



Unsupervised detection and open-set classification of fast-ramped flexibility activation events

Müller, Nils; Heinrich, Carsten; Heussen, Kai; Bindner, Henrik W.

Published in:
Applied Energy

Link to article, DOI:
[10.1016/j.apenergy.2022.118647](https://doi.org/10.1016/j.apenergy.2022.118647)

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Müller, N., Heinrich, C., Heussen, K., & Bindner, H. W. (2022). Unsupervised detection and open-set classification of fast-ramped flexibility activation events. *Applied Energy*, 312, Article 118647. <https://doi.org/10.1016/j.apenergy.2022.118647>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



Unsupervised detection and open-set classification of fast-ramped flexibility activation events

Nils Müller ^{*}, Carsten Heinrich, Kai Heussen, Henrik W. Bindner

Center for Electric Power and Energy, Technical University of Denmark, Elektrovej, Building 325, 2800 Kgs. Lyngby, Denmark

ARTICLE INFO

Keywords:

Flexibility
Event detection
Open-set classification
Active distribution networks
Machine learning
Electrification

ABSTRACT

The continuous electrification of the mobility and heating sectors adds much-needed flexibility to the power system. However, flexibility utilization also introduces new challenges to distribution system operators (DSOs), who need mechanisms to supervise flexibility activations and monitor their effect on distribution network operation. Flexibility activations can be broadly categorized to those originating from electricity markets and those initiated by the DSO to avoid constraint violations. Coinciding electricity market driven flexibility activations may cause voltage quality or temporary overloading issues, and the failure of flexibility activations initiated by the DSO might leave critical grid states unresolved. This work proposes a novel data processing pipeline for automated real-time identification of fast-ramped flexibility activation events. Its practical value is twofold: (i) potentially critical flexibility activations originating from electricity markets can be detected by the DSO at an early stage, and (ii) successful activation of DSO-requested flexibility can be verified by the operator. In both cases the increased awareness would allow the DSO to take counteractions to avoid potentially critical grid situations. The proposed pipeline combines techniques from unsupervised detection and open-set classification. For both building blocks feasibility is systematically evaluated and proofed on real load and flexibility activation data.

1. Introduction

Renewable electricity and electrification are key pillars of global efforts to eliminate fossil fuels in the energy supply. The European goal of carbon neutrality in 2050 is reported to require increased shares of renewable energy and continued electrification of the mobility and heating sectors [1]. This trend will further increase uncertainty and volatility in distribution networks (DNs). Thus, active management of DN's based on the emerging smart solutions for monitoring, control, and communication is seen as a requirement for distribution system operators (DSOs) [2]. With the improving capability and affordability of information and communication technology (ICT) the implicit or explicit utilization of local consumption flexibility, commonly referred to as demand response, is becoming more attractive. By requesting a load deviation of flexible units during a certain time period such as peak hours, referred to as flexibility activation (FA) event, DSOs can use local flexibility to avoid or postpone grid reinforcements [3]. However, applying local end user flexibility for mitigation of potentially critical network states makes grid operation and security partly dependent on the reliability of third parties. Thus, real-time detection of FA events is desirable for DSOs to verify successful activation, and initiate other measures in case of a failure. Moreover, FA events are not exclusive

to DSOs, due to activations originating from electricity markets. On the one hand, balance responsible parties could request and activate flexibility for portfolio optimization. On the other hand, controllable heat pumps and electric vehicles may simultaneously react to price signals with a sudden change of power consumption [4,5]. As a result, DSOs are not aware of all FA events affecting their network. If not detected at an early stage, such fast-ramped FA events could trigger transformer or line protections due to high coincidence [6,7] or load rebound effects [8–11]. The resulting disconnection of customers could lead to high social and financial cost.

The increasing deployment of measuring devices, such as micro phasor measurement units and smart meters (SMs), increase observability of DN's and thus provide the data basis for identification of FA events. However, real-time identification of FA events is challenged by different practical problems. The infrequent occurrence of FA events limits the available data required for implementation of supervised detection methods. Moreover, the operation of active DN's is influenced by a variety of rare or even unseen event classes such as line faults, topology changes or communication failures [12,13]. Thus, FA event identification also requires differentiation between unknown event classes and FA events. This questions the use of widely applied closed-set

^{*} Corresponding author.

E-mail address: nilmu@elektro.dtu.dk (N. Müller).

<https://doi.org/10.1016/j.apenergy.2022.118647>

Received 30 November 2021; Accepted 24 January 2022

Available online 24 February 2022

0306-2619/© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Schematic overview of the proposed EIP for FA events in active DNs.

(CS) classifiers that will falsely classify unknown event classes, due to their inability of rejecting these. Another challenge for real-time FA event identification is seen in the central data processing, e.g. via cloud computing. Already today the integration of SM data in real-time power system operation is limited by communication instead of meter-recording capability [14]. Upgrading communication networks entails a high economic burden. Moreover, long communication paths increase the possibilities for false data injection attacks and other fraudulent modification of data [15,16]. One approach to overcome the drawbacks of central data processing is seen in a distributed event identification architecture based on edge or fog computing [17].

The described challenges set specific requirements to the approach and implemented techniques for FA event identification. However, a formulation of requirements is missing, complicating the development of appropriate strategies and methods for identifying FA events.

In this work, a novel event identification pipeline (EIP) for fast-ramped FA events is proposed. A schematic overview of the proposed pipeline is depicted in Fig. 1. The data processing pipeline is based on unsupervised detection and open-set (OS) classification algorithms, suitable for application in a distributed event identification architecture. The scheme and algorithm selection are based on a thorough requirements analysis. The core contributions of this work are the systematic selection of processing algorithms and their evaluation on real load and FA event data.

1.1. Related work

To the extent of the authors' knowledge, this is the first work on detection and classification of FA events. Works on thematically-related topics are presented first, followed by a presentation of methodologically-related works. In both cases literature on detection and classification is discussed separately.

1.1.1. Thematically-related works

A frequently studied topic in power system literature is unsupervised anomaly detection in energy time series data. To detect anomalies, most works train models predicting normal behavior. A data point is declared an anomaly if deviation between model prediction and ground truth exceeds a predefined threshold. Various models such as variational autoencoder [18], hierarchical temporal memory (HTM) [19], autoregressive integrated moving average (ARIMA) or long short-term memory [20] are applied. None of these works consider FA events as anomaly. Moreover, most works assume anomaly-free training data for learning of the normal behavior. Some works exist on flexibility detection on building or device level, which try to quantify the flexible load potential in load data of individual devices or buildings [21,22]. Although the name suggests similarity, the problem under investigation is different to the present work, as these works actually investigate flexibility potential identification.

The topic of event classification in DNs has been studied intensely [23]. Most works investigate the multi-class CS classification problem. Literature considering OS classification in a power system context is rare. Lazzaretti et al. [24] first applied one-class classifiers for automatic oscillography classification under existence of unseen event classes in two different approaches. The first one considers a single one-class classifier for modeling the boundary of multiple known classes. A separate multi-class classifier is applied to differentiate between the

known classes. In the second approach each class is modeled with a separate one-class classifier. In the following years other works considered one-class classifiers for detection of new classes [25,26]. Although these methods in general can be used for OS classification, the different problem setting results in a comparatively low detection performance [27, 28]. To the best of the authors' knowledge, the literature provides no work applying classifiers inherently made for the OS problem to event classification in active DNs. With respect to the proposed EIP, some works on event classification exist that assume an upstream detection step [29,30]. However, none of these works describe how input samples for the classifier are generated from detector results, rather examining event classification for existing samples.

1.1.2. Methodologically-related works

Similar to literature on anomaly detection in energy time series data, multiple works propose forecasting-based unsupervised anomaly detection [31–33]. In most cases, the euclidean distance between point forecast and ground truth is used to flag anomalies based on a defined threshold. In [34], the authors propose the use of a convolutional neural network (CNN) as the time series forecaster. According to the authors, the proposed method can be trained on comparatively few training data and without removing anomalies from the training dataset. A novel approach on anomaly detection in time series data is proposed in [35]. The authors introduce the use of the spectral residual (SR) algorithm from saliency detection in computer vision for unsupervised anomaly detection in time series.

In contrast to the traditional CS classification problem, less literature exists on OS classification. Scheirer et al. [36] first formalized the OS classification problem and proposed the 1-vs-Set machine as a preliminary solution. Since then various methods such as distance-based [37], margin distribution-based [38] or generation-based [39] OS classifiers were proposed. In [27], the authors provide a systematic categorization of OS classification techniques, and compare a number of OS classifiers on popular benchmark datasets.

1.2. Contribution and paper structure

The main contributions of this work are as follows:

- Proposal of a novel EIP based on unsupervised detection and OS classification.
- First work on detection and classification of FA events.
- First application of an OS classifier to events in active DNs.
- Introduction of a new performance metric for the evaluation of real-time detection of FA events.
- Systematic demonstration and evaluation of unsupervised detection and OS classification of fast-ramped FA events as the main building blocks of the proposed pipeline based on real load and FA event data.

The remainder of the paper is structured as follows: In Section 2 requirements for FA event identification in active DNs are evaluated and strategies are proposed. Section 3 provides a description of models and methods. In Section 4 the experimental setup is presented, including the dataset under investigation, data preparation for model development and evaluation, and applied performance metrics. In Section 5 results are presented and discussed followed by a conclusion and a view on future work in Section 6.

2. Requirements and strategies for FA event identification

Identifying FA events in active DNs comes with specific requirements not only concerning the general approach, but also the implemented algorithms. Additional requirements at algorithmic level are introduced by the consideration of event identification based on a distributed architecture. In the following, the identified requirements are presented. Based on the requirements analysis, the concept of the proposed EIP as well as the selection of specific models for the main building blocks of the pipeline are justified. A detailed explanation of the implemented models and the proposed pipeline follows in Section 3.

The problem is limited to fast-ramped load reduction and load increase FA events with a length of up to 3 hours. Moreover, the aggregated active power load profile is assumed to represent averaged active power measurements on secondary substation level or aggregated SM data collected in a data hub on neighborhood level. In both scenarios a large low-voltage feeder is considered. The core of the concept is the separation of the event identification task into an unsupervised event detection and a supervised classification task, resulting in the proposed EIP. Besides the two main building blocks, an event sampler is required to prepare event observations for the classifier based on the results of the event detector.

2.1. Real-time identification

A key requirement for identification of FA events is real-time capability. Real-time identification allows DSOs to take immediate counteractions in cases where FAs could result in critical situations, such as congestions and under or over-voltages. Supervised detection or classification of time series events usually requires as input the entire time series sample [40]. For real-time event identification this becomes a fundamental problem. Existing early classification techniques come at the cost of decreased accuracy [40], and are not applicable to an OS classification problem. In the proposed pipeline the problem of prediction delay is addressed by separating event identification into two consecutive steps. The use of an unsupervised, point-wise event detector allows for immediate flagging of abnormal data points in real-time. Although this cannot solve the intrinsic problem of supervised classification being dependent on multiple data points of an event, information extraction is improved. Instead of identifying an event at the end of its occurrence, with the presented EIP DSOs will immediately be aware of the existence of a deviation from normal operation, followed by an ex-post classification of the event.

2.2. Model development based on limited and partly-labeled training data

FAs in active DNs constitute rare events. Thus, comprehensive datasets of FA events for supervised methods are difficult to obtain. The heterogeneity of DNs and flexibility portfolios brings additional challenges for acquiring datasets, since characteristics of FAs will differ for different networks. In contrast, unsupervised event detection does not require datasets of FA events. Instead, most works on unsupervised event detection in energy time series data, such as [18], assume event-free training data to learn a representation of the normal behavior. However, existing training data will most likely contain events, since manual removing is a time-consuming and impractical process [41] and DSOs might not be aware of all FAs (see Section 1). In the proposed pipeline a persistence forecast-based detector is considered, which is not dependent on event-free training data. By applying an unsupervised detector with no demand for event-free training data the dependency on labeled training data is reduced to the classification step. Therefore, compared to an one-step event identification approach, the proposed pipeline maintains event detection capability also in scenarios without labeled training data, maximizing information extraction.

2.3. Lightweight models for event identification

Distributed event identification in an edge or fog computing scheme requires models and methods to be lightweight. The vast number and limited processing power of edge devices as well as the continuously growing amount of data sets time and resource constraints to the development, operation and maintenance of models and methods. Therefore, a key requirement for FA event identification is seen in keeping computational and maintenance efforts, such as periodical re-training, at a minimum. This is considered for the proposed concept in several ways. First of all, the proposed pipeline entirely works with delta encoded data. Delta encoding is a technique for data compression based on differencing sequential data, reducing data communication and storage load [42]. By working with differenced data the proposed pipeline can directly be applied in a system which uses delta encoding for data compression. Both the detection and classification model within the pipeline only retrieve features from the univariate load time series. No additional information such as weather or market price data are considered, reducing the requirement for data communication and the dimensionality of the detection and classification problem. The use of a simple persistence forecast keeps size and computational effort of the proposed detector at a minimum, and avoids the need for frequent re-training. In the proposed pipeline the classifier only gets activated if the detector has detected a deviation from normal behavior above a predefined threshold. This event-triggered scheme avoids continuous running of the classifier which reduces the required processing power.

For OS classification the extreme value machine (EVM) model is selected as it comes with several features, making it a comparatively lightweight classifier [38]. EVM is capable of incremental learning which allows for efficient model updating without time and computationally intensive re-training. Moreover, the model reduction strategy of EVM discards redundant data points within a class of training points, allowing for limitation of model size and classification time as dataset size increases.

2.4. Handling multiple and unknown event classes

In active DNs a large variety of events with various backgrounds, such as faults and switching actions, can occur. While in this work the detection performance is evaluated on the basis of fast-ramped FA events, in principle an *unsupervised* detector also allows for detection of other fast-ramped events.

Although traditional CS classifiers can differentiate between multiple known event classes, introducing new unknown classes will lower classification performance drastically, since observations of unknown classes are wrongly assigned to one of the classes the classifier was trained on [27]. Given that many event classes only occur rarely and new ones might emerge, e.g. due to changes in grid topology, assuming that training data includes sufficient observations to describe all existing events is considered an unrealistic assumption. An important requirement for FA event identification is therefore seen in the capability to differentiate between FA events and other event classes by either recognizing known or rejecting unknown event classes. For that purpose, an OS classifier is specifically selected.

2.5. Extension to new event classes

For many other events, such as high-impedance faults or sensor failures, real-time identification would add additional value to DSOs. However, adding additional identification models for every event would again violate the aforementioned time and resource constraints. For this reason, a central requirement is seen in the capability of a model to be extended to identification of additional events while respecting computational and maintenance effort limitations. The use of an unsupervised event detector allows for the detection of other fast-ramped events beyond the considered FA events. To extend the event

identification problem to slow-ramped events, the persistence forecast-based detector needs to be replaced or extended. However, due to the modular fashion of the proposed pipeline an extension to slow-ramped events can be achieved without affecting the subsequent classification step. With regard to the classification step, the implementation of an OS classifier with rejection and incremental learning capability facilitates the extension to new event classes: indicating observations of unknown classes allows for automated collection, facilitating the manual preparation of new event classes. Once sufficient observations of a new class are collected, the EVM model enables efficient incorporation under an incremental update mechanism. The model reduction strategy makes EVM a sparse OS classifier with limited model size also under extension with new classes.

3. Model description

This section first formulates the problem of unsupervised event detection and OS classification and describes the implemented models. Thereafter, the proposed EIP is explained.

3.1. Unsupervised detection of FA events

In this subsection, the unsupervised event detection problem is formulated. Subsequently, the implemented models for unsupervised event detection are described, namely HTM, ARIMA, CNN, SR and Persistence detector.

3.1.1. Problem formulation

The unsupervised detection of FA events is formulated as a point anomaly detection problem in univariate time series data. A point anomaly is considered a data point that significantly deviates from its expected value. Given a univariate time series $X = \{x_1, x_2, \dots, x_N \mid x_i \in \mathbb{R} \forall i\}$, a data point x_t at time t is declared an anomaly if the anomaly score s_t , defined as the distance to the expected value \hat{x}_t , exceeds a predefined threshold τ :

$$s_t = |x_t - \hat{x}_t| > \tau \quad (1)$$

Although all detectors within this work follow different strategies to calculate s_t , they are all either explicitly or implicitly based on fitting a model to the normal behavior. Given the univariate time series X , all detection models either aim at learning a mapping function Φ from historical time steps to the next time step

$$\hat{x}_t = \Phi(x_{t-w}, \dots, x_{t-1}), \quad (2)$$

or a direct mapping function Θ from historical time steps to the anomaly score of the next time step

$$s_t = \Theta(x_{t-w}, \dots, x_{t-1}), \quad (3)$$

where w is the size of the history window, which can vary for the different detectors.

3.1.2. HTM detector

HTM is a machine learning technique that is based on the structural and algorithmic properties of the neocortex [43]. Compared to many other methods, HTM comes with several advantages that simplify handling of the anomaly detection problem. These include continuous online learning capability, robustness to noise, and applicability without case-specific hyperparameter tuning. Thus, HTM can be applied without model training and selection on separate training and validation datasets and frequent re-training. For a detailed description of HTM the reader is referred to [44].

The implementation of HTM includes an internal calculation of anomaly scores such that the HTM detector overall follows (3). HTM comes with approximately 30 model configuration parameters. For the anomaly detection problem, a set of optimal parameters is provided in the supplementary material of [31], which is applied in this work.

Table 1

Summary of hyperparameters of the CNN model.

Hyperparameter	Search space	Selected value
History window size	144, 288, 576, 1152	288
Forecasting horizon	1	1
Learning rate	[0.00001, 0.1]	1e-5
L2 weight regularization	[0.0001, 0.1]	0.01
Dropout rate	[0, 0.2]	0.2
Batch size	10, 50, 100, 500	50
Maximum number of epochs	2000	2000
Number of filters	10, 30, 50, 70	50
Kernel size	2, 3, 4	3
Neurons in the fully connected layer	10, 50, 100, 150	100
Early stopping patience	10, 50, 100	50
Activation function	ReLU, Tanh	ReLU

3.1.3. ARIMA detector

ARIMA models are widely known and applied for time series forecasting [45]. As they use lagged values to forecast future behavior, ARIMA models can be applied to learn a mapping function according to (2). To enable application to seasonal data, this work considers modeling of seasonal patterns based on Fourier terms as proposed in [46]. Model training and selection is based on the first 10 days of the dataset. In a pre-processing step, the distribution of the training data is centered on a mean of $\mu = 0$ and a standard deviation of $\sigma = 1$. Based on the Akaike information criterion (AIC) a $(p, d, q) = (8, 1, 2)$ ARIMA model is chosen. After the initial training, the autoregressive and moving average parameters are updated with every new incoming observation. Every 14 days an entirely new ARIMA model is selected based on the AIC, resulting in a new selection of p , d and q .

3.1.4. CNN detector

CNNs [47] are a specialized class of artificial neural networks that can capture temporal patterns in time series data. Thus, they can be applied to provide a mapping function according to (2). An advantage of CNNs is the good performance also for small training datasets (as opposed to many other deep learning techniques) even without removing anomalies from the training data [34].

To define the architecture and hyperparameters of the CNN, extensive empirical experiments are conducted based on the first 10 days of the dataset. The model is trained on predicting the difference $x_{\Delta,t} = x_t - x_{t-1}$ instead of x_t to avoid learning a local minimum given by the persistence forecast $\hat{x}_t = x_{t-1}$. To enable faster convergence the training data is centered on a mean of $\mu = 0$ and a standard deviation of $\sigma = 1$. While the first 7 days are used as an initial training dataset, the remaining 3 days are used for validation. The resulting CNN architecture consists of three convolutional/max-pooling pairs followed by a fully connected layer. An overview of the main hyperparameters is given in Table 1. After the initial training and model selection phase the CNN is re-trained every 14 days based on the previous data. To avoid overfitting, a combination of early stopping, L2 regularization and dropout is applied.

3.1.5. SR detector

SR [48] is an unsupervised algorithm for visual saliency detection in computer vision. Recently, Ren et al. [35] proposed the use of SR for anomaly detection in time series data, motivated by the similarity to visual saliency detection. An advantage of SR is the comparatively small number of hyperparameters. SR performs a mapping of previous data points to the next time step according to (2). However, the mapping is conducted within a saliency map representation of a sliding sequence $x = \{x_1, x_2, \dots, x_N\}$, which is based on Fourier transformation and calculation of the spectral residual. Within the saliency map a local average of previous data points is compared to the actual value to declare anomalies similar to (1). As this work is concerned with real-time detection, only the most recent data point x_N of sequence x is

evaluated. Since detection performance of SR improves for data points located in the center of x , Ren et al. [35] propose to add estimated data points following x_N . In this work, the hyperparameters are selected according to the selection in [35], which results from empirical investigation of multiple datasets for time-series anomaly detection. The number of estimated points is set to 5, the size for the local average is 21 and the length of the sequence x is set as 1440. It is worth noting that a parameter study on the investigated dataset revealed low sensitivity of the SR detection performance to hyperparameter selection.

3.1.6. Persistence detector

In this work, the use of a persistence forecast for modeling the expected value \hat{x}_t is proposed and compared to the previously introduced methods. The proposed Persistence detector considers a history window $w = 1$ and determines \hat{x}_t according to

$$\hat{x}_t = x_{t-1}. \quad (4)$$

The triviality of the Persistence detector eliminates the need for model selection, training and re-training.

3.2. OS classification of FA events

This subsection first formulates the OS classification problem. Thereafter, the implemented OS classifier is described.

3.2.1. Problem formulation

OS classification is contrasted with CS methods typically applied in the literature. CS classification requires identical event classes in training and test data and thus assumes full awareness of all existing event classes. For example, a CS classifier trained on observations of line failures and tap change operations will declare every observation of an unknown event class as either a line failure or a tap change. For DSOs it might be impractical to obtain data comprising observations off all existing classes, since events such as failures of transmission elements, large loads or sensors rarely occur. An OS classifier rejects observations of event classes not previously seen and declares them as *unknown*. In Fig. 2 CS classification is compared to OS classification. Making CS assumptions leads to regions of unbounded support, as can be seen from Fig. 2(b). This will result in misclassification of observations from unknown classes which can drastically weaken the performance of the classification.

The problem of OS classification can be formulated as follows [49]. Let $D_{\text{train}} = \{(v_i, y_i)\}_{i=1}^{N_{\text{train}}}$ be a training dataset, with $v_i \in \mathbb{R}^d$ being a feature vector instance and $y_i \in Y_{\text{train}} = \{1, 2, \dots, K\}$ the corresponding event class label. During test or application a classifier needs to predict event classes of the open dataset $D_o = \{(v_i, y_i)\}_{i=1}^{\infty}$, where $y_i \in Y_o = \{1, 2, \dots, K, K+1, \dots, M\}$ with $M > K$. The occurrence of unknown classes requires the classifier to learn a mapping function $f(v) : V \rightarrow Y' = \{1, 2, \dots, K, \text{unknown}\}$, with the option *unknown* representing the rejection of classes not seen during training.

Similar to the described situation of a DSO not having comprehensive recordings of event classes, in the present dataset the number of FA events is comparatively small. Thus, the size of the input feature vectors v_i is limited to 6. Reducing the dimension of the feature space that needs to be described by the limited number of training observations allows for better determination of the decision boundaries. All features are derived from the delta encoded time series (see Section 3.3) and are listed in Table 2. The definition of the input features is given based on a delta encoded sequence observation $x_{\Delta} = \{x_{\Delta,1}, \dots, x_{\Delta,N_x}\}$.

Table 2
Overview of features used for OS classification of FA events.

Feature	Definition
Mean $\mu_{x_{\Delta}}$	$\frac{1}{N_x} \left(\sum_{i=1}^{N_x} x_{\Delta,i} \right)$
Standard deviation $\sigma_{x_{\Delta}}$	$\sqrt{\frac{1}{N_x-1} \sum_{i=1}^{N_x} (x_{\Delta,i} - \mu_{x_{\Delta}})^2}$
Minimum value $x_{\Delta,\text{min}}$	$\min(x_{\Delta})$
Maximum value $x_{\Delta,\text{max}}$	$\max(x_{\Delta})$
Number of zeros n_0	$\text{count}(x_{\Delta} \stackrel{!}{=} 0)$
Points between minimum and maximum value n_{minmax}	$ \text{index}(x_{\Delta,\text{min}}) - \text{index}(x_{\Delta,\text{max}}) $

3.2.2. EVM classifier

The EVM is an OS classifier proposed by Rudd et al. [38]. Advantages of the EVM include incremental learning and model reduction capability, which avoids frequent re-training and keeps model size and classification time small. The EVM models known classes within the training dataset by a set of radial inclusion functions (see Fig. 2(c)). Based on the radial inclusion function of a class C_i , the probability $\hat{P}(C_i|v')$ of a new observation v' belonging to C_i can be determined. The decision function of the EVM is given by

$$y^* = \begin{cases} \arg \max_{j \in \{1, \dots, M\}} \hat{P}(C_j|v') & \text{if } \hat{P}(C_j|v') \geq \rho \\ \text{unknown} & \text{otherwise} \end{cases}, \quad (5)$$

where ρ is a threshold defining the boundary between the set of known classes and the unknown open space.

The EVM model training and selection is based on 90% of the available event observations applying 5-fold time series cross-validation. The features are standardized based on training data for every individual split. In order to select an appropriate threshold ρ , a minimum performance requirement on the training dataset is defined based on the F_1 score (Section 4.2). The model with the smallest threshold ρ , which still fulfills the performance requirement $F_1 \geq 0.8$ in the time series cross-validation, is selected. An overview of the selected hyperparameters is given in Table 3.

3.3. EIP

To connect the presented models for unsupervised event detection and OS classification, a data processing pipeline is proposed. In Fig. 1

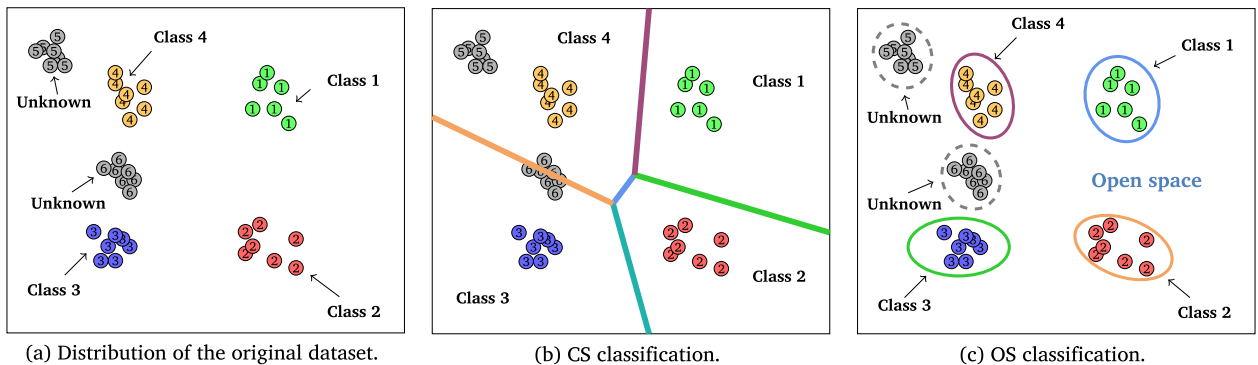


Fig. 2. Comparison of CS and OS classification according to [27].

Table 3
Summary of hyperparameters of the EVM model.

Hyperparameter	Search space	Selected value
Tailsizer	[1,100]	7
Distance multiplier	[0.1,1.1]	0.9
Distance metric	Canberra, Cosine, Euclidean	Canberra
Threshold	[0.1, 0.99999]	0.9

a schematic overview of the EIP is depicted. The functionality of the main building blocks of the pipeline was described in Sections 3.1 and 3.2. In Section 2.3 the use of delta encoded data for the EIP was motivated. Delta encoding exploits the autocorrelation of time series data. In the simplest version of delta encoding, a time series $X = \{x_1, x_2, \dots, x_N\}$ is encoded as difference between successive samples, resulting in the differenced time series $X_\Delta = \{x_1, x_2 - x_1, \dots, x_N - x_{N-1}\} = \{x_1, x_{\Delta,2}, \dots, x_{\Delta,N}\}$. Delta encoding performs best when the values in the original data contain only small changes between adjacent values [50]. By applying the Persistence detector on the differenced time series X_Δ , the anomaly detection problem reduces to comparison of the amplitudes of $x_{\Delta,i}$ to the predefined threshold τ . To connect point-wise unsupervised event detection with OS classification, an event sampler is interposed, exploiting the characteristics of fast-ramped events. As can be seen in Fig. 4 fast-ramped FA events show peaks in the delta encoded data at the beginning and/or end of an event. Based on this property, the detection of an event at data point $x_{\Delta,t}$ can be used to extract a backward sequence sample $x_{\Delta,bw} = \{x_{\Delta,t-w_x}, \dots, x_{\Delta,t+e}\}$ and forward sequence sample $x_{\Delta,fw} = \{x_{\Delta,t-e}, \dots, x_{\Delta,t+w_x}\}$, where w_x is the sample window size and e a window extension, ensuring sampling of the entire event. In case of forward sampling, an early stopping criterion can be introduced which breaks the sampling process in case another event is detected at a data point $x_{\Delta,t+a} \in \{x_{\Delta,t+1}, \dots, x_{\Delta,t+w_x}\}$, resulting in a forward sample $x_{\Delta,fw} = \{x_{\Delta,t-e}, \dots, x_{\Delta,t+a+e}\}$. Such event-triggered early stopping of the forward sampling process reduces the sampling time and thus the time until a sample can be classified. In the Appendix in Algorithm 1 the general procedure of the proposed EIP is described.

4. Experimental setup

This section presents the experimental setup of this study. The considered dataset as well as the preparation of the dataset for investigation of unsupervised event detection and OS classification of FA events are presented in Section 4.1. Section 4.2 introduces metrics for the evaluation of the detection and classification performance.

4.1. Dataset and data preparation

The first part of this subsection is concerned with presenting key information of the dataset under investigation as well as describing the process of activating flexibility. In the second part the preparation of

the dataset is explained, which includes data cleaning and in case of preparation for the investigation of the OS classification the extension of the dataset with two artificial event classes.

4.1.1. Dataset

Within this work, a dataset from EcoGrid 2.0 is used. EcoGrid 2.0 was a demonstration project which examined the use of flexible consumption of residential customers for power system services at transmission system operator and DSO level [9]. The experiments were conducted on the Danish island of Bornholm. The residential customers were equipped with SMs and ICT infrastructure for participating in demand response experiments. The flexible load corresponded to electric heaters and heat pumps that were controlled by adjusting room temperature setpoints or by sending a throttle signal, respectively. An increase in the setpoints results in higher consumption, while lowering the setpoints leads to a reduction of consumption. The throttle signal blocks the operation of the heat pump until the signal is released. Besides flexible load, the installed SMs also capture household consumption and photovoltaic production, when present. As can be seen in Fig. 3(a), the dataset used consists of six and a half months of aggregated load data, beginning from 15th of September 2017. The aggregated active power profile consists of 450 household loads with a 5 min time resolution. The FA events comprise load reduction and load increase experiments (see Fig. 3(b)) realized by two different aggregators and customer portfolios. Activation periods are in the range of 30 – 120 min. Different numbers of customer loads participated in each FA, and activations were conducted under varying conditions of temperature, time of the day and photovoltaic production. In Fig. 3(a) the trend and (b) seasonality of the non-stationary load time series can be noticed.

4.1.2. Dataset preparation for investigation of unsupervised event detection

The dataset contains 325 FA events. Start and end time as well as type of FA event is known for every experiment. In 75 cases two flexibility portfolios were activated simultaneously. To avoid double counting of either true positives (TPs) or false negatives (FNs), parallel experiments are considered as one FA event, reducing the number of FA events to 250. In most cases, experiments result in either load reduction or load increase of a subset of customer loads. However, in some cases little to no flexibility was activated. This can be due to exclusive testing of connectivity, failed activation of flexibility assets, or low flexibility potential due to high temperatures and thus low heating demand. For this work, such FA events are not considered in the performance evaluation, reducing the number of FA event samples to 205.

4.1.3. Dataset preparation for investigation of OS event classification

In Fig. 4 examples of all event classes are depicted in absolute and delta encoded values. Note that only for the OS classification problem all introduced event classes are considered. For the FA event detection problem only FA events are taken into account. The classifier is trained on two known event classes, namely FA and normal operation (NO)

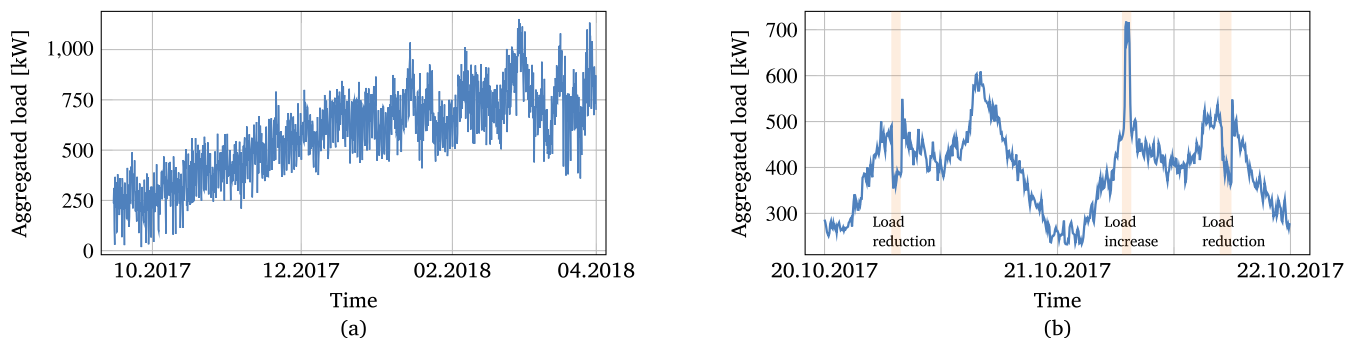


Fig. 3. Aggregated load dataset (a) and aggregated load of two representative days with frequent FAs (b). Periods of FAs are marked with an orange background.

events. All 205 FA events are sampled, including 3 timestamps (15 min) that precede and succeed each event. As the duration of FA events varies, the length of FA event samples varies as well. An equal amount of NO events are randomly sampled from the remaining dataset. The length of a NO event is randomly selected from the distribution of FA event lengths. With this approach, the classifier is prevented from differentiating between FA and NO events based on the sample length. As the sample length of both FA and NO events can vary in the proposed EIP (see Section 3.3), learning a constant sample length is not considered a valid approach.

The problem of OS classification, as formulated in Section 3.2.1, requires additional event classes within the test dataset to investigate the capability of rejecting unknown classes. In this work, 3 unknown event classes are considered. Besides the FA and NO event classes, the EcoGrid 2.0 dataset includes another event class, which in the course of this work will be called Monday peak (MP). On every Monday within the dataset, a load peak occurs at around 8 am. The load peak results from short, collective heating of electric water boilers to 80 °C to inhibit the growth of bacteria. Only MP events that could be manually detected in an extensive ex-post evaluation of the dataset are considered. MP events are not considered a normal operation and thus no overlapping of NO and MP events exists. In total the dataset includes 15 MP events. In order to extend the OS classification problem, two additional artificial event classes are introduced, namely the frozen value (FV) and data unavailability (DU) event class. The FV event class models a data transmission or processing failure in which a measurement at time t_0 remains constant for N_{cons} consecutive steps. At $t_{N_{\text{cons}}+1}$ recording of true measurements is reestablished. The length of FV events is randomly selected from the distribution of FA event lengths. 205 FV events are randomly introduced into the subset of the dataset which is not influenced by FA, NO and MP events. In a DU event a subset of individual measurements, e.g. SM readings, is considered to be unavailable due to device or data transmission failure. The fraction of available measurements is randomly selected from the uniform distribution $U(0.4, 0.8)$. As for NO and FV events, the length of DU events is drawn from the distribution of FA event lengths. 205 DU events are introduced into the subset of the dataset not affected by any of the previously described events.

4.2. Performance metrics

The performance evaluation of both the unsupervised detection and OS classification of FA events is based on a labeling of each

instance of the respective dataset. The definition of an instance for the detection and classification task follows in Sections 4.2.1 and 4.2.2, respectively. One performance metric used for evaluation of both the detection and classification part is the F_1 score. The F_1 score represents the harmonic mean between precision and recall, and is a widely applied performance metric for detection and classification problems with imbalanced classes. Let TP_l , FP_l , TN_l , and FN_l be the number of TPs, false positives (FPs), true negatives (TNs), and FNs for the l th event class, respectively, where $l \in \{1, 2, \dots, M\}$. For a multi-class problem, the macro F_1 score is calculated by

$$F_1 = \frac{1}{M} \sum_{l=1}^M F_{1,l} = \frac{1}{M} \sum_{l=1}^M 2 \frac{Pr_l \cdot Re_l}{Pr_l + Re_l}, \quad (6)$$

where precision Pr and recall Re of the l th event class are defined as

$$Pr_l = \frac{TP_l}{(TP_l + FP_l)}, \text{ and } Re_l = \frac{TP_l}{(TP_l + FN_l)}. \quad (7)$$

4.2.1. Unsupervised event detection metrics

In this work, the problem of unsupervised event detection is considered an anomaly detection problem, reducing the number of classes to $M = 2$. Since anomalies constitute the primary class of interest, the multi-class formulation of the F_1 score in (6) reduces to

$$F_1 = F_{1,1} = 2 \frac{Pr_1 \cdot Re_1}{Pr_1 + Re_1}. \quad (8)$$

Besides the F_1 score, another widely applied performance metric for anomaly detection is the area under the precision–recall curve (AUCPR). Precision–recall curves summarize the trade-off between precision and recall for different thresholds τ . While the consideration of TNs in traditional receiver-operating-characteristic curves may lead to an overly optimistic view on the performance in case of highly imbalanced classes, AUCPR is specifically tailored to problems with imbalanced classes or rare events.

In this work, the entire sequence of an FA event, referred to as event window ω_{FA} , is considered for labeling as TP or FN. The event window of a FA event is defined by the FA start time $t_{\text{FA,start}}$ and end time $t_{\text{FA,end}}$. Since the dataset under investigation is a real-world dataset it contains some inaccuracies in the event labeling. In some cases flexibility was activated before the official start time $t_{\text{FA,start}}$. Since in these cases an early detection would result in falsely FNs, the event window ω_{FA} is extended by 10 min, such that $\omega_{\text{FA}} = \{x_{t_{\text{FA,start}}-2}, \dots, x_{t_{\text{FA,end}}}\}$. The first detection that falls into ω_{FA} is considered as TP, while further detections within the same event window are ignored. If no point of ω_{FA} is detected the event label will be considered a FN. In most cases,

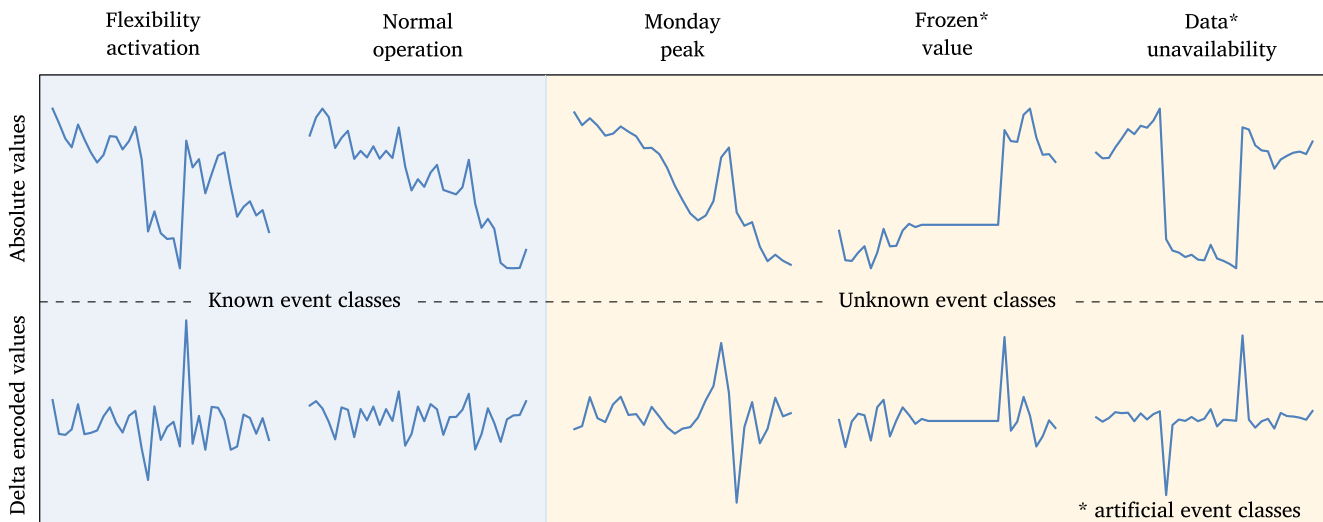


Fig. 4. Exemplary representation of event classes considered in this work.

FAs result in a subsequent load rebound, which can be considered a deviation from the normal load behavior outside of the activation period. While regarding a detection within the rebound area as TP would introduce a positive bias to detection performance, considering them as FP would result in overly pessimistic performance results, as the detector indeed has detected an anomaly. For this reason, a rebound window ω_R is introduced in which detections are ignored and are thus neither considered a TP nor FP. The length of a rebound window is defined as three times the FA event length, resulting in $\omega_R = \{x_{t_{FA,end}+1}, \dots, x_{t_{FA,end}+3 \times N_{FA}}\}$. In contrast to calculation of TPs and FNs, the calculation of FPs and TNs is conducted point-wise. Thus, all detections outside the event and rebound windows are considered FP.

To evaluate the early detection capability, the average detection delay $\bar{\delta}_{det}$ is introduced as the average time between FA start time $t_{FA,start}$ and the first detection time t_{det} in minutes according to

$$\bar{\delta}_{det} = \frac{1}{N_{det}} \sum_{i=1}^{N_{det}} \delta_{det,i} = \frac{1}{N_{det}} \sum_{i=1}^{N_{det}} (t_{det,i} - t_{FA,start,i}), \quad (9)$$

where N_{det} is the number of detected FA events and $\delta_{det,i}$ the detection delay of a detected FA event. Note that the detection delay of detections is assumed to be zero within the subset $\{x_{t_{FA,start}-2}, x_{t_{FA,start}-1}, x_{t_{FA,start}}\}$.

The use of widely applied performance metrics allows for an easy understanding and comparison of the results to other studies. However, in order to take performance requirements of a specific scenario into account an individual performance metric is required. For this purpose, the flexibility activation detection score (*FAD* score) is proposed. In the scenario of real-time detection of FAs in active DNs, the cost of FN is considered to be higher compared to the cost of FP. While a missed critical FA could lead to violation of power or voltage boundaries, a false alarm would result in a moderate additional manual inspection effort. Moreover, in the proposed pipeline, a FP will lead to a sample of normal behavior (NO event class) which can be classified as such by the OS classifier. In this way the classifier relativizes the FP of the unsupervised event detector. For these reasons, the *FAD* score puts more weight on FNs than on FPs. Further, early detection capability is an important requirement in the considered scenario. The earlier a potentially critical FA event is detected, the greater the scope for countermeasures. On the contrary, detection near the end of a critical FA results in almost no benefit. The *FAD* score takes these considerations into account and expresses the performance for the specific scenario of real-time FA event detection in one score. Besides an easier evaluation and comparison of detection models, the *FAD* score also allows for easy selection of an optimal threshold τ for unsupervised detection of FA events. The *FAD* score is constituted by 3 scoring functions ζ_{TP} , ζ_{FN} and ζ_{FP} , representing the contribution of TPs, FNs and FPs, respectively:

$$FAD = \zeta_{TP} - \zeta_{FN} - \zeta_{FP}. \quad (10)$$

Let $x_{t_{FA,start},i}$, $x_{t_{FA,end},i}$ and $x_{t_{det},i}$ be the start point, end point and first detection within the event window $\omega_{FA,i}$ of the i th FA event, respectively. Then ζ_{TP} is given by

$$\zeta_{TP} = \sum_{i=1}^{N_{det}} \sigma_i(x_{t_{det},i}), \quad (11)$$

where $\sigma_i(x_{t_{det},i})$ is the positive score of one TP detection. Between $\sigma_i(x_{t_{FA,start},i}) = \xi$ and $\sigma_i(x_{t_{FA,end},i}) = 0$ the score σ_i follows a linear declining function, where ξ is the maximum positive score of one TP detection. For each missed FA event a negative score of η is considered according to

$$\zeta_{FN} = -\eta \cdot FN_1. \quad (12)$$

Scoring function ζ_{FP} is represented by a moved negative exponential function given as

$$\zeta_{FP} = -\gamma \cdot \exp\left(\frac{-FP_1}{\nu}\right) - \nu, \quad (13)$$

where γ is the maximum negative score of a FP detection. Note that $\gamma \ll \xi$. Parameters ν and ν allow for additional adjustments of the score to the dataset. The negative exponential decline considers that the first FP detections will have a stronger negative impact on performance, while for subsequent FPs the additional negative impact is small. The final *FAD* score is normalized such that $FAD_{norm} \in [0, 1]$, according to

$$FAD_{norm} = \frac{FAD - FAD_{null}}{FAD_{opt} - FAD_{null}}, \quad (14)$$

where FAD_{null} is the *FAD* score without any detection and FAD_{opt} the *FAD* score under optimal detection. In this work, the parameters are selected as $\xi = 1$, $\eta = 1$, $\gamma = 0.05$, $\nu = 10000$ and $\nu = 0$.

4.2.2. OS event classification metrics

The performance evaluation of the OS classification of FA events is based on the F_1 score according to (6). However, for the OS problem the number of classes should only be determined by the known classes. Considering all of the unknown classes as a single additional class in the test dataset would result in a biased performance result. If the problem is treated in the same way as a CS scenario, rejected samples of unknown classes would be considered as TPs - although no training samples of the unknown classes existed. Instead, the calculation of the F_1 score is given by

$$F_1 = \frac{1}{K} \sum_l F_{1,l} = \frac{1}{K} \sum_l 2 \frac{Pr_l \cdot Re_l}{Pr_l + Re_l}, \quad (15)$$

where K is the number of known classes from the training dataset. In Section 5.2 the influence of the number of unknown event classes on the classification performance will be evaluated, which requires the definition of the *openness* of a test dataset. In [36] the authors introduce a formal definition of the openness O of a dataset according to

$$O = 1 - \sqrt{\frac{2 \times |\text{training classes}|}{|\text{testing classes}| + |\text{target classes}|}}, \quad (16)$$

with $O \in [0, 1]$. Large values for O correspond to a higher number of unknown classes in the dataset, while for the CS problem $O = 0$.

5. Results and discussion

In this section the performance evaluation of the proposed models for unsupervised detection and OS classification of FA events is presented. In Section 5.1 the proposed Persistence detector is compared to various other models, introduced in Section 3.1. Section 5.2 investigates the OS classification of FA events based on the introduced EVM model (Section 3.2.2). The classification performance is compared to a CS classifier benchmark. The performance evaluation is conducted based on the dataset and performance metrics from Section 4.

5.1. Unsupervised FA event detection

In Fig. 5 the maximum F_1 score $F_{1,max}$ and *FAD* score FAD_{max} at the optimal threshold $\tau_{opt,F1}$ and $\tau_{opt,FAD}$, respectively, are depicted together with the *AUCPR* for Persistence, HTM, ARIMA, SR and CNN detector. From the comparison of $F_{1,max}$ and *AUCPR* it can be derived that Persistence, ARIMA, SR and CNN detector roughly lie in the same performance range. However, the HTM detector shows a significantly poorer performance. While according to the F_1 score the Persistence, ARIMA and CNN detector achieve the best detection results, with $F_{1,max} = 0.71$, in accordance with the *AUCPR* the SR detector outperforms all other detectors with *AUCPR* = 0.69. Interestingly, with a difference of 4 percentage points the F_1 score shows a poorer detection performance for the SR detector.

Although both the F_1 score and *AUCPR* are performance metrics specifically tailored to scenarios with highly imbalanced classes and higher emphasis on the positive class, they suggest different results. This again motivates the need for a scenario-specific performance




Fig. 5. Maximum F_1 score ($F_{1,max}$), AUCPR and maximum FAD score (FAD_{max}) for Persistence, HTM, ARIMA, SR and CNN detector.

metric. Based on the comparison of the maximum FAD score FAD_{max} in Fig. 5 it can be concluded that both the ARIMA and CNN detector achieve the best result for the problem of real-time detection of FA events in aggregated load data with $FAD_{max} = 0.7$. On the contrary, with $FAD_{max} = 0.6$ the SR detector shows a significant poorer performance in the given case of FA event detection. However, according to $F_{1,max}$ and AUCPR the SR detector can potentially keep up with or even outperform other detection methods in scenarios with other requirements. With $FAD_{max} = 0.69$ the performance of the Persistence detector is only slightly below the best FAD scores achieved by the ARIMA and CNN detector. As can be seen in Fig. 4 FA events in the given dataset are fast-ramped events that are characterized by a steep slope at the beginning and end of an event, resulting in a large deviation between x_t and x_{t-1} . In case of the Persistence detector this large deviation directly translates to a large anomaly score according to (1). Given that the Persistence detector can keep up with the more complex detectors regardless of the considered performance metrics, it can be concluded that the Persistence detector constitutes a trivial but effective method for detection of fast-ramped FA events.

In Fig. 6 the average detection delay $\bar{\delta}_{det}$ is shown as a function of the threshold τ for all detectors. As for $\tau = 0$ all data points are declared an event, the average detection delay is $\bar{\delta}_{det} = 0$ for all detectors. It can be seen that the HTM detector has the lowest detection delay for thresholds $\tau > 0.1$. The comparatively high early detection capability also explains the reduced performance discrepancy between the HTM and the other detectors for FAD_{max} compared to $F_{1,max}$ and AUCPR (Fig. 5). While for the FAD score the contribution of TPs is weighted based on the detection delay, $F_{1,max}$ and AUCPR do not take early detection into account.

The Persistence, ARIMA, SR and CNN detector show a similar $\bar{\delta}_{det}$ for $\tau < 0.4$. For thresholds $\tau > 0.4$ the SR detector shows a significantly higher detection delay, while Persistence, ARIMA and CNN detector continuously show a similar detection delay. Fig. 7 compares the

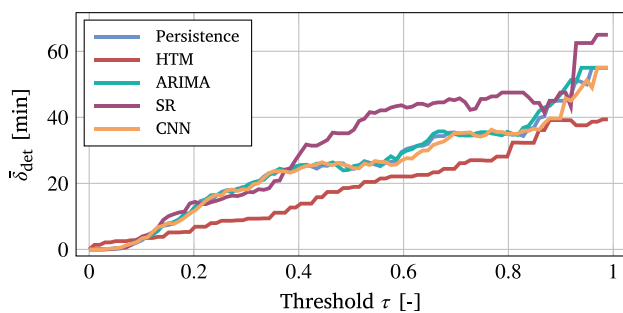


Fig. 6. Average detection delay $\bar{\delta}_{det}$ for Persistence, HTM, ARIMA, SR and CNN detector.

average detection delay $\bar{\delta}_{det}$ of all detectors at the optimal threshold $\tau_{opt,FAD}$ corresponding to FAD_{max} . Although the HTM detector has the lowest average detection delay over the largest range of τ (Fig. 6), at an operating point relevant for the considered scenario (i.e. at FAD_{max}), it has a higher detection delay compared to most other detectors. The HTM detector requires a higher threshold compared to the other detectors (see Fig. 8(c)), due to the particularly strong vulnerability to high FP numbers for low thresholds. The higher threshold in turn explains the higher average detection delay of the HTM detector. The SR detector with $\bar{\delta}_{det} = 11.42$ min has by far the highest detection delay. Persistence, ARIMA and CNN detector show similar delays. In fact, the proposed Persistence detector shows the lowest detection delay with $\bar{\delta}_{det} = 7.41$ min. However, it has to be considered, that the calculation of the average detection delay is only based on detected events according to (9). Thus, detecting additional events close to the end of an event (as done by ARIMA and CNN detector) results in an improved FAD score, as negative FN scores are avoided, even though the average detection delay increases.




Fig. 7. Average detection delay $\bar{\delta}_{det}$ at FAD_{max} for Persistence, HTM, ARIMA, SR and CNN detector.

Fig. 8(a) and (b) show the F_1 score over the threshold τ and the precision–recall curve for the different detectors, respectively. Both on the F_1 score and precision–recall curve a strong similarity of the Persistence, ARIMA and CNN detector can be noticed. This can be explained by the signal-to-noise ratio of the dataset. Although aggregated, the load data under investigation show a comparatively low signal-to-noise ratio due to fluctuations, introduced by unforeseeable customer behavior, and the high resolution of the data. Because of the low signal-to-noise ratio it is difficult for more complex methods, such as the applied ARIMA and CNN model, to extract additional information from the dataset compared to the trivial persistence forecast. Thus, the explainability of the dataset can be exploited by the persistence forecast to a large extent, explaining the similarity of the Persistence, ARIMA and CNN detector. However, the SR detector clearly shows a different behavior. This is due to the different mathematical approach of transforming the dataset from time into the frequency domain.

Fig. 8(b) shows that, compared to all other detectors, the SR detector is able to keep the precision on a higher level for an increasing recall. While a high precision is not seen as an important requirement for the considered scenario, the SR detector may have advantages over the other detectors in scenarios with different requirements. Interestingly, the HTM detector clearly shows a different behavior compared to the Persistence, ARIMA and CNN detectors, even though it is also based on a time series forecast. This can partly be explained by the internal calculation of the anomaly score, which differs from the external calculation used for Persistence, ARIMA and CNN detector (see Section 3.1.2). However, the comparatively poor detection performance also indicates a poor underlying forecast that is even outperformed by a trivial persistence forecast. A potential reason could be insufficient adaption of the various model parameters to the dataset and scenario. Although the authors of HTM claim the provided set of parameters to be the best for anomaly detection, it may not be sufficiently appropriate for the given scenario. However, as explained before, due to the low signal-to-noise ratio, it can be expected that even extensive parameter




Fig. 8. (a): F_1 score, (b): Precision–recall curve, (c): FAD score for Persistence, HTM, ARIMA, SR and CNN detector.

tuning will not result in a significantly better forecast compared to the persistence forecast.

Fig. 8(c) shows the FAD score for all detectors over the threshold τ . By comparing the F_1 score with the FAD score, a shift between the optimal threshold $\tau_{opt,F1}$ and $\tau_{opt,FAD}$ towards smaller values can be noticed. A smaller threshold increases the number of TPs and results in earlier detection of an event. At the same time, the number of FPs increases as well. However, as described in Section 4.2.1 the FAD score emphasizes early event detection and weights FPs low compared to FNs, explaining the decrease of the optimal threshold. Based on the FAD score an optimal threshold for the proposed Persistence detector of $\tau_{opt,FAD} = 0.16$ is determined. At $\tau_{opt,FAD}$ 191 of 205 FA events (93%) are detected by the Persistence detector, while 498 data points outside the event and rebound windows (1.25%) are falsely declared a FA event.

As previously described, for the dataset and scenario under investigation, more sophisticated models such as ARIMA and CNN only achieve minor improvements of the forecast compared to the persistence forecast. This also translates to a similar characteristic of the FAD score curve over the threshold τ . It can be inferred, that for the detection of fast-ramped FA events an upper performance limit should exist at an FAD score of roughly $FAD \approx 0.7$. This performance limit can approximately be reached with the proposed Persistence detector ($FAD_{max} = 0.69$). More advanced detection methods, such as the ARIMA and CNN detectors, only slightly improve the detection performance, but require a significantly higher maintenance and computational effort. The Persistence detector is therefore proposed to avoid frequent time and computation intensive model re-training. As described in Section 2 this is considered a great advantage in a scenario of edge computing-based distributed event detection with time and resource constraints.

5.2. OS classification of FA events

In Fig. 9 the confusion matrix for the EVM model applied on the OS test dataset is depicted. Besides the two known event classes FA and NO, three unknown event classes are included in the test dataset, namely MP, FV and DU, which are summarized as *unknown*. The test dataset in total contains 63 observations with $N_{FA} = 21$, $N_{NO} = 21$, $N_{MP} = 7$, $N_{FV} = 7$, $N_{DU} = 7$. The test dataset has an openness of $O = 24.7\%$. From Fig. 9 it can be derived that the EVM is able to correctly classify 90% of the FA events and 76% of the NO events. Moreover, the EVM successfully rejects 71% of all observations of the unknown classes. It can be concluded that the EVM in principle is able to differentiate between FA and NO observations also in an OS scenario with an acceptable performance. However, the EVM also erroneously classifies 29% of the observations from the unknown classes as FA events, which reduces precision for FA events. Precision for NO events is not affected by the unknown event classes. Also the recall of the

FA and NO event classes is negatively influenced, since 10% of the FA events and 19% of the NO events, respectively, are rejected. In general, Fig. 9 demonstrates that the influence of the unknown classes on precision and recall of a class is higher compared to the influence of the other known class.

True event class	FA	19 90%	0 0%	2 10%
	NO	1 5%	16 76%	4 19%
	unknown	6 29%	0 0%	15 71%
		FA	NO	unknown
		Predicted event class		

Fig. 9. Confusion matrix of the EVM model on the OS classification test dataset with an openness $O = 24.7\%$.

In order to investigate the benefit of applying an OS classifier in the more realistic scenario with presence of unknown classes, performance is compared to a CS classifier. For this purpose, the EVM model is applied on the test dataset as both an OS and CS classifier. In the CS setting the rejection of observations with $\hat{P}(C_i|v') < 0.9$ is deactivated and observations are classified according to $\hat{P}(C_i|v')$. Moreover, in order to investigate the influence of the number of unknown classes, the comparison is conducted on a test dataset with increasing fractions of unknown classes. In Fig. 10 the comparison of the OS and CS EVM for a varying openness O of the test dataset is depicted. Performance is evaluated based on the F_1 score. For the CS problem ($O = 0$) the performance of the CS classifier is slightly better compared to the OS classifier, since the OS classifier wrongly rejects some of the observations of the known classes. By adding MP events as unknown class to the test dataset the openness of the test dataset increases to $O = 10.56\%$. In this scenario the OS classifier outperforms the CS classifier. While the OS classifier is capable of rejecting observations from unknown classes, the CS classifier assigns all observations of unknown classes to one of the known classes, resulting in a decreased precision. However, the performance of the OS EVM decreases as well. This is due to two reasons. First, not all observations from the unknown classes are successfully rejected. Second, being capable of rejecting observations can also lead to falsely rejected observations of known classes. Nevertheless, for the investigated scenario of FA event classification the rejection capability improves the performance compared to the CS classifier already for the existence of only one unknown class. Extending the test dataset with the unknown FV event class ($O = 18.35\%$) has no influence

on the performance of the OS EVM. This can be explained by the specific characteristic of FV events. The number of zeros n_0 constitutes a strong differentiator for observations of the FV event class, making it comparatively easy for the OS EVM to differentiate between FV and the known FA and NO events. Nevertheless, the F_1 score of the CS classifier further decreases from $F_1 = 0.839$ to $F_1 = 0.792$ since all observations of the FV event class are assigned to either the FA or NO event class. In the final scenario ($O = 24.7\%$) all unknown event classes are added to the test dataset, corresponding to the scenario described by Fig. 9. Adding the unknown DU event class further decreases performance for both the OS and CS classifier. However, while for the CS EVM the F_1 score decreases by 6.31%, the OS classifier only shows a decrease of 3.12% percent. In summary, the performance of the OS EVM decreased from $F_1 = 0.896$ ($O = 0\%$) to $F_1 = 0.837$ ($O = 24.7\%$), while the CS classifier performance decreased from $F_1 = 0.905$ ($O = 0\%$) to $F_1 = 0.742$ ($O = 24.7\%$). This demonstrates that the rejection capability of the OS classifier allows maintaining the classification performance on a higher level, for increasing fractions of unknown classes. Nevertheless, also for the OS classifier the performance deteriorates with additional unknown classes.

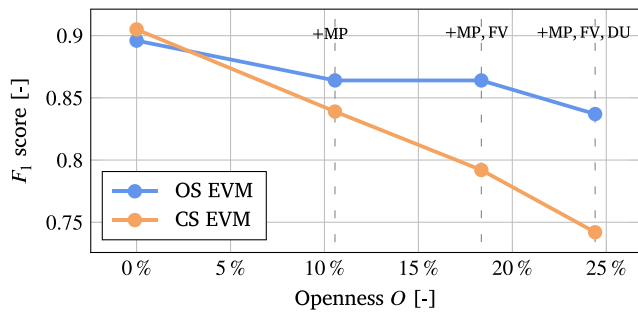


Fig. 10. Comparison of the F_1 score between OS and CS EVM model on the OS classification problem with a varying openness O .

To summarize, FA events can be classified also in the more realistic OS scenario and applying OS classifiers can significantly improve performance under these conditions. However, the classification performance of the EVM is expandable, due to the very limited training data. The size of the dataset constitutes a limitation of the presented study, since the small number of observations and classes prevent more comprehensive investigations. Nevertheless, this study proves the fundamental feasibility of OS classification of FA events on real data.

6. Conclusion and future work

This work demonstrates the fundamental feasibility of unsupervised detection and OS classification of fast-ramped FA events. A data processing pipeline for FA event identification is suggested, which combines both steps. For unsupervised FA event detection, a simple Persistence detector is proposed and implemented. The comparison with more complex and computationally expensive detection models demonstrates a similar performance and the existence of an upper performance limit. Results indicate that the Persistence detector is particularly suitable for the specific class of FA events. As OS classifier the EVM is used. It is shown that the use of an OS classifier significantly improves classification performance in the more realistic OS scenario compared to a traditional CS classifier. Both the Persistence detector and EVM classifier are selected with a view to an application in a distributed event detection architecture with time and resource constraints due to edge computing. Their good performance demonstrates that main building blocks of the proposed pipeline can be realized with comparatively simple and lightweight methods that fulfill important requirements for an application in a distributed event detection architecture.

Given the fundamental proof of the main building blocks, a logical next step is the investigation of the coupling of the Persistence detector

and EVM classifier in the proposed EIP for FA events. Moreover, for both the detection as well as classification step, several possible improvements could be investigated. One direction could be the integration of additional regressors for unsupervised event detection such as temperature and solar radiation. For the OS classification problem principle component analysis or other methods for dimensionality reduction could be applied to reduce the feature space dimension while retaining the majority of the information. Finally, the proposed pipeline could be extended from FA event identification to identification of multiple relevant events in active DNs.

CRedit authorship contribution statement

Nils Müller: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Visualization. **Carsten Heinrich:** Conceptualization, Resources, Data curation, Writing – review & editing. **Kai Heussen:** Conceptualization, Resources, Writing – review & editing, Supervision. **Henrik W. Bindner:** Funding acquisition, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is partly funded by the Innovation Fund Denmark (IFD) under File No. 91363 and by the INTERPRETER project, which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 864360.

Appendix

Algorithm 1 General procedure of the EIP for FA events.

```

1: for new incoming data point  $x_{\Delta,t}$  do
2:   PERSISTENCE_FORECAST( $x_{\Delta,t}$ )  $\triangleright$  Start of unsupervised event detection
3:   return  $\hat{x}_{\Delta,t}$ 
4:   if  $|\hat{x}_{\Delta,t} - x_{\Delta,t}| < \tau$  then:
5:     Declare  $x_{\Delta,t}$  normal behavior
6:   else:
7:     Declare  $x_{\Delta,t}$  an event
8:     BACKWARD_SAMPLING( $x_{\Delta,t}$ )  $\triangleright$  Start of event sampling
9:     return  $x_{\Delta,bw} = \{x_{\Delta,t-w_x}, \dots, x_{\Delta,t}\}$ 
10:    FORWARD_SAMPLING( $x_{\Delta,t}$ )
11:    if another event at  $x_{\Delta,t+a} \in \{x_{\Delta,t+1}, \dots, x_{\Delta,t+w_x}\}$  then:
12:      return  $x_{\Delta,fw} = \{x_{\Delta,t-e}, \dots, x_{\Delta,t+a+e}\}$ 
13:    else:
14:      return  $x_{\Delta,fw} = \{x_{\Delta,t-e}, \dots, x_{\Delta,t+w_x}\}$ 
15:    for  $x$  in  $[x_{\Delta,bw}, x_{\Delta,fw}]$  do:
16:      Calcul. feature vector  $v = [\mu_x, \sigma_x, x_{\min}, x_{\max}, n_0, n_{\minmax}]$ 
17:      EXTREME_VALUE_MACHINE( $v$ )  $\triangleright$  Start of OS classification
18:      return  $P(C_{flexibility}|v), P(C_{normal\_behavior}|v)$ 
19:      if  $P(C_{flexibility}|v) \geq P(C_{normal\_behavior}|v)$  and  $\geq \rho$  then:
20:        Declare sample  $x$  as flexibility activation event
21:      else if  $P(C_{flexibility}|v) \leq P(C_{normal\_behavior}|v)$  and  $\geq \rho$  then:
22:        Declare sample  $x$  as normal behavior
23:      else:
24:        Declare sample  $x$  as unknown event

```


References

- [1] European Commission. Impact assessment on stepping up Europe's 2030 climate ambition. investing in a climate-neutral future for the benefit of our people. 2020. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020SC0176>. [Online; accessed: 23/07/2021].
- [2] Zhao J, Wang C, Zhao B, Lin F, Zhou Q, Wang Y. A review of active management for distribution networks: current status and future development trends. *Electr Power Compon Syst* 2014;42(3-4):280-93. <http://dx.doi.org/10.1080/15325008.2013.862325>.
- [3] Spiliotis K, Ramos Gutierrez AI, Belmans R. Demand flexibility versus physical network expansions in distribution grids. *Appl Energy* 2016;182:613-24. <http://dx.doi.org/https://doi.org/10.1016/j.apenergy.2016.08.145>.
- [4] Muratori M, Rizzoni G. Residential demand response: Dynamic energy management and time-varying electricity pricing. *IEEE Trans Power Syst* 2015;31(2):1108-17. <http://dx.doi.org/10.1109/TPWRS.2015.2414880>.
- [5] Schey S, Scofield D, Smart J. A first look at the impact of electric vehicle charging on the electric grid in the EV project. *World Electr Veh J* 2012;5(3):667-78. <http://dx.doi.org/10.3390/wevj5030667>.
- [6] Richardson P, Flynn D, Keane A. Impact assessment of varying penetrations of electric vehicles on low voltage distribution systems. In: *IEEE PES General Meeting*, 2010, 1-6. <http://dx.doi.org/10.1109/PES.2010.5589940>.
- [7] Lehnhoff S, Krause O, Rehtanz C. Dezentrales autonomes Energiemanagement. *Automatisierungstechnik* 2011;59(3):167-79. <http://dx.doi.org/doi:10.1524/auto.2011.0906>.
- [8] Sperstad IB, Degefa MZ, Kjølle G. The impact of flexible resources in distribution systems on the security of electricity supply: A literature review. *Electr Power Syst Res* 2020;188. <http://dx.doi.org/https://doi.org/10.1016/j.epsr.2020.106532>.
- [9] Heinrich C, Ziras C, Syrri AL, Bindner HW. EcoGrid 2.0: A Large-scale field trial of a local flexibility market. *Appl Energy* 2020;261. <http://dx.doi.org/10.1016/j.apenergy.2019.114399>.
- [10] Ziras C, Heinrich C, Pertl M, Bindner HW. Experimental flexibility identification of aggregated residential thermal loads using behind-the-meter data. *Appl Energy* 2019;242:1407-21. <http://dx.doi.org/10.1016/j.apenergy.2019.03.156>.
- [11] Mishra A, Irwin D, Shenoy P, Zhu T. Scaling distributed energy storage for grid peak reduction. In: *Proceedings of the Fourth International Conference on Future Energy Systems*. 2013, p. 3-14. <http://dx.doi.org/10.1145/2487166.2487168>.
- [12] Zhang H, Liu B, Wu H. Smart grid cyber-physical attack and defense: A review. *IEEE Access* 2021;9:29641-59. <http://dx.doi.org/10.1109/ACCESS.2021.3058628>.
- [13] Labrador Rivas AE, ao TA. Faults in smart grid systems: Monitoring, detection and classification. *Electr Power Syst Res* 2020;189. <http://dx.doi.org/10.1016/j.epsr.2020.106602>.
- [14] Kemal M, Sanchez R, Olsen R, Iov F, Schwefel H-P. On the trade-off between timeliness and accuracy for low voltage distribution system grid monitoring utilizing smart meter data. *Int J Electr Power Energy Syst* 2020;121. <http://dx.doi.org/10.1016/j.ijepes.2020.106090>.
- [15] Müller N, Afzal Z, Eliasson P, Ekstedt M, Heussen K. Threat scenarios and monitoring requirements for cyber-physical systems of energy flexibility markets. 2021. arXiv preprint [arXiv:2111.03300](https://arxiv.org/abs/2111.03300).
- [16] Singh D, Banyal RK, Sharma AK. Cloud computing research issues, challenges, and future directions. In: Rathore VS, Worrning M, Mishra DK, Joshi A, Maheshwari S, editors. *Emerging trends in expert applications and security*. Singapore: Springer Singapore; 2019, p. 617-23.
- [17] Yousefpour A, Fung C, Nguyen T, Kadiyala K, Jalali F, Niakanlahiji A, Kong J, Jue JP. All one needs to know about fog computing and related edge computing paradigms: A complete survey. *J Syst Archit* 2019;98:289-330. <http://dx.doi.org/10.1016/j.sysarc.2019.02.009>.
- [18] Pereira J, Silveira M. Unsupervised anomaly detection in energy time series data using variational recurrent autoencoders with attention. In: *2018 17th IEEE international conference on machine learning and applications (ICMLA)*. IEEE; 2018, p. 1275-82. <http://dx.doi.org/10.1109/ICMLA.2018.00207>.
- [19] Barua A, Muthirayan D, P. Khargonekar P, Al Faruque MA. Hierarchical temporal memory based one-pass learning for real-time anomaly detection and simultaneous data prediction in smart grids. *IEEE Trans Dependable Secur. Comput.* 2020. <http://dx.doi.org/10.1109/TDSC.2020.3037054>.
- [20] Hollingsworth K, Rouse K, Cho J, Harris A, Sartipi M, Sozer S, Enevoldson B. Energy anomaly detection with forecasting and deep learning. In: *2018 IEEE international conference on big data (big data)*. IEEE; 2018, p. 4921-5. <http://dx.doi.org/10.1109/BigData.2018.8621948>.
- [21] Neupane B, Pedersen TB, Thiessen B. Towards flexibility detection in device-level energy consumption. In: *International workshop on data analytics for renewable energy integration*. Springer; 2014, p. 1-16.
- [22] Mocanu E, Nguyen PH, Gibescu M. Energy disaggregation for real-time building flexibility detection. In: *2016 IEEE power and energy society general meeting (PESGM)*. IEEE; 2016, p. 1-5. <http://dx.doi.org/10.1109/PESGM.2016.7741966>.
- [23] Liu Y, Wu L, Li J. D-PMU based applications for emerging active distribution systems: A review. *Electr Power Syst Res* 2020;179. <http://dx.doi.org/10.1016/j.epsr.2019.106063>.
- [24] Lazzaretti AE, Ferreira VH, Neto HV, Toledo LFRB, Pinto CLS. A new approach for event classification and novelty detection in power distribution networks. In: *2013 IEEE power & energy society general meeting*. IEEE; 2013, p. 1-5. <http://dx.doi.org/10.1109/PESMG.2013.6672703>.
- [25] Lazzaretti AE, Tax DMJ, Vieira Neto H, Ferreira VH. Novelty detection and multi-class classification in power distribution voltage waveforms. *Expert Syst Appl* 2016;45:322-30. <http://dx.doi.org/10.1016/j.eswa.2015.09.048>.
- [26] Huang N, Fang L, Cai G, Xu D, Chen H, Nie Y. Mechanical fault diagnosis of high voltage circuit breakers with unknown fault type using hybrid classifier based on LMD and time segmentation energy entropy. *Entropy* 2016;18(9). <http://dx.doi.org/10.3390/e18090322>.
- [27] Geng C, Huang S-J, Chen S. Recent advances in open set recognition: A survey. *IEEE Trans Pattern Anal Mach Intell* 2021;43(10):3614-31. <http://dx.doi.org/10.1109/TPAMI.2020.2981604>.
- [28] Mahdavi A, Carvalho M. A survey on open set recognition. 2021. arXiv preprint [arXiv:2109.00893](https://arxiv.org/abs/2109.00893).
- [29] Niazazari I, Livani H. Disruptive event classification using PMU data in distribution networks. In: *IEEE Power & Energy Society, editor. 2017 IEEE power & energy society general meeting*. New York: IEEE; 2018, p. 1-5. <http://dx.doi.org/10.1109/PESGM.2017.8274154>.
- [30] Phillips D, Overbye T. Distribution system event detection and classification using local voltage measurements. In: *2014 Power and Energy Conference at Illinois (PECI)*. IEEE; 2014, p. 1-4. <http://dx.doi.org/10.1109/PECI.2014.6804576>.
- [31] Ahmad S, Lavin A, Purdy S, Agha Z. Unsupervised real-time anomaly detection for streaming data. *Neurocomput* 2017;262:134-47. <http://dx.doi.org/10.1016/j.neucom.2017.04.070>.
- [32] Yaacob AH, Tan IK, Chien SF, Tan HK. Arima based network anomaly detection. In: *2010 Second International Conference on Communication Software and Networks*. IEEE; 2010, p. 205-9. <http://dx.doi.org/10.1109/ICCSN.2010.55>.
- [33] Ergen T, Kozat SS. Unsupervised anomaly detection with LSTM neural networks. *IEEE Trans Neural Netw Learn Syst* 2019;31(8):3127-41. <http://dx.doi.org/10.1109/TNNLS.2019.2935975>.
- [34] Munir M, Siddiqui SA, Dengel A, Ahmed S. DeepAnT: A Deep learning approach for unsupervised anomaly detection in time series. *IEEE Access* 2019;7:1991-2005. <http://dx.doi.org/10.1109/ACCESS.2018.2886457>.
- [35] Ren H, Xu B, Wang Y, Yi C, Huang C, Kou X, Xing T, Yang M, Tong J, Zhang Q. Time-series anomaly detection service at microsoft. In: *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. ACM Digital Library, New York, NY, USA: Association for Computing Machinery; 2019, p. 3009-17. <http://dx.doi.org/10.1145/3292500.3330680>.
- [36] Scheirer WJ, de Rezende Rocha A, Sapkota A, Boulte TE. Toward open set recognition. *IEEE Trans Pattern Anal Mach Intell* 2013;35(7):1757-72. <http://dx.doi.org/10.1109/TPAMI.2012.256>.
- [37] Bendale A, Boulte T. Towards open world recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, 1893-1902.
- [38] Rudd EM, Jain LP, Scheirer WJ, Boulte TE. The extreme value machine. *IEEE Trans Pattern Anal Mach Intell* 2018;40(3):762-8. <http://dx.doi.org/10.1109/TPAMI.2017.2707495>.
- [39] Geng C, Chen S. Collective decision for open set recognition. *IEEE Trans. Knowl. Data Eng.* 2022;34(1):192-204. <http://dx.doi.org/10.1109/TKDE.2020.2978199>.
- [40] Gupta A, Gupta HP, Biswas B, Dutta T. Approaches and applications of early classification of time series: A review. *IEEE Trans Artif Intell* 2020;1(1):47-61. <http://dx.doi.org/10.1109/TAI.2020.3027279>.
- [41] Ziras C, Heinrich C, Bindner HW. Why baselines are not suited for local flexibility markets. *Renew Sustain Energy Rev* 2021;135. <http://dx.doi.org/10.1016/j.rser.2020.110357>.
- [42] Suel T. Delta compression techniques. *Encycl Big Data Technol* 2018;63:1-8. http://dx.doi.org/10.1007/978-3-319-63962-8_63-1.
- [43] Hawkins J, Blakeslee S. *On intelligence*. USA: Macmillan; 2004.
- [44] Wu J, Zeng W, Yan F. Hierarchical temporal memory method for time-series-based anomaly detection. *Neurocomput* 2018;273:535-46. <http://dx.doi.org/10.1016/j.neucom.2017.08.026>.
- [45] Deb C, Zhang F, Yang J, Lee SE, Shah KW. A review on time series forecasting techniques for building energy consumption. *Renew Sustain Energy Rev* 2017;74:902-24. <http://dx.doi.org/10.1016/j.rser.2017.02.085>.
- [46] Hyndman RJ. *Forecasting with long seasonal periods*. 2010. <https://robjhyndman.com/hyndsight/longseasonality/>. [Online; accessed: 23/07/2021].
- [47] O'Shea K, Nash R. An introduction to convolutional neural networks. 2015. arXiv preprint [arXiv:1511.08458](https://arxiv.org/abs/1511.08458).
- [48] Hou X, Zhang L. Saliency detection: A spectral residual approach. In: *2007 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE; 2007, p. 1-8. <http://dx.doi.org/10.1109/CVPR.2007.383267>.
- [49] Yu Y, Qu W-Y, Li N, Guo Z. Open-category classification by adversarial sample generation. 2017. arXiv preprint [arXiv:1705.08722](https://arxiv.org/abs/1705.08722).
- [50] Smith SW. *The scientist and engineer's guide to digital signal processing*. USA: California Technical Publishing; 1997.