Mr. Manjunath Raikar. International Journal of Engineering Research and Applications
www.ijera.com
ISSN: 2248-9622, Vol. 13, Issue 11, November, 2023, pp: 247-253

RESEARCH ARTICLE OPEN ACCESS

Utilizing Very Low-Frequency Smart Meter Time Series for Appliance Detection

Mr. Manjunath Raikar
Lecturer / Department of Electrical & Electronics
Government Polytechnic – Mundgod - 581349, Dist U.K.
manjunathraikarn@gmail.com

Date of Submission: 25-10-2023 Date of acceptance: 03-11-2023

ABSTRACT

Electricity providers have been using smart meters extensively in recent years to enhance the smart grid system's administration. These meters often gather information on energy use at a relatively low frequency (every 30 minutes), which helps utilities provide more accurate bills to their consumers. Due to the extremely low frequency of meter readings, identifying the appliances that consumers possess is the next difficult task that must be completed to make more individualized suggestions. While there are several classifiers for time series classification that have been presented in the literature, no research has applied and compared them to the appliance detection problem, even though it may be seen as a time series classification problem. This work provides a comprehensive analysis and comparison of the most recent time series classifiers used to identify different appliances in extremely low-frequency smart meter data.

We present our findings using five actual datasets. Utilizing 30-minute sampling data, we first examine the effects of 13 distinct appliances' detection quality. Then, we suggest analyzing the potential improvement in detection performance that might result from utilizing a greater meter reading frequency. The findings show that there are large variations in the performance of the time series classifiers in use today. Even with 30 minutes of sampled data, some of them—specifically, deep learning-based classifiers—show promise in terms of accuracy (particularly for specific appliances) and are scalable to the sizable smart meter time series collections of energy consumption data that electricity providers currently have access to.

Keywords: Appliance Detection, Smart Meter Data, Time Series Classification

I. INTRODUCTION

The desire for a more secure and sustainable energy supply is the main factor driving the massive changes that the energy sector is experiencing. Gaining a deeper understanding of our consumption is one method to better control it. Millions of smart meters have been placed globally by power providers in the past ten years in an effort to better control the electrical grid [10]. With the help of these meters, which capture comprehensive time-stamped data on power use, enterprises and individual consumers alike may more effectively analyze and justify their usage [6]. Suppliers can also benefit from this data as they can more precisely predict energy consumption. All things considered, the widespread use of smart meters is essential to the shift to a more efficient and sustainable energy system.

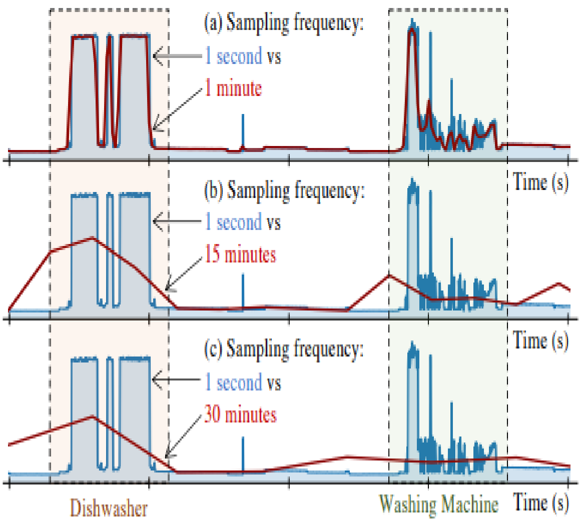

Figure No. 1 - load curve comparisons between a washing machine and a dishwasher at various sample frequency

We see that knowledge about the electrical items owned by clients has become crucial for power providers. With the use of this information, providers are better able to segment their clientele [3] and, as a result, provide tailored offerings and services that boost client retention and satisfaction. They can also assist consumers in rationalizing their electricity usage, which will aid in the energy shift. A consumption questionnaire can be used to directly ask clients for this information. Customers might not be willing to accept such a large time and resource commitment, and this strategy is prone to mistakes. Electricity providers must thus devise more effective and non-intrusive methods of obtaining this data, such as by employing sophisticated data analytics methods to identify the appliances directly from the information gathered from smart meters [2].

Many methods are used to identify the presence of devices, and appliance detection has grown to be an important field of study [3]. Non-Intrusive Load Monitoring (NILM), which seeks to determine an individual appliance's power usage, pattern, or on/off status activation using just the overall consumption series, is strongly linked to this issue [9]. Although several ways have been presented in the literature [5], and appliance detection may be seen as a phase in NILM-based systems [7], they are not the same as our goal. In fact, rather than determining if a home has a certain device, these studies primarily concentrate on determining when that item is "ON," and in many situations, such knowledge existed prior to the use of these techniques. Furthermore, most NILM research employ signature-based approaches [36, 48] since they rely on data collected at ≥1Hz, which necessitates training on each appliance power usage or knowledge of how each appliance works. However, most smart meter installations that are already in place only sample data once every 10 to 60 minutes (or even less often in some circumstances).

As a result, certain appliance pattern information is lost or smoothed out in the signals. Figure 1 shows how this information is lost. We see that when the sample frequency decreases, it gets harder to tell the dishwasher's (on the left) and washing machines (on the right) signatures apart. Therefore, at the sample frequencies that are really utilized in practice, it becomes impractical to reliably detect appliances using signature-based approaches. In this work, we present a benchmark of various cutting-edge classification techniques for the issue of detecting appliances in extremely low-frequency electricity consumption time series. We use several time series classifiers to conduct our experimental assessment on five genuine smart meter datasets. We first concentrate on identifying appliances in extremely low-sampled smart meter data (30-minute level), as this is now one of the common sampling rates used by power providers. The growing detection quality is then thoroughly examined utilizing higher frequency smart meter readings of 15 minutes, 10 minutes, and 1 minute. To the best of our knowledge, this is the first research to thoroughly compare 11 cutting-edge techniques for various sampling frequencies using five distinct actual datasets and 13 different types of appliances.

Even at the 30-minute resolution, the experimental assessment shows that the time series classifiers in use today are capable of reliably detecting several appliances. More specifically, when big datasets of smart meters are used, deep learning approaches prove to be the most accurate and scalable. Furthermore, we show how utilizing time series classifiers to significantly improve appliance recognition is possible when the smart meter reading frequency is set to 1 minute. The following is a summary of our contributions.

- We present a publicly accessible framework for evaluating the effectiveness of several time series classification techniques for the appliance detection problem.
- Using five different actual datasets and eleven time series classifiers—including both conventional machine learning and deep learning techniques—we conduct a thorough experimental evaluation.
- We present our comparison's findings, which show that: (i) deep learning classifiers are the most accurate and scalable solution; (ii) electricity suppliers should aim for a minimum smart meter reading frequency of 15 minutes; and (iii) current time series classifiers can only detect specific appliances at the 30-minute resolution.
- The results of this study can assist power providers in making well-informed choices about the features of upcoming smart meter rollouts. Furthermore, these results provide intriguing (and still difficult) avenues for future study in time series analysis of power usage, and specifically appliance identification.

II. RELATED WORK

Smart Meter Data, A univariate time series X = ($\boldsymbol{x}$ 1, $\boldsymbol{x}\ T$) of ordered items $\boldsymbol{x}\ \ j$ ∈ R 1 +} after (i 1, $i\ T$) time consumption indexes (i.e., timestamps) is what is known as an electrical consumption load curve. The definition of sampling frequency is the difference in time between two recordings, index Δt B $i\ j\ -\ i\ j$ −1. Every element $\boldsymbol{x}\ \ j$, which is often expressed in Watt, denotes the average electric power called throughout the interval time Δt or the actual power at time $i\ j$. Another way to express the value is in watt-hours. The definitions of high-frequency and low-frequency smart meter data might vary in the literature [4]. We discuss low-frequency data collected between one second and one minute, and high-frequency data recorded at less than one second in this paper. Very low-frequency smart meter data is indicated by data samples taken longer than one minute. [Load curve for individual appliances] The consumption load curve for each individual appliance in a home may be obtained by using individual meters to monitor electric devices. But the cost of instrumenting every household item is unaffordable. [The load curve aggregated] A smart meter gadget that is mounted on the household's electrical meter typically records the primary power use of a home. The total power usage of all the household's individual appliances makes up this aggregate signal.

Non-Intrusive Load Monitoring (NILM) and Appliance Detection, using just the total aggregated load curve [29], Non-Intrusive Load Monitoring (NILM) [2], also known as load disaggregation, is based on determining the individual power consumption, pattern, or on/off state activation of individual appliances. Initially, methods for NILM were designed to estimate the percentage of total power consumption used by different active appliances at each time step, treating the issue as one containing linear combinations [2]. Combinatorial optimization approaches were used in earlier studies on this subject [9]. Subsequently, Hidden Markov Models took over as the most used method, and deep learning models have become the standard for doing disaggregation in recent years [9]. Furthermore, based on whether labeled data is used to train the models, NILM techniques can be further classified as supervised or unsupervised learning. With supervised learning, events such as appliances turning on or off are classified by comparing extracted attributes [3]. Unsupervised NILM techniques, on the other hand, do not require labeled data; instead, they identify events by examining feature similarities or correlations [6].
Various techniques have been proposed in the literature to identify appliances in load curves utilizing high- or low-frequency smart meter data, since device recognition may be viewed as a stage of NILM-based systems [7]. Several studies that use low-frequency pattern recognition need an understanding of how each device functions. A small number of recent studies [5] have employed deep learning representations or time series characteristics to identify patterns of appliance activation or occurrences. We point out that this research, which uses contemporary machine learning techniques and shows encouraging results, are limited to high-frequency data—that is, data collected at a minimum rate of one sample per second.

III. PROBLEM DEFINITION AND BENCHMARK

The appliance detection problem is approached as a supervised binary classification problem in this work. Regardless of how many times an appliance has been activated, our goal is to determine whether the activation signature of that appliance is present or absent in a smart meter data set. The fact that the gadget has been turned "ON" at least once serves as a straightforward definition of its presence. Formally, the issue is defined as follows:

We now give a summary of the many methods that have been suggested in the literature to address the TSC problem (see Figure 2). Comparing how well different approaches work with the appliance detection challenge is the goal. Each classifier uses the ground-truth labels and the univariate consumption time series (i.e., 1D signal as training data.

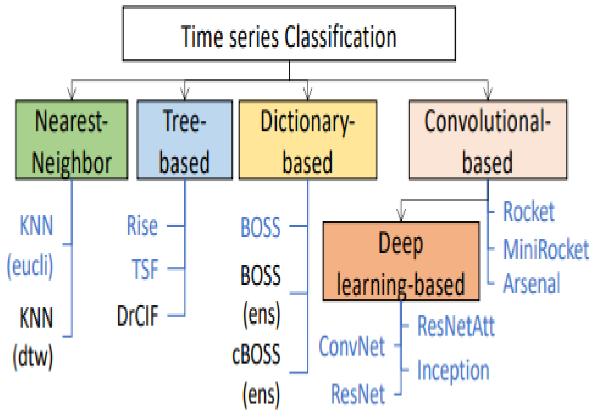

Figure No. 2 – Taxonomy of Classifier considered in our benchmark

KNN, based on the idea of time series similarity, classifiers are the most basic and understandable classifiers. Each new instance is categorized by assigning the same label as the majority label of the K nearest samples in the training set, after a selected distance measurement (K = 1 in our experiments, i.e., we use 1-NN classifiers). Euclidean distance is the most widely used measure of distance because it makes point-to-point comparisons between two instances possible. Nevertheless, this distance doesn't account for potential temporal axis deviations. A distance metric called Dynamic Time Warping (DTW) [49] is used to calculate how similar two time series are, even if the underlying patterns may change in pace over time. Because of its high processing cost, DTW is difficult to use with big data sets.

Tree Based Classifier, in classification challenges, classifiers such as Random Forest [8] have shown encouraging results. [Forest of Time Series] A random forest-based classifier called TSF [6] employs characteristics taken from the raw data series at randomly chosen intervals as input. The mean, standard deviation, and slope are the three basic properties that are taken from each of the r intervals that the algorithm has originally chosen, together with their random length and start point. Ultimately, a traditional random forest classifier is trained using the $3\boldsymbol{U}$ additional features. By default, the number of intervals is set to $\sqrt{T}$, where T is the length of the input time series. The decision tree's number of estimators is set at 200.

Deep Learning Based Classifier, Over the past several years, there has been a noticeable increase in interest in deep learning techniques for time series categorization [8]. These models have outperformed the state-of-the-art in terms of performance. [ConvNet] Convolutional neural networks (CNNs) [4] are a subset of deep learning neural networks that are frequently employed in image recognition applications. They are specifically made to extract patterns from input that has a grid-like structure, such time series or photographs. Convolution is the method used by CNN to apply a filter to a sliding window covering a time series. The three-layered convolutional blocks in the ConvNet architecture described in [10] are followed by global average pooling [7] and a Softmax activation algorithm.

IV. ENERGY CONSUMPTION DATASET

The literature has many datasets on energy use [9], some of which have been used as references for NILM research [13]. Nevertheless, with a high sample frequency, these datasets usually only include aggregated and appliance-level load curves for a small number of homes. They may be resampled extremely often, which significantly reduces the amount of data. To accommodate a wider variety of appliances and be consistent with previous research, we have incorporated two NILM datasets, UK-DALE [1] and REFIT, into our studies. We have included two private datasets from EDF (The main French electricity supplier. We consider five diverse datasets in our experiments.

NILM Collections. Two well-known high-frequency Smart Meters datasets that are utilized in NILM investigations are UKDALE and REFIT [1]. [UK-DALE] The UK-DALE dataset [31] comprises data from five UK homes, including whole-house aggregate data series captured at 16kHz and appliance-level load curves taken every six seconds. While the fifth home was documented for 655 days, the other four were recorded for more than a year and a half. [REFIT] Using smart home technology, the REFIT project (Personalized Retrofit Decision Support Tools for UK Homes) was conducted from 2013 to 2015. Twenty UK homes were observed with several sensors and smart meters throughout this time, and the data was collected. Appliance load curves, both total and individual, are provided in this dataset with sample intervals of every 8 seconds.

CER Information, to evaluate smart meter functioning and customer energy usage, the Commission for Energy Regulation of Ireland [1] recorded the aggregate load curve consumption for more than 5,000 Irish homes and businesses every 30 minutes. Participants answered questions on the makeup of the family, how they use power, and what kinds and numbers of appliances they have at home or at work. The study's residential sub-group, or 4225 homes that collected data between July 15, 2009, and January 1, 2011, is what we employ in this work. In all, 4225 series, each with a length of 25728 data points, were recorded.

EDF Collections. Electricity De France (EDF) surveys consumer samples in order to gain a better understanding of its customer base and power consumption patterns. Only the total power usage of the home is recorded, and these customers have given EDF permission to utilize their data and examine their consumption patterns. Customers respond to a questionnaire, much to the CER research, providing details on the equipment in their homes as well as their usage patterns. For our investigations, two EDF datasets from two distinct research were employed.

V. EXPERIMENTAL SETUP

Every experiment is run on a cluster of high-performance computers. The default settings given by the authors in the original publications are used for each classifier in the Python 3.7 source code. For methods other than deep learning, we employ the sk time library [3]. Every experiment is run on a server equipped with two Intel Xeon Gold 6140 CPUs and 190 GB of RAM. We use the 1.8.1 version of the PyTorch framework [11] to create all models for deep learning, and we conduct experiments on a server that has two NVidia V100 GPUs and 16GB of RAM.

We consider every classifier that is described in Section 3.2. We use distinct random trains, validation, and test splits to run each method five times, and we present the average of these runs. It should be noted that the error bars in Figures 3, 7, and 8 represent the classifiers' average variability throughout the course of these five runs. We also impose a 10-hour time restriction on each work. The models are evaluated only once they have completed a run (training + inference). We observe that the dilation convolution of the residual block in the ResNet with Attention model was not consistent with the tiny size of the time series of the UKDALE and REFIT datasets, so this model was not assessed using these datasets.

Data Preprocessing, we preprocess the datasets for the experiments as detailed below because they were prepared using varying sample frequencies for this investigation. Table 1's left section lists the number of time series and their related lengths for each dataset, broken down by sampling frequency. NILM preprocessing of datasets. Appliance level and total consumption load curves for a limited number of homes—five and twenty, respectively—are provided by the REFIT and UKDALE databases. Furthermore, it's possible that the homes' electrical appliances are identical. We divided each household's whole consumption curve into smaller sub-sequences as part of the preprocessing of the datasets, which was inspired by the data processing phase in NILM research [9].

Datasets	Tot. TS	TS Length					Appliance case	Datasets									
								REFIT		UKDALE		CER		EDF 1		EDF 2	
		1min	10min	15min	30min			♯TS	IB Ratio	♯TS	IB Ratio	♯TS	IB Ratio	♯TS	IB Ratio	♯TS	IB Ratio
REFIT	9091	1440	144	96	48	Tech	Desktop Computer	5190	0.56	/		3286	0.47	1402	0.38	3740	0.62
							Television	1134	0.92	/		/		/		/	
UKDALE	4767	1440	144	96	48	Kitchen	Cooker	/		/		1682	0.76	/		/	
							Kettle	4790	0.72	1222	0.84	/		/		/	
							Microwave	7434	0.55	1678	0.77	/		324	0.91	/	
							Electric Oven	/		/		/		510	0.85	1152	0.91
CER	4225	/	/	/	25728	Washer	Dishwasher	7798	0.44	2378	0.32	2350	0.66	224	0.93	2846	0.75
							Tumble Dryer	3466	0.22	/		2214	0.68	1534	0.41	3470	0.42
							Washing Machine	7422	0.54	2830	0.38	/		/		/	
EDF 1	2611	/	/	/	17520	Heating	Water Heater	/		/		3070	0.56	1336	0.66	548	0.86
							Electric Heater	/		/		1348	0.19	1624	0.58	1538	0.56
							Convector/Heat Pump	/		/		/		506	0.69	/	
EDF 2	1553	/	26208	17472	8736	Other	Electric Vehicule	/		/		/		140	0.3	/	

Table 1 – Time Series Database

We first resample the data for each trial to a predetermined sampling rate, and then we use linear interpolation to fill in the gaps of less than an hour. Next, we analyze the datasets by removing any missing values and dividing the consumption load curve for each home into smaller sub-sequences of a single day. A general balance between positive (i.e., containing the device) and negative variables is achieved by selecting one day as the sub-sequence duration. This is because, on average, the appliances in these datasets are highly utilized devices—using them once every two or three days. Using the matching disaggregated appliance load curve, we may determine if the appliance is present in a subsequence by assigning a positive or negative label.

VI. RESULT AND DISCUSSION

The findings from our experimental examination are shown in this section.

To get general findings for every example, we first normalize the various datasets to the same sample frequency, or 30 minutes. Next, we conduct an experimental assessment to determine how sample frequency affects the classifiers' ability to recognize patterns. We also examine how the amount of data affects the quality of detection. We conclude with a review of the overall outcomes.

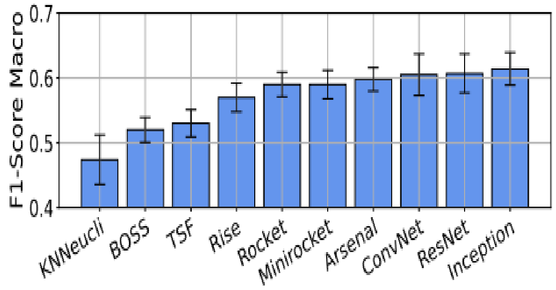

Figure No. 3 – Classifier Detection Score

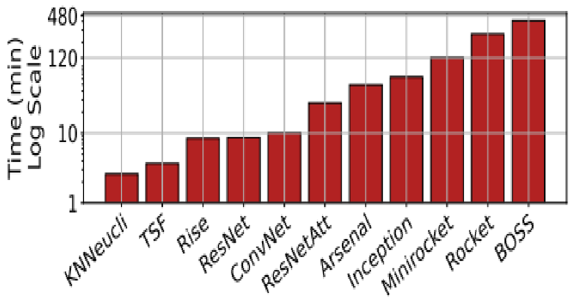

Figure No. 4 – Running Time Per Run

Table 2 provides a summary of the classifiers' appliance detection findings for a 30-minute sample frequency. We note that the UKDALE dataset yields subpar results from all classifiers. (We go into more depth about these findings in Section 5.3.) Furthermore, we see that some appliances are simpler to identify than others, regardless of the dataset. The findings are analyzed based on the kind of appliances in the sections that follow.

Tech Appliance, In the REFIT dataset, desktop computers and televisions appear to be well recognized; the best classifiers have a Macro F1-Score over 0.7. While not as good, the Desktop Computer score on other datasets is in line with the quantity of time series offered. It makes sense since in longer load curves for smart meters, the pattern is obscured by other appliance activation signals, making it difficult for classifiers to identify.

Kitchen Appliance, Initially, it appears that identifying Kettle use is rather difficult, as all classifiers produced subpar results, with a Macro F1-Score of less than 0.45. Since a kettle only runs for brief periods of time, it makes sense that 30 minutes of collected data would not fully reflect its activity. In the EDF datasets, microwave and conventional ovens are not very well spotted. But because there is more data available for this scenario in REFIT, the detection score that the top two classifiers acquire is higher than 0.7. Lastly, the CER dataset shows that the Cooker is accurately spotted.

Washier Appliance, Classifiers using CER and EDF 2 datasets show good performance in recognizing dishwashers and tumble dryers. The smaller number of labeled examples provided for these circumstances explains the inferior performance shown with the EDF 1 dataset. Nevertheless, the lackluster REFIT scores for the three washing appliances are not attributable to the volume of time series data. We think that the reason for this low detection score is because the classifiers have a hard time telling these three devices apart since they are utilized in tandem and have similar activation patterns.

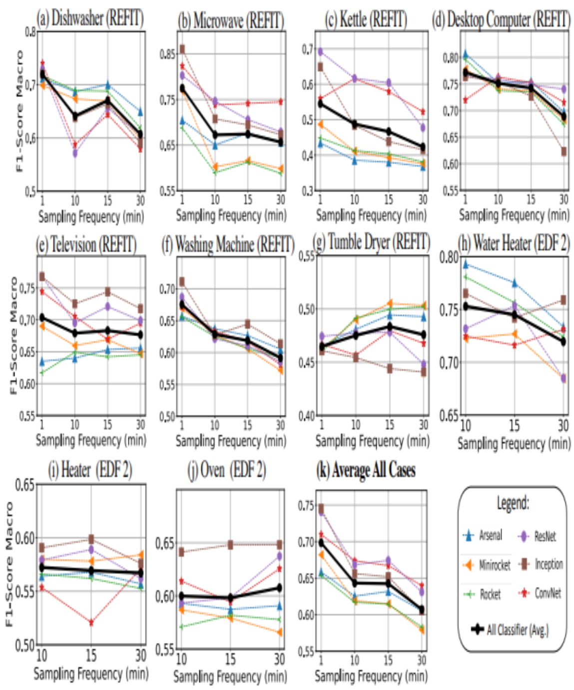

Figure No. 7 – Influence of sampling frequency on different appliance detection cases

We examine how the number of unique households affects the performance of the classifier. These studies show that when the smart meter data sampling frequency is relatively low, classifiers cannot efficiently learn the patterns of an appliance using only a limited number of households (which explains the poor findings provided in Section 5.1 for the UK-DALE dataset). Furthermore, we show that the quantity of data provided for each home is not as significant as the number of households in terms of training the machine learning models.

We contrasted the two training methods listed below: The models may be trained in two ways: (i) by randomly selecting a subset of the homes and using all the data from these houses, or (ii) by selecting all the houses and using a random fraction of the time series from each house. Using the REFIT dataset, we ran the tests on the appliance detection situations. Furthermore, we conducted the studies using 4 different sampling frequencies: 1 minute, 10 minutes, 15 minutes, and 30 minutes to account for the influence of the smart meter reading on these results.

Figure 6 presents an overview of the test's findings. The graphs display each classifier1's average performance for every sampling rate.

VII. CONCLUSIONS

In-depth analysis of the most recent time series classifiers used for appliance recognition in extremely low-frequency smart meter data is presented in this research. We utilize five distinct real datasets of extremely low-frequency power usage with diverse time series lengths to construct the first benchmark of time series classifiers for appliance identification. The findings show that existing time series classifiers perform inconsistently well; only long-running appliances can be reliably identified with 30 minutes of sampling data. However, the detection accuracy of tiny appliances may be significantly improved by employing 1 minute sample data. Additionally, the accuracy of deep learning-based classifiers has shown promise, especially for certain appliances. All things considered, this work offers significant assistance to energy providers, analysts, and practitioners in selecting the proper classifier for precisely identifying appliances in extremely low-frequency smart meter data.

REFERENCE

[1]. Adrian Albert and Ram Rajagopal. 2013. Smart Meter Driven Segmentation: What Your Consumption Says About You. IEEE Transactions on Power Systems 28, 4 (2013), 4019–4030. https://doi.org/10.1109/TPWRS.2013.2266122

[2]. Hunt Allcott. 2011. Social norms and energy conservation. Journal of Public Economics 95, 9 (2011), 1082–1095. https://doi.org/10.1016/j.jpubeco.2011.03.003 Special Issue: The Role of Firms in Tax Systems.

[3]. Muzaffer Aslan and Ebra Nur Zurel. 2022. An efficient hybrid model for appliances classification based on time series features. Energy and Buildings 266 (2022), 112087. https://doi.org/10.1016/j.enbuild.2022.112087

[4]. Anthony Bagnall, Aaron Bostrom, James Large, and Jason Lines. 2016. The Great Time Series Classification Bake Off: An Experimental Evaluation of Recently Proposed Algorithms. Extended Version. https://doi.org/10.48550/ARXIV.1602.01711

[5]. T. Cover and P. Hart. 1967. Nearest neighbor pattern classification. IEEE Transactions on Information Theory 13, 1 (1967), 21–27. https://doi.org/10.1109/TIT. 1967.1053964

[6]. G.W. Hart. 1992. Nonintrusive appliance load monitoring. Proc. IEEE 80, 12 (1992), 1870–1891. https://doi.org/10.1109/5.192069

[7]. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. https://doi.org/10.48550/ARXIV.1512.03385

[8]. Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep Learning for Time Series Classification: A Review. Data Min. Knowl. Discov. 33, 4 (jul 2019), 917–963. https://doi.org/10. 1007/s10618-019-00619-1

[9]. Matthias Kahl, Daniel Jorde, and Hans-Arno Jacobsen. 2022. Representation Learning for Appliance Recognition: A Comparison to Classical Machine Learning. https://doi.org/10.48550/ARXIV.2209.03759

[10]. J. Zico Kolter. 2011. REDD: A Public Data Set for Energy Disaggregation Research.

[11] Yu Liu, Congxiao Liu, Yiwen Shen, Xin Zhao, Shan Gao, and Xueliang Huang. 2021. Non-intrusive energy estimation using random forest based multi-label classification and integer linear programming. Energy Reports 7 (2021), 283–291. https://doi.org/10.1016/j.egyr.2021.08.045 2021 The 4th International Conference on Electrical Engineering and Green Energy.

[12] Keiron O'Shea and Ryan Nash. 2015. An Introduction to Convolutional Neural Networks. CoRR abs/1511.08458 (2015). arXiv:1511.08458 http://arxiv.org/abs/ 1511.08458

[13] ick Schäfer and Mikael Högqvist. 2012. SFA: A symbolic fourier approximation and index for similarity search in high dimensional datasets. ACM International Confer

[14] Bochao Zhao, Minxiang Ye, Lina Stankovic, and Vladimir Stankovic. 2020. Nonintrusive load disaggregation solutions for very low-rate smart meter data. Applied Energy 268 (2020), 114949.

[15] Chunyu Deng, Kehe Wu, and Binbin Wang. 2022. Residential Appliance Detection Using Attention-based Deep Convolutional Neural Network. CSEE Journal of Power and Energy Systems 8, 2 (2022), 621–633.