Spatio-Temporal Context Anomaly Detection for Residential Power Consumption

Nurhidayat Sisworahardjo and Akram A. Saad

The University of Tennessee at Chattanooga, Tennessee, United States
Nur-Sisworahardjo@utc.edu, yfv586@mocs.utc.edu

Abstract: Monitoring energy consumption and diagnosing abnormal behavior will enable utilities to introduce strategies to improve system resiliency, stability, and to meet energy efficiency targets. The deployment of advanced metering infrastructure (AMI) enables utilities to collect various raw data from its customers and networks. This paper presents contextual anomaly detection algorithm to detect irregular power consumption and visualize anomaly scores using unsupervised learning algorithm and temporal context generated from meter readings. The proposed algorithm computes an anomaly score for each user by considering historical consumption data. The anomaly score for a user is then adjusted by analyzing other contextual variables such as seasonal variation day of the week and other users with the same historical pattern. The implementation of real-world data set provided by power utility company shows a high performance of the proposed algorithm.

Index Terms: contextual anomaly detection, unsupervised learning algorithm, temporal context

1. Introduction

According to the U.S. Green Building Council [1], buildings account for 39% of CO_2 emission, which most of this emission is caused by electricity consumption. The World Energy Outlook report states that residential and commercial buildings have the most undiscovered energy efficiency possibility [2]. One way to achieve building energy efficiency is to monitor energy usage, detect, and diagnose abnormal consumption behavior. There has been rapid growth in electrical energy monitoring infrastructure. For example, more than 65 million smart meters have been installed in the United States by 2015 [3]. The abnormal energy usage monitoring (anomaly detection) requires prior knowledge of typical energy consumption. Therefore, data-analytics techniques are applied to the readings gathered from smart meters to identify normal consumption patterns. This will enable the detection of patterns that do not conform to recognizable patterns.

An anomalous condition is defined as abnormal power consumption usage whereas an anomalous state is defined as a deviation from the normal consumption. Depending on their nature, anomalies can be characterized as point, contextual, or collective anomalies [4]. Point anomalies refer to data instances that appear to be anomalous to all other instances in the dataset. For example, a sudden increase in energy consumption for a building on a given day might be anomalous compared to previously recorded daily consumptions in the related data set. Collective anomalies refer to a set of related data instances that have different behavior than the entire data set. However, these instances when considered individually are not anomalies. As energy consumption of a building might be normal considering previously recorded instances. Nevertheless, considering these instances over a certain window of time might represent a collective anomaly. Finally, contextual anomalies refer to data instances that might appear anomalous in a given context but normal in a different one. Contextual anomalies take contextual attributes into account, for instance, in residential buildings high consumption might be anomalous during the weekday, but not on the weekends.

The intuition behind anomaly detection is to identify deviation from normal consumption pattern [5]–[7]. Yet, this intuition suffers from the lack of global semantic view; as consumption can vary based on time of the day, the day of the week, and many other contextual variables such season of the year and day of the year. This paper proposes a method that uses context information which is attached to the meter identity and can be obtained from meter readings, i.e., timestamp and location. Temporal context information such as daily and weekly cycles, extracted from time stamps are used to refine the anomaly detection selectivity and accuracy by selecting only the relevant historical data for baseline estimation. This captures the effect of factors such as high and low consumption periods (for residential buildings), operation hours (for commercial buildings), and seasonal changes (change in heating/cooling loads). Finally, correlation factor indicates the extent to which customer's power consumption time series are correlated. It is computed based on daily power consumption time series between individual meters. This factor is used for adjusting the anomaly score to account for unknown context variables that influence correlated consumption patterns in the same way such as holidays or special events.

The remaining sections of the paper are organized as follows: Section 2 describes related anomaly detection work which followed by a deeper understanding of energy consumption in Section 3. Section 4 outlines and discusses the proposed approach in this research. Section 5 presents a case study and discusses the results, and finally, Section 6 presents the conclusion of this work.

2. Power Anomaly Detection - Related Work

Anomaly detection is an important topic that has been investigated by researchers in various endeavors such as fraud detection, intrusion detection, and industrial damage detection. In this section, an analysis of related work in power anomaly detection is presented.

Existing studies of anomaly detection methods in the energy consumption focused on identifying point anomalies [5]–[8]. Bellala et al. [5] used an unsupervised anomaly detection technique based on clustering a low-dimensional dissimilarity matrix. It based on the assumption that data points which represent normal behavior will form tight clusters while anomalous points will lie outside these clusters. Chou and Telaga [6] proposed a two-stage method to detect anomalous consumption within a building space. One-step-ahead daily consumption is predicted by using ARIMA model. Anomalies are flagged if the difference between real and predicted value is above a certain threshold. Janetzko et al. [7] proposed a time-weighted prediction using historical power consumption data to identify the anomalies. Chen et al. [8] transformed time series energy into symbol sequence, then clustering algorithm was used to detect outlier energy patterns.

Other researchers have adopted contextual information to detect anomalous behavior within a certain context. Zhang et al. [9] proposed several methods to detect anomalous days including regression, clustering, and entropy. They used temperature as an extra contextual information to increase accuracy. However, the model is static and doesn't adapt the change in customer's behavior. Araya et al. [10] proposed sliding window framework to integrate historical sensors data and contextual features. They created 17 and 26 features model using *autoencoder* to identify consumption patterns for a building. Gupta et al. [11] proposed a framework for monitoring and detecting anomalies by integrating the information from system logs and time series measurement data.

In this work, a hybrid Clustering technique was used to extract context and metric patterns and anomalies are detected based on these context and metric patterns. The contextual information used are obtained from meter readings. Other contextual information such as (seasonal changes and temperature) can be integrated using correlation factor information. The algorithm is flexible regarding identifying new contextual features such holidays. While the approach proposed in this paper is similar to the one used in [5], it exploits contextual information as well as historical consumption to effectively identify anomalies as in the [11] and uses a novel technique to identify anomalous patterns.

3. Energy Consumption – A Deeper Look

A. Energy Consumption Visualization

Before describing the proposed methodology for anomaly detection in detail, let's take a first look at one customer energy consumption provided in Figure 1.
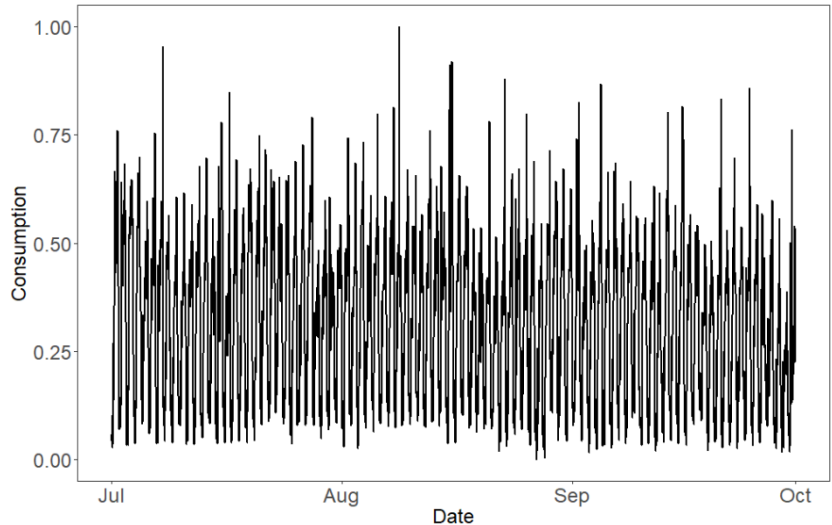

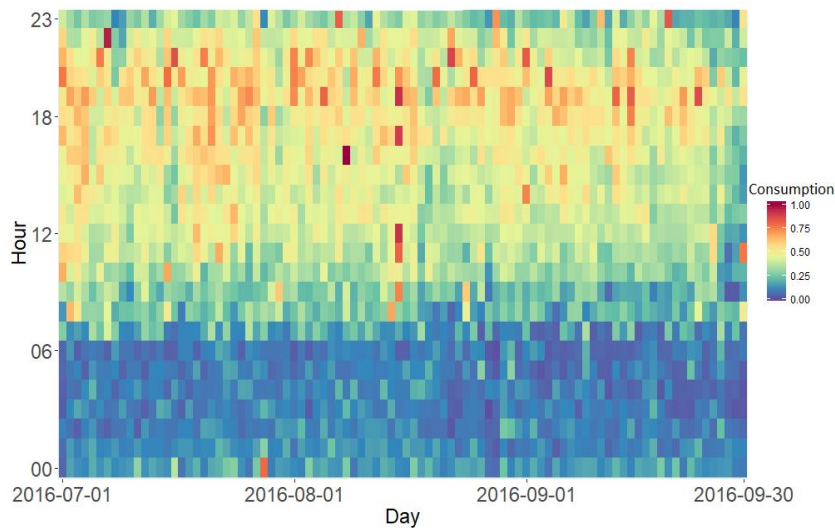Figure 1. Energy consumption visualization with line charts

Figure 2. Energy consumption visualization with pixel-based technique

The line chart given in Figure 1 shows the normalized energy profile for one of the customers. The line charts are the most common visualization technique of time series. However, with the fine-grained consumption readings provided by smart meters, this method is no longer effective to uncover the underlying patterns in energy consumption.

In this paper, a unique visualization technique is employed. It is built upon the work presented in [7]. The pixel-based technique is used to visually encode numerical values into colors. It gives an overview of large time span and provides the possibility to recognize patterns or exceptions in the power consumption time series and to analyze the anomaly scores. Figure 2 shows the

same energy consumption visualized with pixel-based technique. This technique is showing that this customer is exhibiting high consumption period from 7 A.M to 23 P.M, and low consumption period from 23 P.M to 7 A.M.

B. Temporal Context

Contextual variables such as occupancy and appliance usage can influence energy consumption considerably. To be able to analyze such contextual variables, extra sensing infrastructure should be used. Nonetheless, a human activity typically follows regular temporal. For instance, appliance usage is highly correlated to temporal context [12]. Therefore, in this paper, occupancy and appliance usage are indicated by temporal context sets as follows:

- High consumption period: this set contains readings during weekdays from (7 a.m. to 12 a.m.).
- Low consumption period: this set contains reading during weekdays from (12 a.m. to 7 a.m.).
- Weekend: this set contains reading taken during the weekend when residential buildings are likely to have increased occupancy, and commercial premises are unoccupied.

C. Geographic Area Effect

The geographic area is defined as the set of buildings that are expected to be influenced similarly by the same context variables. Figure 3. represents six customers from the same geographic area. For each customer, the area marked in red represents a deviation from the normal pattern. Even so, in overall, this pattern represents a trend among different customers, hence, flagging it as an anomaly is not necessarily correct. For example, during holidays, when residential buildings have increased occupancy, that will result in higher energy consumption. Another example is temperature variation for a specific area can influence energy consumption considerably.
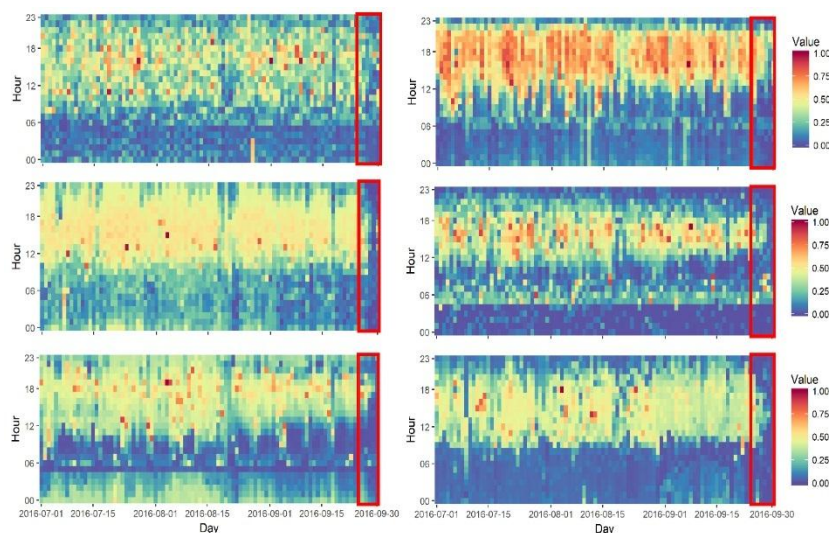

Figure 3. Energy consumption of six customers from the same geographic area

4. Proposed Methodology

In this section, the proposed approach for computing anomaly scores for energy consumption data is discussed. The anomaly score is scaled in the range of 0 to 1 (low to high severity), which implies the likelihood of a given data instance to be anomalous. The approach consists of three steps for computing this score:

1. *Data preprocessing*: Split meter reading into disjoint sets based on defined temporal context information.

2. *Initial anomaly score*: Apply anomaly detection for each context set separately to compute an initial anomaly score.
3. *Correlation factor adjustment*: Adjust the initial anomaly score for each meter based on correlation information.

A. Data preprocessing

The data from smart meters usually suffer from incompleteness due to momentary connection failures. On average, the data missing represents less than 1.5% of the whole dataset. For each individual customer, if the missing value exceeded 3%, then the data set of that customer is discarded. The rest of the dataset is considered be representative since values are missing at random. Linear interpolation was used to generate missing values. Once the missing values were generated, the dataset is then normalized in an attempt to avoid suppressing of small values. The data is normalized by rescaling the meter reading to range in [0 1].

The completed, normalized data set is then split based on temporal context. Each meter reading is classified into one of three temporal context sets as defined in Section 3.

B. Initial anomaly score

For each given temporal context subset, anomaly detection algorithm is applied, this step is referred to as initial anomaly score, which is summarized in Algorithm 1.

Algorithm 1: Initial anomaly score

Input: M meter readings spanning D days.

Output: $A_{c,d}^m$: initial normalized anomaly score for each instance, $i_{c,d}^m$.

1. Split meter readings (M) into C disjoint temporal context
2. **for all** the $c \in C$ one temporal context **do**
3. Compute dissimilarity matrix Δ_{cm} for all pair of instances $i_{c,i}^m$, $i_{c,j}^m$ using Euclidean distance measure.
4. Reduce the dimensionality of Δ_{cm} to 2 dimensions using Multi dimension scaling (MDS) algorithm.
5. Use PAM partition around medoid algorithm to cluster Δ_c^m into O clusters, compute population of each c and save in S_c^m.
6. **for all** the instances $i_{c,d}^m \in D$ **do**
7. Compute the O-sized vector of Euclidean distances between each cluster medoid and data instance $i_{c,d}^m$.
8. The initial anomaly score $A_{c,d}^{m*}$ is the dot product between O-sized vector and S_c^m cluster size.
9. Compute the initial normalized anomaly score, $A_{c,d}^m$, for each instance $i_{c,d}^m$:
$$A_{c,d}^m = \frac{A_{c,d}^{m*}}{\max A_{c,d}^{m*}}$$

Each meter temporal context subset represents a separate input to the algorithm, (e.g. high consumption period), the algorithm outputs is an anomaly score associated with each day in the time series. In this stage, only historical data is required to calculate the initial anomaly score.

Each series of consecutive meter readings within a single context set is denoted as an instance, $i_{c,d}^m$, where $m \in \{1, ..., M\}$ is a meter, $c \in \{1, ..., C\}$ is an additional contextual information, and $d \in \{1, ..., D\}$ is the index of the intance, respectively. In this paper, three classification sets are used, those are high, low and weekend periods. For example, all meter readings between 7 a.m. to 12 a.m. will form one instance in high consumption period set.

The dissimilarity matrix Δ_m^c is calculated for all pairs of instances $i_{c,i}^m$ and $i_{c,j}^m$ which are in the same temporal class. With the dissimilarity matrix, a low-dimensional matrix is obtained using multidimensional scaling (MDS) [12]. This step is important since clustering algorithms perform better in lower dimensions. The next step is to perform clustering algorithm on the low-dimensional dissimilarity matrix. An unsupervised *k*-medoid clustering algorithm, in particularly, partitioning around medoids (PAM) is used to for this purpose [13]. The PAM

algorithm also estimates the optimum number of clusters O in the dissimilarity matrix Δ_m^c. Unlike k-mean which is sensitive to outlier since mean easily influenced by extreme values, PAM uses an actual data point in the cluster to represent it, hence, its more robust to noise and outliers. Figure 4 shows a 2-dimensional representation of energy consumption data, note that normal patterns form tight clusters, while outliers lie outside these clusters.

The initial anomaly score is then calculated for each instance $i_{c,d}^m$ based on the Euclidean distance between this instance and clusters medoids. The set of Euclidean distance between each medoid and an instance $i_{c,d}^m$ is stored in O-size vector. For each instance in the same context, the initial anomaly score $A_{c,d}^{m*}$ is the dot product between the O-sized vector and cluster size S_c^m. Finally, the initial anomaly score is normalized to range in [0, 1] to determine the level of severity of each anomaly.
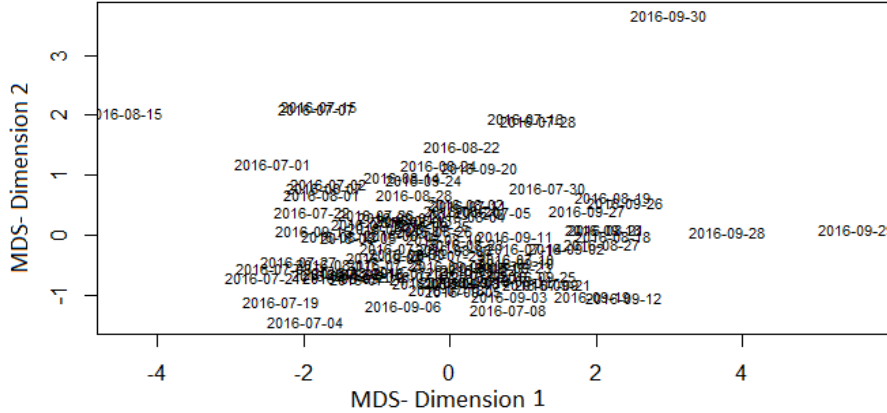

Figure 4. Two-dimensional representation of energy consumption data.

In the existence of multiple context classifications, in this paper, the number of context classification sets is $C = 3$. The anomaly score for each context can then be merged to give A^m, the initial anomaly score for meter m. The next step is to feed the initial anomaly scores to correlation factor adjustment algorithm.

C. *Correlation factor adjustment*

To incorporate geographic area effect, baseline correlation factor between meters is used. It reflects how customer's patterns are related and used to adjust the initial anomaly score.

The adjustment factor $a_{c,d}^m$ for each instance $i_{c,d}^m$ is computed based on the correlation of power consumption between individual meters within the same geographic area and the initial normalized anomaly score $A_{c,d}^m$.

Algorithm 2: Correlation factor Adjustment

Input: $A_{c,d}^m$ initial anomaly scores for each m meter.
    daily correlation matrix list R_d spanning D days.

Output: $\hat{A}_{c,d}^m$: adjusted anomaly score.

1. **for all** the day $d \in D$ **do**
2. **for all** the meter $m_1 \in M$ **do**
3. Select $A_{c,d}^{m_1}$, that corresponds to meter m_1.
4. **for all** the meter $m_2 \in M$ **do**

$$a_{c,d}^{m_1} = \sum_{m_1 \neq m_2} R_d(m_1, m_2) * A_{c,d}^{m_2}$$

$$\hat{A}_{c,d}^{m_1} = A_{c,d}^{m_1} - w \times \text{mean}\left(a_{c,d}^{m_1}\right)$$

As shown in Algorithm 2, for each instance $i_{c,d}^m$, the adjustment factor $a_{c,d}^m$ is defined as the sum of the correlation-weighted initial anomaly scores from the same instance, for all meters within the same geographic area of meter m. Then, the adjusted anomaly score $\hat{A}_{c,d}^m$ is calculated as the difference between the initial normalized anomaly score $A_{c,d}^m$ and the mean of the weighted adjustment factor $w \times a_{c,d}^m$. The optimal value of w can be calculated experimentally or provided by domain experts.

5. Case Study

The dataset contains energy meter readings from 100 customers from the same distribution circuit in the distribution network. Hence, this group of customers responds in the same way to contextual variables such as temperature variations and special event days. The data was collected from July 1st to September 30th, 2016 in 15 minutes interval. Figure 5 represents the baseline correlation between customers for the period of collection. Figure 5 shows a high and positive correlation between these customers in the same geographic area. This indicates that these customers behave the same way to different seasonal changes and special event days. Therefore, this additional correlation data from other meters within the same geographic area should help the anomaly detection algorithm to distinguish between patterns that are reported from multiple meters which reflect unseen contextual variables (such as holidays and special event days) and unique anomalous behavior for a single meter.
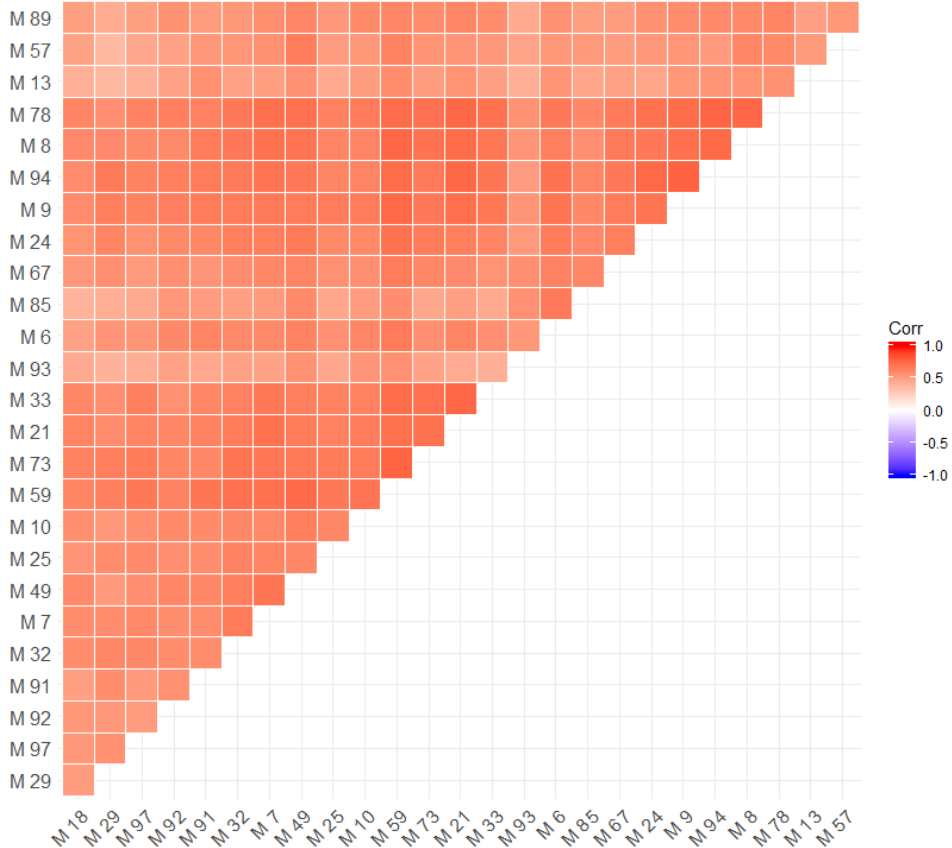

Figure 5. Correlation between 26 customers.

A. Evaluation Criteria

The proposed algorithm identifies anomalies through analyzing historical data, therefore, the analysis is focusing on identifying particular injected events that represent abnormal energy usage. For example of such events is heat pump inadvertently set to an emergency mode of residential buildings or operation of HVAC loads in the weekends for commercial buildings.

Table 1. Injected anomalous consumption patterns.

Case #	Anomaly Type	# of instances	Duration
1	High consumption on low consumption period	5 (red rectangular in Figure 6)	3-7 hours
2	High consumption on weekend days	3 (blue rect. & arrow in Figure 6)	24 hours
3	High consumption all day	3 (green rect. & arrow in Figure 6)	24 hours

The evaluation criteria for the proposed algorithm is carried out in twofold:
- Compare the performance of an existing algorithm Bellala et al. [5] that shares a similar idea with the proposed initial anomaly score algorithm.
- Manually inject anomalous consumption patterns into the raw power readings and marked it as a ground truth for the evaluation. Table 1 summarizes the injected anomaly events.

B. Experimental Results

The experimental results are focusing on analyzing the injected anomalous consumption with respect to temporal context.

A subset of 26 smart meter reading was used for experimental results. To indicate the importance of using context and correlation factor, different versions of the proposed algorithm as well as the existing algorithm were compared.

a. Anomaly score with no temporal context (NTC): This method is used to demonstrate the importance of using temporal context information to increase the accuracy of the anomaly detection method. This method is described in Algorithm 1. However, meter readings were not classified into temporal context.
b. Anomaly score with temporal context (WTC): This method represents the proposed method where meter readings were classified into three temporal context sets.
c. Adjusted anomaly score with temporal context (ATC): This method represents the modified version of WTC where correlation factor was used to adjust the anomaly score as in Algorithm 2.
d. Cluster-based algorithm (CA): The method was proposed in [5], where the outputs are the probability of a day being anomalous.

Figure 6 represents meter readings for one of the customers with different anomaly scores computed by different anomaly detection methods. Considering severe anomaly score is above 0.75, among the 11 injected anomalies, NTC assigned a high score for 4 out of 11 injected anomalies. Meanwhile, WTC succeeded to assign a high score for 9 out of 11 injected anomalies. WTC identified the other two injected anomalies, however, they were not flagged as severe anomalies. With correlation information was used in ATC to adjust the anomaly score, out of 11 injected anomalies, 7 were identified as unique anomalies. These unique anomalies have clear distinction compare to the pattern which shared among the customers in the same geographic area. On the other hand, 4 out of 11 injected anomalies were considered normal when compared to other customers' consumption. Lastly, the cluster-based algorithm (CA) succeeded to assign a high score for the injected anomalies, however, it also assigned a high score to days where consumption seems to be normal.
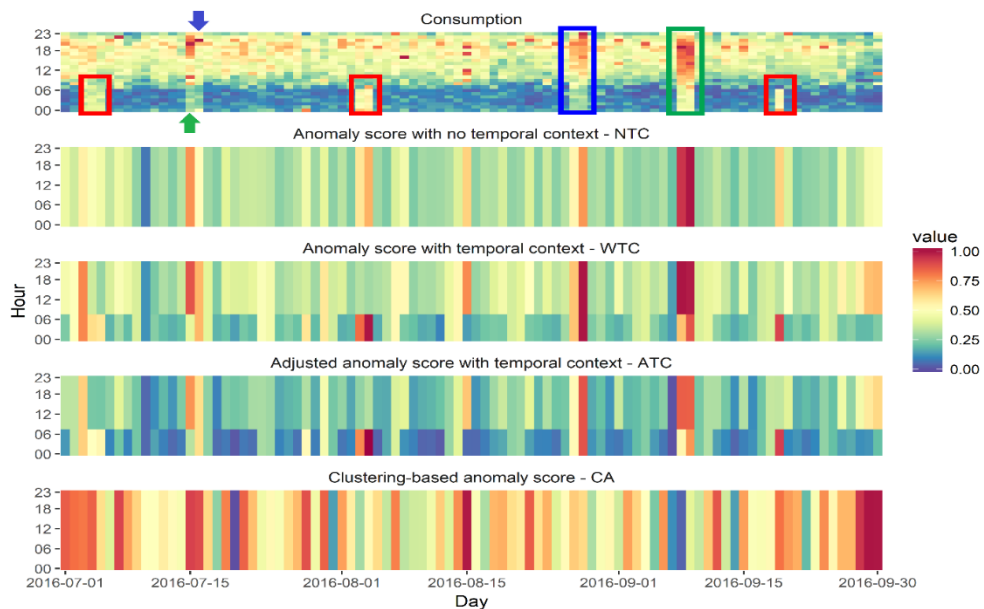
Figure 6. Comparison between different anomaly detection methods

Adjusted anomaly score (ATC) algorithm, although seems to detect less number of injected anomalies compare to WTC's, it provides more robust results and fewer false positive, because it relies not only on customer's consumption data like WTC but also it incorporates additional external information to identify anomalies. The value of weighting factor w used in ATC was calculated experimentally. The optimal value of w must result in no negative anomalies, and due to high correlation between customers in the sub set chosen, this value was determined to be 0.5.

6. Conclusion

In this paper, a contextual anomaly detection algorithm has been presented. The algorithm extracts the context information from the meter readings and combines other contextual variables using correlation factor to better identify anomalous consumption behavior. Experimental results have shown that the proposed algorithm successfully identified the injected anomaly cases and thereby outperform the existing method.

Future work will focus on integrating other contextual variables such as temperature variations to the algorithm and identify new ways to correlate between customers such as customer clustering based on consumption.

7. Acknowledgement

8. References

[1]. K. Fowler, E. Rauch, J. Henderson, and A. Kora, "Re-assessing green building performance: A post occupancy evaluation of 22 GSA buildings," *Re-assessing green building performance: A post occupancy evaluation of 22 GSA buildings*, 2010.

[2]. M. Van der Hoeven, "World Energy Outlook 2012," *International Energy Agency: Tokyo, Japan*, 2013.

[3]. U.S. Department of Energy, "Smart Grid Investment Grant Program Final Report," Dec-2016. [Online]. Available: https://www.smartgrid.gov/ files/Final_SGIG_Report_ 20161220.pdf. [Accessed: 30-Jan-2017].

[4]. V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, Jul. 2009.

[5]. G. Bellala, M. Marwah, M. Arlitt, G. Lyon, and C. E. Bash, "Towards an understanding of campus-scale power consumption," in *Proceedings of the Third ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings - BuildSys' ' '11*, New York, New York, USA, 2011, p. 73.

[6]. J.-S. Chou and A. S. Telaga, "Real-time detection of anomalous power consumption," *Renewable and Sustainable Energy Reviews*, vol. 33, pp. 400–411, May 2014.

[7]. H. Janetzko, F. Stoffel, S. Mittelstädt, and D. A. Keim, "Anomaly detection for visual analytics of power consumption data," *Comput Graph*, vol. 38, pp. 27–37, Feb. 2014.

[8]. C. Chen, D. Cook, C. Chen, and D. Cook, "Energy Outlier Detection in Smart Environments," *Artificial Intelligence and Smarter Living*, vol. 11, 2011.

[9]. Y. Zhang, W. Chen, and J. Black, "Anomaly detection in premise energy consumption data," in *2011 IEEE Power and Energy Society General Meeting*, 2011, pp. 1–8.

[10]. D. Araya, K. Grolinger, H. ElYamany, M. Capretz, and G. Bitsuamlak, "Collective contextual anomaly detection framework for smart buildings," *Neural Networks (IJCNN), 2016 International Joint Conference on*, 2016.

[11]. M. Gupta, A. Sharma, H. Chen, and G. Jiang, "Context-aware time series anomaly detection for complex systems," *Workshop Notes*, 2013.

[12]. Z. Ning and D. Kirschen, "Preliminary analysis of high resolution domestic load data," *Part of Supergen Flexnet Project, The University of Manchester*, 2010.

[13]. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York: Wiley, 2001.

Nurhidayat Sisworahardjo (S'94, M'05) received his B.S. and M.S. from Institut Teknologi Bandung, Bandung, Indonesia and his Ph.D. from the University of Alabama all in electrical engineering. He is an Associate Professor of Electrical Engineering at the University of Tennessee at Chattanooga. He has been conducting research in the areas of data-analytics in power distribution, generating unit asset valuation, load forecasting, distributed generation (fuel cells & microturbine) modeling, distributed generation penetration, and post-disturbance network reconfiguration. His general research interest including smart grid, distributed generation, renewable/alternative energy, power systems operation, optimization, simulation, and planning.
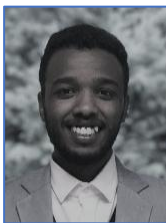
Akram A. Saad received his B.SC in Electrical Engineering from University of Khartoum, Sudan, in 2014. He is a candidate to receive an M.SC in Electrical Engineering in May 2018 from University of Tennessee at Chattanooga- USA, where he works as a graduate assistant at the Department of Electrical Engineering.
Mr. Saad current interest includes data analytics application in the context of smart grids, power system stability analysis, and protective relaying.