

Detecting Anomalous Electrical Appliance Behavior based on Motif Transition Likelihood Matrices

Andreas Reinhardt
Technische Universität Clausthal
Clausthal-Zellerfeld, Germany
reinhardt@ieee.org

Delphine Reinhardt
University of Bonn & Fraunhofer FKIE
Bonn, Germany
delphine.reinhardt@cs.uni-bonn.de

Abstract—With the rise of the Internet of Things, innumerable sensors and actuators are expected to find their way into homes, office spaces, and beyond. Electric power metering equipment, mostly present in the form of smart meters and smart plugs, can thus be anticipated to be installed widely. While smart plugs (i.e., individual power monitors attachable to wall outlets) primarily cater to the user comfort by enabling legacy devices to be controlled remotely, their power measurement capability also serves as an enabler for novel context-based services. We realize one such functionality in this paper, namely the recognition of unexpected behavior and appliance faults based on observed power consumption sequences. To this end, we present a system that extracts characteristic power consumption transitions and their temporal dependencies from previously collected measurements. When queried with a power data sequence collected from the appliance under consideration, it returns the likelihood that the collected power data indicate normal behavior of this appliance. We evaluate our system using real-world appliance-level consumption data. Our results confirm the individual nature of consumption patterns and show that the system can reliably detect errors introduced in the data within close temporal proximity to their time of occurrence.

I. INTRODUCTION

Sensing equipment to monitor electrical power consumption can be expected to become ubiquitous, both due to legislative reasons (e.g., the European Union Energy Services Directive [1], enforcing the installation of smart meters) and the benefits they provide to utilities (e.g., reduced labor cost for taking readings). In addition to smart meters that capture the aggregate energy demand of a home, wireless power monitoring devices that attach to wall outlets (often referred to as *smart plugs*) are also gaining popularity among users who want to better understand their power demand. Consequently, a large volume of power and/or energy consumption data from residential and industrial buildings can be anticipated to be collected in the future.

Many services to enrich the users' understanding of their energy consumption have been proposed in current research. For example, the disaggregation of smart meter data is an active research domain which assists users in determining the energy consumptions of individual loads from the aggregate [2–4]. Once appliance-level data is available for at least parts of a building (e.g., through disaggregation or the use of smart plugs), the underlying appliance types can be detected [5], energy wastage be identified [6], or the expected consumption for the entire dwelling be extrapolated [7]. While these energy-based services can be leveraged to increase user awareness and/or comfort, the vast majority of them do not allow for the analysis of real-

time data streams, making the realization of services like the live detection of anomalous consumption behavior impossible.

In this manuscript we propose a novel data processing method to detect anomalies in appliance-level power consumption data. It relies on the unique combination of several data processing steps, catering to both detection performance and accuracy. Initially, all collected power consumption data are converted to their *symbolic representations*. Secondly, consecutive symbols are combined into *motifs*, the primary data entities used by our approach. At the heart of our solution is the 3-dimensional *motif transition likelihood matrix* (MTLM), which captures the likeliness for the transition between any two motifs at a given temporal distance. When queried for a motif transition observed in real-time input data, the MTLM returns the corresponding likelihood and thus indicates whether an appliance is in a normal state of operation.

We introduce an overview of our system architecture and its fundamental design considerations in Sec. II, followed by providing more details on the core feature of our approach, the MTLM, in Sec. III. We evaluate the system's capability to detect erroneous appliance behavior in Sec. IV and summarize related work in Sec. V. We conclude this paper in Sec. VI.


II. SYSTEM OVERVIEW AND IMPLEMENTATION DETAILS

The core contribution of this paper is a system design to detect anomalies in electrical power consumption data streams. We derive the system's major design drivers from a real-world example in this section and introduce its overall architecture as well as highlighting how its components are realized in practice.


A. Fundamental Observations and System Overview

Many electrical appliances show recurring power consumption patterns across their operation cycles. Let us, e.g., consider a dishwasher appliance. In Fig. 1a, we have plotted the device's power demand for seven of its operation cycles. A high degree of similarity between the traces is apparent shortly after the beginning of the appliance's operation. However, it can also be observed that the device's individual operational phases differ in their durations across the seven dishwashing cycles and thus the deviations between the individual curves grow larger as time progresses. This can be predominantly attributed to variations in the appliance's environmental conditions (e.g., different inlet water temperatures or varying numbers of dishes in the appliance).

Firstly, this observation serves as evidence for the existence of repetitive temporal patterns in appliance's power consumptions.



(a) Superimposed power consumption traces of 7 dishwasher operation cycles



(b) Visualization of the repetitive nature of the dishwasher's cycle through temporally aligning the different phases (dashed lines mark re-alignment points)

Fig. 1. Repetitive power consumption across operation cycles (data from [8])

Secondly, it also highlights that slight temporal deviations are common and the similarity of appliance power consumption traces decreases for longer analysis windows. However, through temporal re-alignment of the segments of the power consumption activity (at offsets where major changes to the power consumption have occurred), a high amplitude coherence can be regained for the individual phases, as can be observed in Fig. 1b. Consequently, we exploit this recurring nature of similar power consumption patterns across appliance operation cycles (although possibly with slight differences in their durations) to detect anomalous consumption. The data flow of our conceived system has been purposefully designed to this end; it is visualized in Fig. 2, and we explain it in more detail as follows.

B. Dimensionality Reduction

One major obstacle in the development of energy-based functionalities is the computational power needed to process collected power consumption data. Smart meters and smart plugs often feature amplitudinal resolutions of fractions of a watt, and most hardware platforms can collect readings once per second or even more often. In conjunction with the fact that many appliance types feature operation cycles of several minutes or hours, enormous amounts of training data may be required to extract characteristic templates for an appliance's operation. Large computing capabilities are required to process this large corpus of high-resolution data, possibly even rendering the processing of streaming data impossible in real-time.

We have thus decided in favor of operating on symbolic representations instead of raw data, because they speed up all subsequent processing and simultaneously reduce the required memory overhead drastically. Dimensionality reduction techniques are traditionally being applied to extract key features from large volumes of input data [9, 10]. However, the specifics of the scenario under consideration render some widely used dimensionality reduction techniques (such as Fourier or Wavelet transforms) inapplicable because they eliminate temporal dependencies from the data. As the temporal sequence of events is essential to detect anomalies in power consumption patterns, we have decided in favor of an approach that alters the values of the

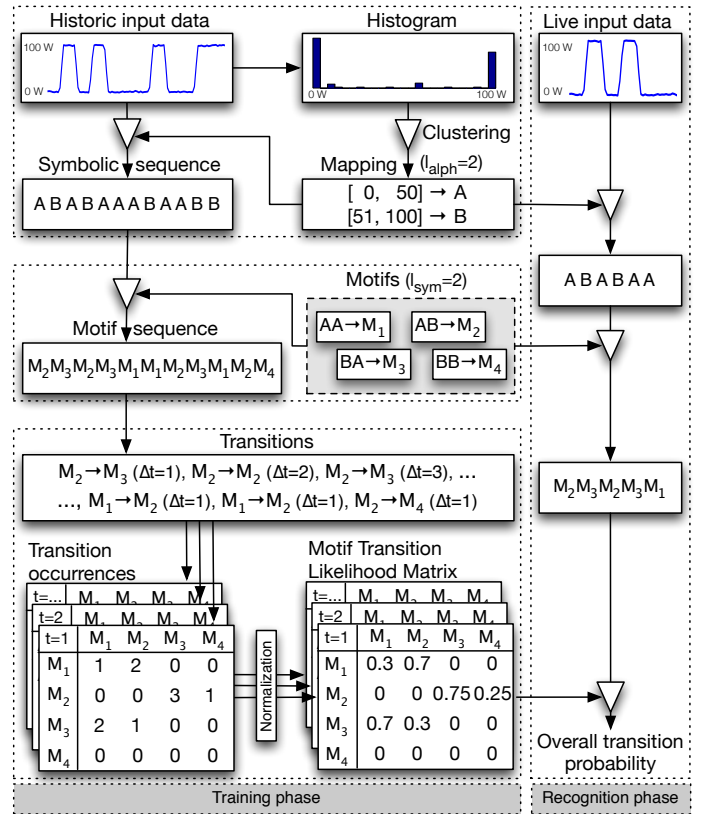


Fig. 2. Overall data flow of the system's training and recognition phase

signal while retaining their temporal order. More precisely, we convert the input data to their *symbolic representations*, i.e., we map the n_{pow} distinct values returned by the power consumption sensor to a set of l_{alph} symbols (with $l_{alph} \ll n_{pow}$).

To ensure that the symbol mapping leads to an accurate representation of the input time series, we compute the histogram across previously collected data. We then extract clusters comprising the most frequently occurring values from the histogram. Similar to existing work that relies on symbolic sequences of power data [11], we apply the DBSCAN algorithm [12] to the normalized histogram and consecutively allocate symbols to all resulting clusters as well as all value ranges between the determined cluster boundaries. By mapping all power consumption values between the determined boundaries to individual symbols, the data dimensionality is drastically reduced and measurement noise is furthermore attenuated as a side effect. The resulting cluster boundaries when computing and clustering the histogram of the dishwasher traces from Fig. 1 are shown as dashed lines in Fig. 3a.

The actual conversion of raw time series input data to their symbolic representation is realized by means of a look-up table derived from the extracted symbol mappings. The mapping table for the dishwasher example is shown in Fig. 3b. Note that the symbol alphabet is defined for each appliance type individually based on its collected historical data. Applying a symbol mapping to a mismatching appliance type can result in very atypical symbol sequences; an analysis of such appliance-specific fitting effects is presented in Sec. IV-B.




Fig. 3. Symbolic representations are allocated based on clustering the histogram and assigning consecutive symbols to each cluster and each gap

C. Motif Extraction

Following the reduction of the value range, we apply a second preprocessing step and convert groups of consecutively occurring symbols into *motifs*. Motifs are symbol sequences of a pre-determined length l_{sym} (their *symbolic length*) and represent the key data entities used through the remainder of this paper. As motifs incorporate both amplitude information and a signal's temporal progress, they are well suited to reflect the observations made in Sec. II-A in symbolic data while incurring lower computational efforts for their processing.

In order to speed up subsequent motif processing operations, motif identifiers i are pre-allocated for each potential motif M_i in advance. For example, in a setting with $l_{sym}=2$ and $l_{alph}=2$, the symbol sequence AA in the input data will be mapped to motif M_1 , $AB \rightarrow M_2$, $BA \rightarrow M_3$, and $BB \rightarrow M_4$. For reference, the used nomenclature is also given in Table I. Once a lookup table for these mappings has been created, symbolic data are converted to their motif representations. This is realized by shifting a window of length l_{sym} across the historic data sequence and replacing all subsequences by their corresponding motif identifiers.

III. THE MOTIF TRANSITION LIKELIHOOD MATRIX

Once the input data have been converted to their corresponding motif representation, the motif transition likelihood matrix (MTLM) is extracted from the motif sequence. The MTLM models temporal dependencies between motif transitions on significantly longer time scales than the motif length l_{sym} , up to an upper limit of w seconds. Thus, it enables our system to detect anomalous consumption patterns in this window. Choosing a value for w depends on the appliance under consideration. For best accuracy, at least one full operational cycle has to be contained within the window, e.g., any value in excess of 2200 seconds for the dishwasher shown in Fig. 1a. Note that the chosen window size has an immediate effect on the system's storage and computational demand. The following three processing steps are executed when populating the MTLM.

TABLE I
NOMENCLATURE USED IN THIS PAPER

Symbol	Explanation
l_{alph}	Number of entries in the symbol alphabet
l_{sym}	Number of symbols in each motif
M	Motif, i.e., a sequence of l_{sym} symbols
$T_{M_i \rightarrow M_j}(t)$	Number of transitions $M_i \rightarrow M_j$ for time difference t
$P_{M_i \rightarrow M_j}(t)$	Probability of transition $M_i \rightarrow M_j$ at time difference t
w	Maximum history depth stored in the MTLM

A. Motif Transition Occurrence Collection

Once the input data is present as a sequence of motifs, our systems counts the transitions between any two motifs M_i and M_j (with $i, j = 1, \dots, l_{alph}^{l_{sym}}$) within time window w , as well as the temporal distance between the motif occurrences. Counted transitions are stored in a three-dimensional transition occurrence matrix, whose first dimension is the starting motif M_i (with $i=1, \dots, l_{alph}^{l_{sym}}$). In the second dimension we store the ending motif M_j , and in the third dimension of the matrix we annotate the time difference between the motif occurrences. In essence, the occurrence matrix thus contains the number of transition occurrences between starting and ending motif for all time differences Δt between 1 and w epochs (as no transition can occur within the same epoch, no entry is needed for $\Delta t=0$).

It may appear that generating the occurrence matrix exhibits a high computational and memory demand, especially when long motifs or many entries in the symbol alphabet are being used. While this is correct, the occurrence matrix is only extracted once from historical data, and significantly less memory and computations are required throughout the remaining operation of the system. Furthermore, the matrices in our evaluations have always been populated very sparsely; when computed for the dishwasher used in our previous example for an alphabet length of $l_{alph}=6$, a symbolic length of $l_{sym}=4$, and a history depth of $w=60$ seconds, only 21,262 out of more than 100 million theoretically possible entries (i.e., only 0.02%) have been populated in the occurrence matrix. Even when increasing the history depth to several hours, no more than 0.2% of the potentially available occurrence matrix entries were populated.

B. Motif Transition Likelihood Matrix Generation

The next step in the process of generating the MTLM is realized by applying a column-wise normalization to the matrix. This transformation is formulated in Eq. (1) and executed for all origin motif identifiers i and destination identifiers j . A visualization of the resulting normalized transition probabilities for source motif $AAAB$ and $w=600s$ for the dishwasher appliance is shown in Fig. 4.

$$P_{M_i \rightarrow M_j}(t) = \frac{T_{M_i \rightarrow M_j}(t)}{\sum_{k=1}^{l_{alph}^{l_{sym}}} T_{M_i \rightarrow M_k}(t)} \quad \forall i, j = 1, \dots, l_{alph}^{l_{sym}} \quad (1)$$

This normalized matrix is already sufficient to determine the likeliness of an appliance's likeliness of normal operation, as will be described in Sec. III-C. However, one more processing step is being applied to eliminate the dominance of self-transitions during device inactivity, i.e., transitions between two occurrences

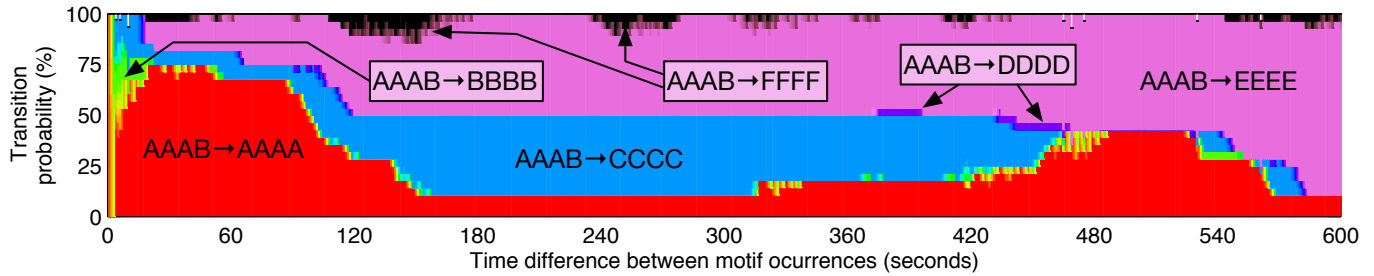


Fig. 4. Transition likelihood plot for the dishwasher's transitions originating from motif AAAB (annotations omitted for transitions with low probabilities)

of the first motif (AAAA→AAAA). Our approach considers their occurrence frequencies when normalizing the transition probabilities, but subsequently disregards all transitions between these motifs in the likelihood computation. The reason for suppressing these transitions is based on the fact that they carry little informational value, as these self-transitions almost exclusively occur when an appliance is in standby mode or off.

C. Motif Sequence Likelihood Determination

With the MTLM established, our system is ready to detect the likelihood of power consumption motif sequences observed in live data. This process is visualized in the right-hand part of Fig. 2. It shows that the same preprocessing steps (described in Secs. II-B and II-C) are being applied to input data. Once its series of motifs has been extracted, however, they are being compared against the transition probabilities modeled in the MTLM in order to determine the likeliness that the encountered motif sequence originated from the trained appliance.

For the sake of computational simplicity, we have decided in favor of a multiplicative computational model. It operates on an input sequence composed of L motifs (without loss of generality we assume $L > w$ here), indexed by their time of collection with $M(t=L)$ denoting the most recently encountered and $M(t=0)$ the oldest collected motif. The probability computation for a single time instance λ is then shown in Eq. (2).

$$P(\lambda) = 1 - \prod_{\tau=1}^{w-(L-\lambda)} (1 - P_{M(t=\lambda-\tau) \rightarrow M(t=\lambda)}(\tau)) \quad (2)$$

The rationale behind computing the converse probability is simple, and best motivated by comparing it to a multiplicative model (such as $P = \prod P_{M(t=w-\tau) \rightarrow M(t=w)}(\tau)$). In this case, a single unlikely transition (i.e., with very small probability) would immediately bring the overall output close to zero and thus classify the appliance operation as unusual. In the chosen converse model, the opposite case is true: More likely observed transitions lead to a smaller product term and thus the result of the computation increases. For the computation of the overall likeliness of anomalous operation, we use Eq. (2) to compute the aggregate likeliness for all motifs leading up to a given motif $M(t=\lambda)$ and repeat the computation for all time instances between $L-w+1$ and L . The equation to compute the overall likeliness as the arithmetic mean across the window w is formulated in Eq. (3).

$$P_{total} = \frac{1}{w} \sum_{\lambda=L-w+1}^L P(\lambda) \quad (3)$$

IV. EVALUATION

Once our system has pre-computed the transition matrices for all appliance types under consideration, it is ready for operation. By supplying it with appliance power consumption motif sequences, the estimated likelihood of normal device operation, P_{total} is computed based on Eq. (3). We demonstrate the system's capability of detecting unusual power consumption patterns in the following subsections.

A. Input Data Selection

As our system has been developed to operate on appliance-level data collected by smart plugs, we have sourced our data from the Tracebase [8]. This data set features more than 1,500 power consumption traces collected from more than 100 different appliances. Before we use the traces in our evaluation, we apply an interpolation step in order to ensure that all data share the same temporal resolution of 1 Hz sampling rate. A second filtering step is subsequently applied to remove erroneous outliers, to facilitate the operation the histogram clustering algorithm.

B. MTLM Adaptation to the Underlying Appliance Type

Before using the data for our analysis of the system's capacity to detect anomalous behavior, we investigate how closely MTLMs can adapt to the underlying appliance type. To this end, we have supplied the system with the input data from eight typical household appliances. We have chosen $l_{sym}=4$ symbols per motif and selected w such that all transitions of at least one complete activity cycle were available when constructing the MTLM. Subsequently, we have computed the average P_{total} values when testing further traces against each stored MTLM. The results are shown in Table II, where the appliance types used to extract the MTLM are listed in the leftmost column, and the data sets for which we have requested the system to return a likelihood estimation are presented in the top row. The table entries show the P_{total} value for each MTLM, showing the selectiveness of our systems for some trained appliances (e.g., the bean-to-cup coffee maker). However, the results in the table also indicate that MTLM entries are not always specific to one given appliance type, but corresponding motifs can partially also be found in the power consumptions of other appliances.

C. Anomalous Consumption Detection

Next, we assess the system's performance when it is confronted with anomalous appliance behavior. To this end, we have manually introduced errors to two trace files in order to quantify the detection delay of our system. We visualize the traces in

TABLE II
 P_{total} VALUES WHEN MATCHING APPLIANCE TRACES AGAINST OTHER APPLIANCE'S MTLMS

MTLM of appliance type	Average P_{total} score when comparing this appliance type with the MTLM							
	CM	DW	FR	KE	MW	RE	SM	WM
Bean-to-cup coffee maker (CM)	99.9 %	10.0 %	0.9 %	0.0 %	0.6 %	0.5 %	0.0 %	2.9 %
Dishwasher (DW)	4.3 %	21.2 %	97.3 %	4.3 %	8.9 %	98.6 %	5.5 %	8.8 %
Freezer (FR)	5.1 %	18.2 %	97.5 %	8.5 %	9.0 %	97.5 %	5.6 %	7.9 %
Kettle (KE)	0.0 %	13.1 %	81.9 %	8.5 %	8.4 %	71.2 %	5.3 %	4.3 %
Microwave oven (MW)	2.6 %	18.2 %	97.4 %	8.4 %	9.1 %	98.5 %	5.6 %	8.4 %
Refrigerator (RE)	1.9 %	23.4 %	44.4 %	67.8 %	60.2 %	98.7 %	4.9 %	77.3 %
Sandwich maker (SM)	0.0 %	9.9 %	81.2 %	8.1 %	8.4 %	71.1 %	5.6 %	4.1 %
Washing machine (WM)	99.9 %	21.2 %	97.5 %	4.3 %	9.0 %	98.7 %	5.6 %	10.2 %

which the anomalies shall be recognized in the top graphs of the diagrams in Fig. 5. In our example, we use two excerpts of a refrigerator's power consumption, to which we have added two types of faults. Firstly, we simulate a continuous motor operation as shown in Fig. 5a. This situation could, e.g., occur as a result of poor insulation or leaving the refrigerator door open. Secondly, we assess the temporary disconnection of the appliance from mains, shown in Fig. 5b, which can be observed in practice when, e.g., de-frosting the device. The values of P_{total} for the two test cases are plotted in the lower parts of the figures using continuous lines. Successively, as more and more motif transitions match the entries in the MTLM, the value of P_{total} changes correspondingly. The clearly visible drop in P_{total} is an indicator for anomalous consumption, and by simply applying a threshold value (e.g., 0.5), the automated detection of anomalies is facilitated.

D. Performance Comparison to Value Averaging

At last, we compare our proposed anomaly detection scheme to the alternative approach of monitoring an appliance's mean power consumption. To minimize fluctuations of the resulting averaged value, the averaging window length needs to be aligned with the duration of the appliance's regular activity period; we have consequently used a sliding averaging window of 12,000 seconds length in our evaluation.

The mean values of the refrigerators' power consumptions are shown as dashed lines in the lower subgraphs of Figures 5a and 5b. While they show similar trends to the MTLM-based probability, it needs to be noted that they only allow for anomaly detection when threshold consumption levels to discriminate normal from anomalous appliance operation have been defined. As these values must include a safety margin (e.g., due to the variations of the refrigerator's period length), their careful selection is not always easy. Moreover, the averaging function inherently introduces a delayed response to changes, whereas our MTLM-based approach quickly recovers; this is particularly visible in Fig. 5b when the refrigerator recommences operation.

V. RELATED WORK

The computational analysis of household energy consumption data is looking back on more than 20 years of history. In his early work [13], George Hart has shown how characteristic features can be extracted from household aggregate power levels, and linked to the (de)activation of specific appliances. Despite its limitation to identify the presence of ON/OFF appliances (i.e.,

devices that result in an equal change of their power consumption when switched on and off), it has inspired manifold new research results on disaggregating home energy consumption and providing new services. Several approaches have been presented to extract and efficiently store patterns in power consumption data in both the time and the frequency domain in order to infer contextual information. For example, [14] and [15] convert load data to the frequency domain in order to infer load types. While the analysis of the signal's spectral components speeds up the processing times, characteristic temporal dependencies are inherently eliminated when choosing this approach. As a result, time series analysis has gained an increasing impact, e.g., [16, 17], because it allows for the exploitation of the rich temporal dependencies in the data. The large computational overhead of operating on raw time series data has been addressed in research as well, e.g., by means of converting data to their symbolic approximations [18, 19], reducing their temporal resolution by means of chunking [11], or solely operating on events extracted from the data [20]. While our presented approach also relies on symbolic representations of time series data, we do not reduce the temporal resolution any further to maintain data fidelity and the possibility to quickly react to anomalies.

Orji et al. have presented another approach to detect anomalies in power consumption data in [21], but require data collected at high sampling rates to analyze higher-frequency harmonics in the signal. The FailureSense system [20] models the typical reporting interval of sensors and their dependencies in order to detect unexpected data transmission characteristics. Atypical sensor data patterns are subsequently being used to determine sensor faults, although anomalies in the actual data values are not considered in the paper. Operating on a daily scale, the work by Seem [22] shows how the extraction of features from 24-hour long traces can be leveraged to determine abnormal energy consumption in buildings. Shorter time scales, like the ones used in this paper, are however not supported. Finally, our approach can be seen as an instance of pattern mining as well. Sequential pattern mining [23, 24] assumes that pre-known patterns are contained in the signal, however without temporally bounding their potential occurrences based on historical data. Our approach extracts occurrence sequences from historical data precisely, and thus confines the potential location of a normal event occurrence more tightly than standard pattern mining approaches. It furthermore differs from [25], in which motif mining is successfully applied to support load disaggregation algorithms, but does not target the detection of anomalies.

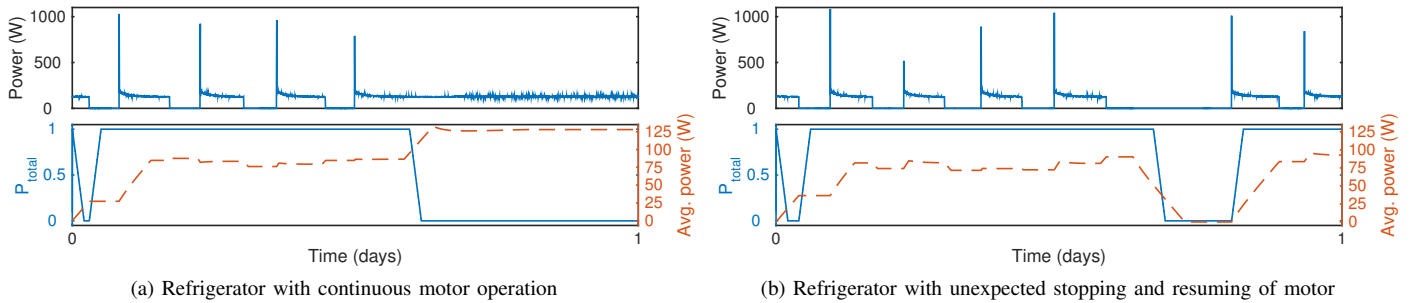


Fig. 5. Correlation between input data and value of P_{total} for two sample anomaly cases

VI. CONCLUSIONS

We have presented an approach to detect anomalies in power consumption data collected by smart plugs. By operating on motifs, its computational burden is significantly lower than when operating on raw time series data. The core idea behind our approach lies in determining the dependencies and typical time differences between motif transitions and storing these values in a motif transition likelihood matrix. Subsequently, any power sensor data stream can be compared against the expected motif transitions for a given appliance type, and anomalies can be quickly detected at low computational overhead. We have confirmed the practical viability of our approach and evaluated its performance using appliance-level data collected from real-world refrigerating devices. Its application in emerging scenarios like smart homes can thus be envisioned to contribute to the safety and security of users.

REFERENCES

- [1] The European Parliament and the Council of the European Union, "Directive 2012/27/EU on Energy Efficiency," *Official Journal of the European Union*, no. L 315, pp. 1–56, 2012.
- [2] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava, "NILMTK: An Open Source Toolkit for Non-intrusive Load Monitoring," in *Proceedings of the 5th Conference on Future Energy Systems (e-Energy)*, 2014, pp. 265–276.
- [3] D. Egarter, V. P. Bhuvana, and W. Elmenreich, "PALDi: Online Load Disaggregation via Particle Filtering," *IEEE Transactions on Instrumentation and Measurement*, vol. 64, no. 2, pp. 467–477, 2014.
- [4] O. Parson, S. Ghosh, M. Weal, and A. Rogers, "An Unsupervised Training Method for Non-intrusive Appliance Load Monitoring," *Artificial Intelligence*, vol. 217, pp. 1–19, 2014.
- [5] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz, "On the Accuracy of Appliance Identification Based on Distributed Load Metering Data," in *Proceedings of the 2nd IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*, 2012, pp. 1–9.
- [6] M. Kazandjieva, O. Gnawali, B. Heller, P. Levis, and C. Kozyrakis, "Identifying Energy Waste through Dense Power Sensing and Utilization Monitoring," Stanford University, Tech. Rep. CSTR 2010-03, 2010.
- [7] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. E. Culler, "Design and Implementation of a High-Fidelity AC Metering Network," in *Proceedings of the 8th International Conference on Information Processing in Sensor Networks (IPSN)*, 2009, pp. 253–264.
- [8] A. Reinhardt, P. Baumann, D. Burgstahler, M. Hollick, H. Chonov, M. Werner, and R. Steinmetz, "On the Accuracy of Appliance Identification Based on Distributed Load Metering Data," in *Proceedings of the 2nd IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*, 2012, pp. 1–9.
- [9] R. Agrawal, C. Faloutsos, and A. N. Swami, "Efficient Similarity Search In Sequence Databases," in *Proceedings of the 4th International Conference on Foundations of Data Organization and Algorithms (FODO)*, 1993, pp. 69–84.
- [10] E. J. Keogh and M. J. Pazzani, "A Simple Dimensionality Reduction Technique for Fast Similarity Search in Large Time Series Databases," in *Proceedings of the 4th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2000, pp. 122–133.
- [11] S. Rollins and N. Banerjee, "Using Rule Mining to Understand Appliance Energy Consumption Patterns," in *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2014, pp. 29–37.
- [12] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD)*, 1996, pp. 226–231.
- [13] G. Hart, "Nonintrusive Appliance Load Monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.
- [14] M. Petrescu and M. Scutariu, "Load Diagram Characterisation by Means of Wavelet Packet Transform," in *Proceedings of the 2nd Balkan Power Conference (BPC)*, 2002, pp. 15–19.
- [15] E. Carpaneto, G. Chicco, R. Napoli, and M. Scutariu, "Electricity Customer Classification Using Frequency-Domain Load Pattern Data," *Electrical Power and Energy Systems*, vol. 28, no. 1, pp. 13–20, 2006.
- [16] G. Chicco, R. Napoli, F. Piglion, M. Scutariu, P. Postolache, and C. Toader, "Emergent Electricity Customer Classification," *IEEE Proceedings on Generation, Transmission & Distribution*, vol. 152, no. 2, pp. 164–172, 2005.
- [17] V. Figueiredo, F. Rodrigues, Z. Vale, and J. Gouveia, "An Electric Energy Consumer Characterization Framework based on Data Mining Techniques," *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 596–602, 2005.
- [18] A. Notaristefano, G. Chicco, and F. Piglion, "Data Size Reduction with Symbolic Aggregate Approximation for Electrical Load Pattern Grouping," *IET Generation, Transmission & Distribution*, vol. 7, no. 2, pp. 108–117, 2013.
- [19] A. Reinhardt and S. Koessler, "PowerSAX: Fast Motif Matching in Distributed Power Meter Data Using Symbolic Representations," in *Proceedings of the 9th IEEE International Workshop on Practical Issues in Building Sensor Network Applications (SenseApp)*, 2014, pp. 18–25.
- [20] S. Munir and J. A. Stankovic, "FailureSense: Detecting Sensor Failure using Electrical Appliances in the Home," in *Proceedings of the 11th IEEE International Conference on Mobile Ad hoc and Sensor Systems (MASS)*, 2014, pp. 1–9.
- [21] U. A. Orji, Z. Remscrim, C. Laughman, S. B. Leeb, W. Wichakool, C. Schantz, R. Cox, J. Paris, J. L. Kirtley Jr., and L. K. Norford, "Fault Detection and Diagnostics for Non-intrusive Monitoring Using Motor Harmonics," in *Proceedings of the 25th Annual IEEE Applied Power Electronics Conference and Exposition (APEC)*, 2010, pp. 1547–1554.
- [22] J. E. Seem, "Using Intelligent Data Analysis to Detect Abnormal Energy Consumption in Buildings," *Energy and Buildings*, vol. 39, pp. 52–58, 2007.
- [23] R. Agrawal and R. Srikant, "Mining Sequential Patterns," in *Proceedings of the 11th International Conference on Data Engineering (ICDE)*, 1995, pp. 3–14.
- [24] N. R. Mabroukeh and C. I. Ezeife, "A Taxonomy of Sequential Pattern Mining Algorithms," *ACM Computing Surveys*, vol. 43, no. 1, pp. 3:1–3:41, 2010.
- [25] H. Shao, M. Marwah, and N. Ramakrishnan, "A Temporal Motif Mining Approach to Unsupervised Energy Disaggregation: Applications to Residential and Commercial Buildings," in *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013, pp. 1327–1333.