

Anomaly detection in small-scale industrial and household appliances

Niccolò Zangrando¹[0000-0002-4796-5649], Sergio Herrera¹[0000-0002-8903-0622],
Paraskevas Koukaras^{2,4}[0000-0002-1183-9878], Asimina
Dimara^{2,5}[0000-0001-9372-7070], Piero Fraternali¹[0000-0002-6945-2625], Stelios
Krinidis^{2,3}[0000-0002-9666-7023], Dimosthenis Ioannidis²[0000-0002-5747-2186],
Christos Tjortjis⁴[0000-0001-8263-9024], Christos-Nikolaos
Anagnostopoulos⁵[0000-0002-4496-275X], and Dimitrios
Tzovaras²[0000-0001-6915-6722]

¹ Politecnico di Milano, 20133, Milano, Italy

{niccolo.zangrando, sergioluis.herrera, piero.fraternali}@polimi.it

² Information Technologies Institute, Centre for Research & Technology, 57001,
Thessaloniki, Greece

{p.koukaras, adimara, krinidis, djoannid, c.tjortjis,
dimitrios.tzovaras}@iti.gr

³ Dept. of Management Science and Technology, International Hellenic University,
Kavala, Greece

⁴ School of Science and Technology, International Hellenic University, 57001,

⁵ Dept. of Cultural Technology and Communication, University of the Aegean,
Intelligent Systems Lab, Greece

canag@aegean.gr

Abstract. Anomaly detection is concerned with identifying rare events/ observations that differ substantially from the majority of the data. It is considered an important task in the energy sector to enable the identification of non-standard device conditions. The use of anomaly detection techniques in small-scale residential and industrial settings can provide useful insights about device health, maintenance requirements, and downtime, which in turn can lead to lower operating costs. There are numerous approaches for detecting anomalies in a range of application scenarios such as prescriptive appliance maintenance. This work reports on anomaly detection using a data set of fridge power consumption that operates on a near zero energy building scenario. We implement a variety of machine and deep learning algorithms and evaluate performances using multiple metrics. In the light of the present state of the art, the contribution of this work is the development of a inference pipeline that incorporates numerous methodologies and algorithms capable of producing high accuracy results for detecting appliance failures.

Keywords: Anomaly detection · Time series analysis · Machine learning · Deep learning.

1 Introduction

Predictive maintenance (PdM) aims to optimize the trade off between run-to-failure and periodic maintenance, by empowering manufacturers to improve the remaining useful life of their machines while at the same time avoiding unplanned downtime and decreasing planned downtime. At the core of PdM lies Anomaly Detection task (AD) whose primary focus is to find anomalies in the operation of working equipment at early stages and alert the supervisor to carry out maintenance activity. In addition, anomaly detection may stand as a core component for prescriptive maintenance (PsM) being a type of maintenance that gains popularity lately and poses as the evolution of PdM.

In recent years, AD has proved beneficial in different application scenarios and has acquired a prominent stance in the unsupervised machine learning research. AD finds use in different fields such as healthcare, where it applies to the analysis of clinical images [1] and of ECG data [2], in the cybersecurity field, where it is applied for malware identification [3] and in the energy field. In this latter area AD may be combined with energy load forecasting to improve accuracy [4], or integrated as a component for detecting non nominal energy fluctuations for enhancing decision making in energy transfer between micro-grids [5]. AD has also been successfully employed for banking fraud detection [6].

However, the lack of public data sets for small scale industrial devices and household appliances makes it difficult to understand the applicability of the anomaly detection methods used for large industrial devices in other contexts such as domestic appliances or common service system in residential buildings (e.g. heating or air-conditioning systems).

This paper summarizes and evaluates the current status of the art on anomaly detection approaches with a focus on their applicability to the context of household appliances. The primary objective is to provide a comprehensive survey of the most important contributions, developments, and experimental approaches in the field. By implementing some of them for the specific use case of a fridge energy behavior, we assess the most relevant techniques and highlight the outstanding research problems for the specific target of house appliances and residential building systems.

The rest of the article is organised as follows: Section 2 overviews the state of the art in anomaly detection. Section 3 surveys the research design including the used data sets, the identified methods and the most common evaluation metrics. Section 4 summarizes the obtained results. Finally, Section 5 provides the conclusions.

2 Related Work

AD refers to the identification of rare events or observations which significantly deviate from the majority of the data [7].

This task spans different disciplines and is primarily applied in industrial IoT applications where data are collected as time series [8]. Time series data

sets collect observations sampled at different times: recording can be continuous, when data are collected continuously in a given interval, or discrete, when data are recorded at set time intervals [9]. In the literature, anomalies for time series are classified into three different types [8,10,11]:

- *PointAnomaly*: represents data that abruptly deviates from the normal ones. With these anomaly types, the time series usually returns to its previous normal state within a very short time of only a few observations.
- *ContextualAnomaly*: represents a data point or sequence observed over a short time but does not deviate from the normal range of values of the time series.
- *CollectiveAnomaly*: collection of observations that are anomalous with respect to the rest of the data. Individual observations within a collective anomaly may or may not be anomalous, but considered as a group they appear suspicious.

Based on the number of observations at each timestamp, the time series can be univariate or multivariate. Univariate time series log values generated by a single sensor, whereas multivariate time series record signals from multiple sensors simultaneously. Depending on the nature of the time series different algorithms and approaches have been applied, ranging from more classical and statistical techniques to the deep learning ones. Moreover, the different approaches can be categorized into three different types depending on their anomaly identification criteria as follow:

- *Reconstruction error*: this criterion applies to all those models whose objective is to generate an output as close as possible to the input. An example is the Autoencoder-based models, which reconstruct input data by extracting features from them. Anomalous data are identified based on the residuals between the input and the generated data: the higher the difference, the higher the probability of an anomaly.
- *Prediction error*: prediction models are used to identify anomalies based on the difference between the predicted value and the expected one. Like the models based on input reconstruction, the larger the residual, the higher the probability of anomalous data.
- *Dissimilarity*: this criterion consists of identifying outliers based on the difference between the input data and the distribution or clusters obtained from the analysis of normal data.

Statistical methods based on regressive models are used for the identification of the outliers in univariate time series, such as Autoregressive Moving Average models (ARMA) used for stationary time series [12] (i.e. time series whose properties do not depend on the time at which the series is observed), or Autoregressive Integrated Moving Average models (ARIMA) preferred for nonstationary time series [13,14].

By exploiting a sliding window on the input data clustering methods have been applied for anomaly detection on time series, such as K-Means clustering

[15], DBSCAN [16] and Local Outlier Factor (LOF) [17]. Also machine learning based approaches have been employed such as Isolation Forest [18] and One-Class Support Vector Machine [19]. All these techniques rely on a dissimilarity criterion to identify anomalies.

With the advent of Deep Learning (DL), several algorithms have been applied to time series to identify anomalies. The results highlight that DL approaches overcome the difficulties of the more classical techniques [11]. Since time series data are related to a temporal context, the Recurrent Neural Networks (RNNs) [20] is one of the most widely used approaches. Due to the vanishing or exploding gradient problem that limits the ability of the network to model long temporal relationships between data, two variants are preferred, Long-Short Term Memory (LSTM) [21] and Gated Recurrent Unit (GRU) [22]. RNN-based techniques are used in two different ways for anomaly detection. The first consists in the prediction error criterion [23,24,25], and the second one is based on the reconstruction error criterion [26,27,28]. CNN-based methods have also been applied to time series analysis, despite they are not designed to identify temporal relationships they still manage to extract meaningful information in the data sequences. The methodologies applied to identify anomalies are, as for RNNs, based on the prediction error [29,30] or the input reconstruction error [31].

Although some data sets with the consumption of household appliances have been published, there are few works concerning anomaly detection in this field, also due to the absence of labelled anomalous data acting as ground truth. In [32], for example, the authors manually analysed and annotated potential anomalies concerning the energy consumption of appliances in the REFIT data set [33]; in [34], instead, the authors have annotated the anomalies in the AMPds2 data set [35] through an ensemble method and then have evaluated their LSTM-autoencoder implementation.

In this paper, we compare the performance of nine different anomaly detection techniques, summarized in Table 1, using a data set of fridge power consumption samples.

Table 1. The AD techniques assessed in this paper and their anomaly identification criterion.

Technique	Anomaly Criterion
Local Outlier Factor (LOF)	dissimilarity
One-Class SVM	dissimilarity
Isolation Forest	dissimilarity
CNN	prediction error
GRU	prediction error
LSTM	prediction error
CNN-Autoencoder	reconstruction error
GRU-Autoencoder	reconstruction error
LSTM-Autoencoder	reconstruction error

3 Research Design

3.1 Experimental Data set

The CERTH data set represents the power consumption of a fridge in a household over a 4 month period, from 10th July 2019 until 3rd November 2019. The data were collected every minute, sampling in total 164,795 consumption values. The raw data were then analyzed and resampled every 10 minutes to remove sensor noise, obtaining the regular power consumption shown in Figure 1 and reducing the total number of observations to 16,710.

To be able to evaluate and compare the performance of the different algorithms, we have manually analyzed the data set and we have annotated all the potential anomalies, by following the criterion used in [32]: data are flagged anomalous if the appliance’s consumption has been found significantly different from its historical normal consumption. This analysis of the data set revealed two recurrent anomalous behaviors:

- An instant increment in power consumption (point anomaly).
- A continuous power consumption over time (contextual anomaly).

As shown in Figure 2, these potential anomalies can occur simultaneously.

For the evaluation, the data until the 30-09-2019 has been cleaned from the anomalies and used for training, and the data from 01-10-2019 to 31-10-2019 has been used for testing the performance of the different algorithms.

Table 2 summarizes the CERTH data set information, regarding the anomalous data and the train-test split.

3.2 Methods and Algorithms

Based on the main criterion of anomaly identification, we implemented nine different techniques:




Fig. 1. Fridge normal consumption pattern.




Fig. 2. Fridge potential anomaly consumption.

Table 2. CERTH data set: data points and anomalies.

	Total data	Anomaly data	% anomalies
Total	16710	1737	10.39%
Train set	10460	not needed	not needed
Test set	4464	488	10.93%

Local Outlier Factor (LOF) clustering algorithm based on the identification of the nearest neighbors and local outliers. We have used a sliding window size of 10 and a k value of 400.

One-Class Support Vector Machine support vector machine used for novelty detection. In the implementation we have used a sliding window length of 5, the RBF kernel with a gamma value of 0.001 and a nu value of 0.025.

Isolation Forest ensemble method that creates different binary trees isolating data points. Anomaly points are more likely to be isolated and closer to the root of an isolation tree. We have used a sliding window length of 5 and 100 trees in the ensemble.

Convolutional Neural Networks mainly used for computer vision tasks. We have used a window size length of 10, a convolutional block with a ReLU activation function, with 2, 4, 8 filters and the kernel size of 2, a max pooling layer and a fully connected layer with 50 neurons. The network has been trained for 300 epochs with a 64 batch size.

Gated Recurrent Unit RNN variant network. We have used a sliding window size of 10, 2 GRU layers with 8 hidden layers and a dropout of 0.2 respectively. The network was trained for 300 epochs with a batch size of 64.

Long Short Term Memory Networks RNN variant network. We have used a sliding window size of 10, 2 LSTM layers with 8 hidden layers respectively and a dropout of 0.2. The network was trained for 300 epochs with a batch size of 64.

CNN-autoencoder hybrid implementation with autoencoder and CNN network. We used a sliding window size of 12 and, for the encoding-decoding phase two convolutional block with 16 and 8 filters and a kernel size of 2. The network has been trained for 300 epochs with a batch size of 64.

GRU-autoencoder hybrid implementation with autoencoder and GRU network. We used a sliding window size of 10 and, for the encoding-decoding phase two GRU layers with 16 and 8 hidden layers. The network has been trained for 300 epochs with a batch size of 64.

LSTM-autoencoder hybrid implementation with autoencoder and LSTM network. We used a sliding window size of 10 and, for the encoding-decoding phase two LSTM layers with 16 and 8 hidden layers. The network has been trained for 300 epochs with a batch size of 64.

3.3 Evaluation Metrics

In order to compare the implemented methods, we have evaluated them with the most widely used machine learning metrics, based on the true positives (TP), false positives (FP), false negatives (FN), and true negatives (TN): precision, recall, F1 score, false alarm rate (FAR) and miss alarm rate (MAR).

$$precision = \frac{TP}{TP + FP} \quad (1)$$

$$recall = \frac{TP}{TP + FN} \quad (2)$$

$$F1 = 2 * \frac{precision * recall}{precision + recall} \quad (3)$$

$$FAR = \frac{FP}{FP + TN} \quad (4)$$

$$MAR = \frac{FN}{FN + TP} \quad (5)$$

4 Results

Table 3 summarizes the results. The metrics described in the previous section (i.e., precision, recall, F1 score) are used to compare the performances of the tested algorithms. One-Class SVM and CNN have the highest precision (i.e., 0.76) while GRU-Autoencoder has the lowest (i.e., 0.62). Local Outlier Factor (LOF) showcases the highest recall (i.e., 0.92) and CNN-Autoencoder the lowest (i.e., 0.58).

Moreover, Isolation Forest attains the highest F1 score (i.e., 0.78), while the lowest one (i.e., 0.64) is achieved by CNN-Autoencoder. It may be observed that precision is only moderately good for all the tested algorithms, varying from 0.62 to 0.76. On the contrary, recall exhibits better performances and ranges from 0.58 to 0.92. The F1 score performs similarly to the precision metrics. Table 3 highlights the best algorithm for each group based on the adopted anomaly criterion.

From a quantitative evaluation standpoint, all the algorithms seem to achieve comparable performances. But if we analyze their behavior from a qualitative perspective, two observations emerge.

The first one is related to the CNN-Autoencoder, which has a high precision compared to the other approaches based on input reconstruction, but has a recall value rather low with respect to all the other algorithms. As shown in Figure 3, this is due to the lack of identification of the anomaly in case the device stops working (i.e. the power consumption is zero), which can be a problem especially for household appliances such as the one analysed.

The second aspect refers to the overall performances of the implemented methods, which achieve good results. However, they have been evaluated considering only each single labelled anomalous point separately, but, as shown in Figure 4, all algorithms correctly identified almost all the time windows in which the anomalies occur. The precise identification of the starting or ending point of them is the main difference among the different techniques.

Table 3. Results summary of implemented methods. For each one is specified the anomaly criterion (Dissimilarity, Prediction, Reconstruction).

Technique	Miss Alarm Rate	False Alarm Rate	Precision	Recall	F1 score
One-Class SVM (D)	0.24	0.03	<u>0.76</u>	0.76	0.76
Local Outlier Factor (D)	0.08	0.06	0.66	<u>0.92</u>	0.77
Isolation Forest (D)	0.16	0.04	0.73	0.84	0.78
CNN (P)	0.4	0.02	<u>0.76</u>	0.60	0.67
GRU (P)	0.38	0.03	0.73	0.62	0.67
LSTM (P)	0.38	0.03	0.73	<u>0.62</u>	0.67
CNN-Autoencoder (R)	0.42	0.03	<u>0.72</u>	0.58	0.64
GRU-Autoencoder (R)	0.17	0.06	0.62	<u>0.81</u>	0.70
LSTM-Autoencoder (R)	0.24	0.05	0.67	0.76	0.71

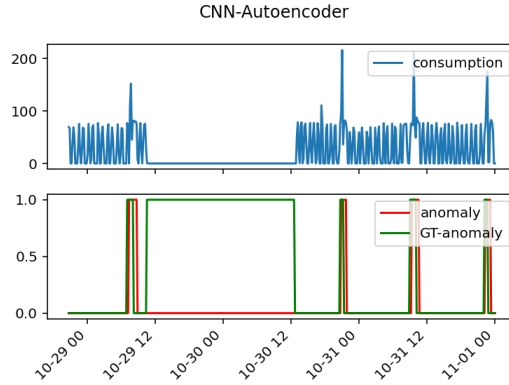


Fig. 3. CNN-Autoencoder anomaly identification lack example. The green line refers to the ground truth, while the red one to the model predictions.

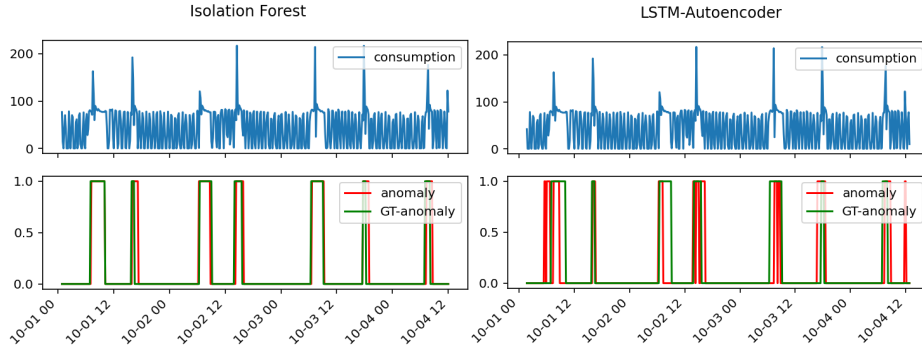


Fig. 4. Qualitative overall performances. The green line refers to the ground truth, while the red one to the model predictions. On the left are shown the Isolation Forest anomaly predictions, while on the right the LSTM-Autoencoder ones.

5 Conclusion

The identification of anomalies focuses on recognizing unusual events/observations that deviate significantly from the rest of the data. Being able to recognize non-standard device operation is seen as a significant responsibility in the energy industry. AD in small-scale residential and industrial settings can benefit the insight into the device health, the maintenance requirements, and the experienced downtime and thus has the potential to reduce maintenance costs significantly. Several alternative techniques have been proposed and evaluated for AD in a variety of application contexts, also in the prescriptive maintenance scenario.

In this work, we applied alternative anomaly detection methods to data collected from a fridge power usage in a real-world zero-energy building prototype. We have implemented several machine and deep learning techniques and assessed their respective performances using multiple metrics. The primary contribution

of this study is to compare the extent at which the various approaches and algorithms are capable of delivering high accuracy results for identifying device/-machine/appliance faults. We obtained promising results with several methods, among which Isolation Forest and LSTM-Autoencoder algorithms stand out.

Limitations of this work may be attributed to the fact that the evaluation of methods and algorithms takes place utilizing a single appliance (fridge). Also, the historical data of this device are limited to 4 months, something that may have negative impact in the process of training the deep learning algorithms. More data should be employed for a more thorough testing phase but also for producing more reliable and generic results.

All in all, anomaly detection seeks to identify anomalous behavior in data observations or highlight data outliers. In terms of appliance or device maintenance it seeks to identify non nominal operation generating prospects for preventing various types of failure completely. In the future we aim to expand this work by researching on the following points.

- Implement the described approach as a stand-alone component being able to function with any data input. This will allow this work to be incorporated as part of an analytics engine or any energy related framework [36].
- Continue tracking the state of the art in anomaly detection focusing on maintenance for proactive buildings in the domain of households or small industrial setups.
- Expand the evaluation phase with more data sets including small scale industrial units or household clusters containing more devices.

Acknowledgements This work is supported by the project PRECEPT - A novel decentralized edge-enabled PREsCRIPTivE and ProacTive framework for increased energy efficiency and well-being in residential buildings funded by the EU H2020 Programme, grant agreement no. 958284.

References

1. Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019.
2. Sucheta Chauhan and Lovekesh Vig. Anomaly detection in ecg time signals via deep long short-term memory networks. In *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–7. IEEE, 2015.
3. Borja Sanz, Igor Santos, Xabier Ugarte-Pedrero, Carlos Laorden, Javier Nieves, and Pablo García Bringas. Anomaly detection using string analysis for android malware detection. In *International Joint Conference SOCO'13-CISIS'13-ICEUTE'13*, pages 469–478. Springer, 2014.
4. Paraskevas Koukaras, Napoleon Bezas, Paschalis Gkaidatzis, Dimosthenis Ioannidis, Dimitrios Tzovaras, and Christos Tjortjis. Introducing a novel approach in one-step ahead energy load forecasting. *Sustainable Computing: Informatics and Systems*, 32:100616, 2021.

5. Paraskevas Koukaras, Christos Tjortjis, Paschalis Gkaidatzis, Napoleon Bezas, Dimosthenis Ioannidis, and Dimitrios Tzovaras. An interdisciplinary approach on efficient virtual microgrid to virtual microgrid energy balancing incorporating data preprocessing techniques. *Computing*, pages 1–42, 2021.
6. Xun Zhou, Sicong Cheng, Meng Zhu, Chengkun Guo, Sida Zhou, Peng Xu, Zhenghua Xue, and Weishi Zhang. A state of the art survey of data mining-based fraud detection and credit scoring. In *MATEC Web of Conferences*, volume 189, page 03002. EDP Sciences, 2018.
7. Douglas M Hawkins. *Identification of outliers*, volume 11. Springer, 1980.
8. Andrew A Cook, Göksel Mısırlı, and Zhong Fan. Anomaly detection for iot time-series data: A survey. *IEEE Internet of Things Journal*, 7(7):6481–6494, 2019.
9. Peter J Brockwell and Richard A Davis. *Time series: theory and methods*. Springer Science & Business Media, 2009.
10. Raghavendra Chalapathy and Sanjay Chawla. Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407*, 2019.
11. Kukjin Choi, Jihun Yi, Changhwa Park, and Sungroh Yoon. Deep learning for anomaly detection in time-series data: Review, analysis, and guidelines. *IEEE Access*, 2021.
12. Brandon Pincombe. Anomaly detection in time series of graphs using arma processes. *Asor Bulletin*, 24(4):2, 2005.
13. H Zare Moayedi and MA Masnadi-Shirazi. Arima model for network traffic prediction and anomaly detection. In *2008 international symposium on information technology*, volume 4, pages 1–6. IEEE, 2008.
14. Asrul H Yaacob, Ian KT Tan, Su Fong Chien, and Hon Khi Tan. Arima based network anomaly detection. In *2010 Second International Conference on Communication Software and Networks*, pages 205–209. IEEE, 2010.
15. Tsuyoshi Idé. Why does subsequence time-series clustering produce sine waves? In *European conference on principles of data mining and knowledge discovery*, pages 211–222. Springer, 2006.
16. Mete Çelik, Filiz Dadaşer-Çelik, and Ahmet Şakir Dokuz. Anomaly detection in temperature data using dbscan algorithm. In *2011 international symposium on innovations in intelligent systems and applications*, pages 91–95. IEEE, 2011.
17. Stefan Oehmcke, Oliver Zielinski, and Oliver Kramer. Event detection in marine time series data. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 279–286. Springer, 2015.
18. Zhiguo Ding and Minrui Fei. An anomaly detection approach based on isolation forest algorithm for streaming data using sliding window. *IFAC Proceedings Volumes*, 46(20):12–17, 2013.
19. Junshui Ma and Simon Perkins. Time-series novelty detection using one-class support vector machines. In *Proceedings of the International Joint Conference on Neural Networks, 2003.*, volume 3, pages 1741–1745. IEEE, 2003.
20. Larry R Medsker and LC Jain. Recurrent neural networks. *Design and Applications*, 5:64–67, 2001.
21. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
22. Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.

23. Pankaj Malhotra, Lovekesh Vig, Gautam Shroff, Puneet Agarwal, et al. Long short term memory networks for anomaly detection in time series. In *Proceedings*, volume 89, pages 89–94, 2015.
24. Adrian Taylor, Sylvain Leblanc, and Nathalie Japkowicz. Anomaly detection in automobile control network data with long short-term memory networks. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 130–139. IEEE, 2016.
25. Mohsin Munir, Shoaib Ahmed Siddiqui, Andreas Dengel, and Sheraz Ahmed. Deepant: A deep learning approach for unsupervised anomaly detection in time series. *Ieee Access*, 7:1991–2005, 2018.
26. Ruei-Jie Hsieh, Jerry Chou, and Chih-Hsiang Ho. Unsupervised online anomaly detection on multivariate sensing time series data for smart manufacturing. In *2019 IEEE 12th Conference on Service-Oriented Computing and Applications (SOCA)*, pages 90–97. IEEE, 2019.
27. Daehyung Park, Yuuna Hoshi, and Charles C Kemp. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. *IEEE Robotics and Automation Letters*, 3(3):1544–1551, 2018.
28. Yifan Guo, Weixian Liao, Qianlong Wang, Lixing Yu, Tianxi Ji, and Pan Li. Multidimensional time series anomaly detection: A gru-based gaussian mixture variational autoencoder approach. In *Asian Conference on Machine Learning*, pages 97–112. PMLR, 2018.
29. Tailai Wen and Roy Keyes. Time series anomaly detection using convolutional neural networks and transfer learning. *arXiv preprint arXiv:1905.13628*, 2019.
30. Yeji Choi, Hyunki Lim, Heeseung Choi, and Ig-Jae Kim. Gan-based anomaly detection and localization of multivariate time series data for power plant. In *2020 IEEE international conference on big data and smart computing (BigComp)*, pages 71–74. IEEE, 2020.
31. FU Xuyun, LUO Hui, Shisheng Zhong, and Lin Lin. Aircraft engine fault detection based on grouped convolutional denoising autoencoders. *Chinese Journal of Aeronautics*, 32(2):296–307, 2019.
32. Haroon Rashid, Vladimir Stankovic, Lina Stankovic, and Pushpendra Singh. Evaluation of non-intrusive load monitoring algorithms for appliance-level anomaly detection. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8325–8329. IEEE, 2019.
33. David Murray, Lina Stankovic, and Vladimir Stankovic. An electrical load measurements dataset of united kingdom households from a two-year longitudinal study. *Scientific data*, 4(1):1–12, 2017.
34. Yu Weng, Ning Zhang, and Chunlei Xia. Multi-agent-based unsupervised detection of energy consumption anomalies on smart campus. *IEEE Access*, 7:2169–2178, 2018.
35. Stephen Makonin, Bradley Ellert, Ivan V Bajić, and Fred Popowich. Electricity, water, and natural gas consumption of a residential house in canada from 2012 to 2014. *Scientific data*, 3(1):1–12, 2016.
36. Paraskevas Koukaras, Paschalis Gkaidatzis, Napoleon Bezas, Tommaso Bragatto, Federico Carere, Francesca Santori, Marcel Antal, Dimosthenis Ioannidis, Christos Tjortjis, and Dimitrios Tzovaras. A tri-layer optimization framework for day-ahead energy scheduling based on cost and discomfort minimization. *Energies*, 14(12), 2021.