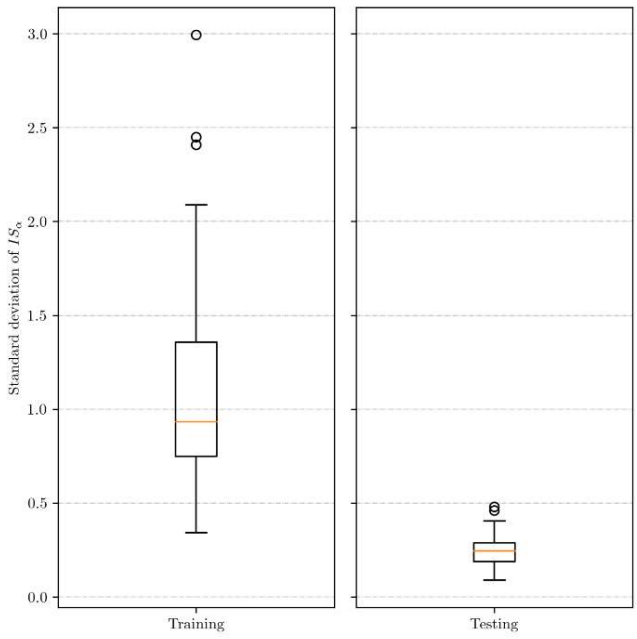

Figure 13. Boxplots of the standard deviation of the interval score obtained during stratified k-fold cross validation for the training and testing datasets.

This seems odd that the coverage of the prediction model is better for the unseen data by the algorithm than for its training data. The training and testing datasets were split respectively 80% and 20% of the available data. There might be an unforeseen interaction with the scoring method.

Discussion

Further investigation into the use of the interval score for FDD application seems to be necessary given the presented results. It is not a commonly used scoring method in the energy field and in the machine learning field from the author's experience. Another scoring metric might be more appropriate for this type of work and a broader literature review of scoring metric applied to interval prediction might prove useful.

Training the model over the whole dataset including all of the measured houses instead of a single one may provide a unique FDD model with acceptable accuracy. This is an important practical aspect because while the accuracy of the model might be better when trained for each customer, it is harder to manage as the number of models will grow with the user base. There's also less data used for model training and model use for a new customer is impossible. Customer segmentation and transfer learning might be interesting paths to explore to address this problem.

The base used in computing the HDH should be added as a parameter of the whole model in a subsequent version of the algorithm to further respect the PRISM method methodology. Instead of using the HDH, the PRISM model parameters and base temperature could be used as inputs of the neural network conjointly with the outdoor temperature.

While the machine learning method used in this paper yielded interesting results, it is quite simple. Many more types of machine learning algorithms could be paired with the customized loss function and might yield better results. Still, one advantage of using feed-forward instead of a time-series oriented type of neural network is that the effect of missing data is mitigated. Indeed, missing data points due to numerous possible reasons is a reality that must be considered in model development. The quantity of data needed to train a satisfactory model has not been studied in this paper. However, it is a very important practical aspect that must be addressed.

The presented results illustrate the model identified over and under consumption at the whole house level. Still, the way this information would be communicated to the consumer can differ from the plots shown in the previous section. A careful analysis of how and when to provide personalized feedback should be undertaken to maximize the effect of this information on the customer's energy consumption and peak power demand.

In Figure 8, Figure 10 and Figure 11, the increase of over consumption hours during March coincides with a sharp decline in HDH as shown in Figure 14. Further investigation is needed to identify if this is a result of common overheating during swing seasons or an artifact of the model. This type of strange behaviour is not unknown to the energy analysts as discrepancies in swing seasons for regression models and calibrated energy simulations are common.

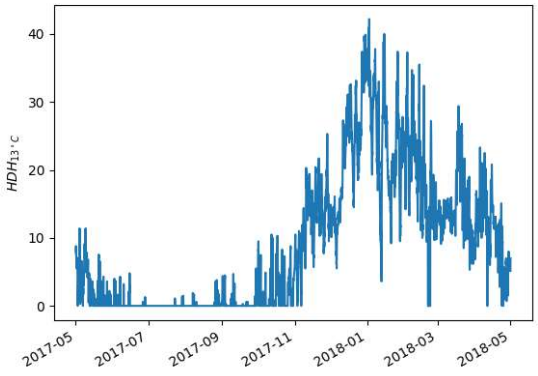

Figure 14 Heating Degree-Hours (HDH) in base 13°C for Trois-Rivières, QC

The underlying assumption in the proposed model is that the training data contains representative data of the residential building, energy systems and occupant's behaviour. A limitation of the proposed model is that a change in the building energy systems, number of occupants or any energy-related change might be attributed to a fault status. Hence, consumption change detection at another level (e.g. thermostat, major appliance) should be investigated to provide a retraining logic to the proposed model.

Model performance enhancements could arise by using directly other tools like TensorFlow (Abadi et al., 2016) and the nested cross-validation should be performed for algorithm hyperparameter tuning. In this paper, performance was not considered as a critical aspect. Thus only a few hyperparameters combinations were tested yielding acceptable performance for research use. Further data and validation are definitely required to evaluate the robustness of the proposed black box FDD model. Customer feedback will provide valuable information to tweak threshold levels and add to the proposed framework.

Finally, validation of such FDD algorithm is critical. Intentional faults could be used to evaluate the performance of the proposed modelling approach. The intentional abnormal behaviours or operation could provide an estimate for false positives and negatives yielded by the algorithm.

Conclusion

In this paper, a black box modelling approach is applied to whole building electrical energy data in a FDD framework using a quantile loss function. A feedforward neural network is trained on the whole house energy consumption 15-minute interval data, aggregated hourly. This modelling approach was developed using the sub-metered data from 30 houses in the vicinity of Trois-Rivières, Québec. While further work is needed to validate the proposed FDD method, the obtained results are promising, and the addition of other data sources and more advanced machine learning methods could be an improvement on the proposed method.

As other home and occupant characteristics become available, the number of inputs of the neural network can be adjusted. The same approach could be applied to the residence electric baseboard thermostats and other major electric loads. Furthermore, the interaction between identified faults at the thermostat level and identified faults at the whole house level could provide further insights to the customers.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., . . . Isard, M. (2016). *{TensorFlow}: A System for {Large-Scale} Machine Learning.* Paper presented at the 12th USENIX symposium on operating systems design and implementation (OSDI 16).

Dong, B., Prakash, V., Feng, F., & O'Neill, Z. (2019). A review of smart building sensing system for better indoor environment control. *Energy and Buildings, 199*, 29-46. doi:https://doi.org/10.1016/j.enbuild.2019.06.025

Ehrhardt-Martinez, K., Donnelly, K. A., & Laitner, S. (2010). *Advanced metering initiatives and residential feedback programs: a meta-review for household electricity-saving opportunities*.

Fels, M. F. (1986). PRISM: An introduction. *Energy and Buildings, 9*(1-2), 5-18. doi:10.1016/0378-7788(86)90003-4

Fournier, M., Leduc, M.-A., Sansregret, S., & Poulin, A. (2018). *Making the Connection: Testing Line-Voltage Communicating Thermostats for Baseboard Heaters in DR and EE Experiments*.

Frank, S., Jin, X., Studer, D., & Farthing, A. (2018). Assessing barriers and research challenges for automated fault detection and diagnosis technology for small commercial buildings in the United States. *Renewable and Sustainable Energy Reviews, 98*, 489-499. doi:https://doi.org/10.1016/j.rser.2018.08.046

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association, 102*(477), 359-378.

Guo, F., Rogers, A. P., & Rasmussen, B. P. (2021). Multivariate fault detection for residential HVAC systems using cloud-based thermostat data, part I: Methodology. *Science and Technology for the Built Environment*. doi:10.1080/23744731.2021.2005375

Katipamula, S., & Brambley, M. R. (2005). Review Article: Methods for Fault Detection, Diagnostics, and Prognostics for Building Systems—A Review, Part I. *HVAC&R Research, 11*(1), 3-25. doi:10.1080/10789669.2005.10391123

Khani, S. M. R., Haghighat, F., Panchabikesan, K., & Ashouri, M. (2021). Extracting energy-related knowledge from mining occupants' behavioral data in residential buildings. *Journal of Building Engineering, 39*, 102319. doi:https://doi.org/10.1016/j.jobe.2021.102319

Kim, W., & Katipamula, S. (2018). A review of fault detection and diagnostics methods for building systems. *Science and Technology for the Built Environment, 24*(1), 3-21. doi:10.1080/23744731.2017.1318008

Mitchell, J. W., & Braun, J. E. (2012). *Principles of heating, ventilation, and air conditioning in buildings*: John Wiley & Sons.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., . . . Dubourg, V. (2011). Scikit-learn: Machine learning in Python. *the Journal of machine Learning research, 12*, 2825-2830.

Rogers, A. P., Guo, F., & Rasmussen, B. P. (2019). A review of fault detection and diagnosis methods for residential air conditioning systems. *Building and Environment, 161*, 106236. doi:https://doi.org/10.1016/j.buildenv.2019.106236

Whitmore, J., & Pineau, P.-O. (2021). État de l'énergie au Québec 2021. In. Chaire de gestion du secteur de l'énergie: HEC Montréal.

Zhao, Y., Li, T., Zhang, X., & Zhang, C. (2019). Artificial intelligence-based fault detection and diagnosis methods for building energy systems: Advantages, challenges and the future. *Renewable and Sustainable Energy Reviews, 109*, 85-101. doi:10.1016/j.rser.2019.04.021