Journal of Industrial Information Integration 50 (2026) 101040

Contents lists available at ScienceDirect

Journal of Industrial Information Integration

journal homepage: www.elsevier.com/locate/jii

Full length article

Automating appliance verification in facilities management using a denoised Voltage-Current feature extraction and classification pipeline

Socretquuliqaa Lee [a,b], Faiyaz Doctor [b,*], Mohammad Hossein Anisi [b], Shashank Goud [c], Xiao Wang [c]

[a] EKYC Solutions Co. Ltd, Phnom Penh, Cambodia
[b] School of Computer Science and Electronic Engineering, The University of Essex, Colchester, United Kingdom
[c] Cloudfm Group Ltd, Colchester, United Kingdom

ARTICLE INFO

Dataset link: https://github.com/Socret360/denoised-vi-pipeline

Keywords:
Appliance load monitoring
V–I trajectories
Feature extraction
Machine learning
Appliance classification
Facilities management

ABSTRACT

Facilities Management (FM) companies can use load monitoring of electrical appliances (assets) to track energy consumption and predictive maintenance. Reliable algorithms are needed to automatically identify or verify appliances through their energy signatures to improve efficiencies during installation and inspection tasks. Most approaches rely on Voltage-Current (V-I) trajectory. These features are extracted from steady-state current and voltage signals. However, these methods often assume signals are uniformly sampled. In real-world conditions, this assumption does not always hold, leading to misclassified steady-state events when signals are noisy. This paper introduces a novel feature extraction and classification pipeline to ensure the validity of detected steady-state events. The approach measures the approximate entropy of current signals and their correlation with voltage to extract denoised features for appliance type classification. The proposed pipeline is evaluated on a large-scale real-world operational dataset spanning multiple appliance categories. We demonstrate that the extracted denoised features significantly improve the performance of Machine Learning (ML) models used for appliance type classification. Finally, we present a deployment framework for FM settings, enabling digital cataloguing of appliances informing businesses on sustainable choices for appliance requirements.

1. Introduction

Global energy demand is projected to rise by 25%, with electricity demand increasing by 58% by 2040 [1,2]. To meet sustainability targets, businesses must adopt tools for managing their power usage and improving energy efficiency [3]. Advances in Appliance Load Monitoring (ALM) [4], supported by Artificial Intelligence (AI) and Internet of Things (IoT) devices, enables usage of appliance-level data for monitoring energy consumption, predicting maintenance needs and inducing behavioural changes towards more efficient energy use [5].

Facilities Management (FM) companies lead this effort by deploying large scale digital infrastructures across sites such as restaurant chains, warehouses, hospitals, offices and casinos. These systems collect and analyse electronic appliance-level data from a wide range of appliances—from lighting systems and HVAC units to kitchen equipment. However, verifying that these appliances are correctly identified during the installation of ALM units remains a major challenge. Manual verification is time-consuming and prone to human error, particularly at sites containing hundreds of assets. Automated algorithms capable of accurately verifying and classifying appliances based on their load signatures are therefore essential for improving installation efficiency and data reliability.

Developing such automated verification capabilities, however, presents significant technical hurdles. Most existing approaches perform appliance identification using Voltage-Current (V-I) trajectory features extracted from steady-state signals [6,7]. These methods assume that data are uniformly sampled and noise-free. In practice, however, signals collected from real-world FM systems are often non-uniform due to compression or missing points. As a result, traditional steady-state detection algorithms can misclassify noise as valid steady-state events, reducing classification accuracy. Similar to how advances in civil engineering employ self-sensing composite materials to provide reliable real-time data for structural health monitoring [8], our work addresses the need for dependable data validation in complex energy monitoring systems.

* Corresponding author.
E-mail addresses: socretlee@ekycsolutions.com (S. Lee), fdocto@essex.ac.uk (F. Doctor), m.anisi@essex.ac.uk (M.H. Anisi), s.goud@cloudfmgroup.com (S. Goud), x.wang@cloudfmgroup.com (X. Wang).

https://doi.org/10.1016/j.jii.2025.101040
Received 3 February 2025; Received in revised form 10 November 2025; Accepted 15 December 2025
Available online 16 December 2025

To overcome these limitations, this study introduces a new denoised feature extraction and classification pipeline for automatically identifying appliance types. The proposed approach measures the Approximate Entropy (ApEn) of steady-state current signals and their correlation with voltage to assess the regularity of load patterns. This enables the removal of noisy or irregular segments that could otherwise distort the extracted V–I trajectory features. Due to ApEn's computational efficiency [9], this proposed extension of current steady-state detection algorithms enables an additional layer of effective noise filtering with minimal impact on pipeline performance.

This study therefore decreases the gap between methodological innovation and real-world deployment within the FM sector. Our approach complements recent advancements in AI and IoT driven energy-efficient infrastructure, where deep learning and intelligent monitoring systems are used to enhance sustainability and operational decision-making in smart buildings [10,11]. These studies emphasise the importance of robust data preprocessing and feature validation for reliable deployment—objectives that our denoised feature extraction and classification pipeline directly address. Our contributions in this paper are as follows:

1. Introducing a novel application of ApEn for detecting and removing highly unstructured current and voltage based steady-state events for extracting asset V-I trajectory features.
2. Developing a denoised V-I trajectory feature extraction approach as a new preprocessing method for handling non-uniformly sampled signals in real world load monitoring data from assets operating in commercial environments.
3. Using the developed denoised asset feature extraction approach to enhance the performance of ML classification models uniquely trained and evaluated on large-scale FM asset load data spanning multiple commercially deployed appliances.
4. Proposing a new automated FM asset verification pipeline using denoised V-I trajectory feature based appliance load classification supporting on-site ALM unit installation.

This paper is structured as follows. Section 2 outlines related work. In Section 3 the concepts of V–I Trajectories and Approximate Entropy are discussed. The FM commercial use case and challenges and described in Section 4. Section 5 details the proposed feature extraction and classification pipeline. We discuss our experimental setup and provide analysis of the results in Section 6. In Section 7 we discuss the wider contributions of the proposed approach, potential deployment scenarios and industrial impacts within FM. The paper ends with conclusions and a discussion on future work in Section 8.

2. Related work

Appliance classification approaches can be found in the literature of NILM which refers to the task as Load Identification [7]. These approaches mostly rely on V-I Trajectories based features, typically extracted from the voltage and current steady-state waveforms [12] of the appliance's operating load, and used as inputs to ML and DL algorithms for classifying its type.

Previous feature extraction approaches have been based on extracting labelled hand-crafted features from the V-I Trajectory plots [13,14], meshing the plots into binary images to identify contour features using methods such as Elliptical Fourier Descriptors, Principle Component Analysis [15–18] or the use of mathematical feature extraction methods such as 2-D Fourier series [12]. The extracted features have then been used to train supervised ML classifiers based on Artificial Neural Networks, Linear Quadratic Discriminant Analysis, Support Vector Machine, Decision Trees, Naïve Bayes classifiers and Ensemble approaches such as Random Forests, Adaptive boosting and Gradient Boosting.

Recent developments in the field of DL, especially its successes in Computer Vision, have led to further interest in the treatment of V-I Trajectories as 2D binary, colour images or 3D temporal spacial trajectory images that can be used for training Convolutional Neural Network (CNN) based models [18–22]. The consensus view across these works is that CNNs are able to automatically extract useful latent features from V-I Trajectories without the need to perform manual feature extraction. While the results from supervised based approaches are promising, the introduction of a new/unknown type of appliance into the system would require these model to be periodicity retrained. Work done on unsupervised methods looks to tackle this issue through applications of clustering based ML and DL algorithms. Here approaches such as hierarchical clustering have been used to classifying appliance V-I Trajectories [23] and the combination of neural networks with density based clustering for extracting low dimensional features of V-I Trajectories which can be grouped as known or undefined appliance types [24]. More recently semi-supervised learning method have been explored in [25] to help identify appliances with limited labelled data. Across the aforementioned works, most of the approaches have been trained and tested on relatively small open source load data mostly gathered from consumer electronic appliances in controlled environments or simulated settings. Examples of these publicly available datasets include Reference Energy Disaggregation Data Set (REDD) [26], Plug Load Application Identification Dataset (PLAID) [27,28], Worldwide Household Industry Transient Dataset (WHITED) [29], together with others such as Controlled On/Off Loads Library (COOLL) [30], LIT [12] and BLOND [31]

The V-I trajectory extraction approaches adopted in previous works make an assumption that the signals data are uniformly sampled and of high frequency [7]. Based on this assumption, steady-state detection algorithms can rely on sliding window approaches to determine sharp increase and decrease in a signal (e.g. apparent energy) as the start and end of an event. In real-world applications, the signals data may not be uniformly sampled. This is often the results of data compression/reduction algorithms that record only the most important data points. In such environments, steady-state detection algorithms can misclassify noise as an event since the points that mark the sharp increase and decrease do not always signify a steady-state. Furthermore, because the publicly available open source datasets used in these works are small, the raw signals gathered are often short (< 10 s). This results in limited steady-state events that can be detected in a signal, leading to fewer V-I trajectories being extracted. The limited amount and diversity of samples therefore may not reflect the variety of appliance types and operating conditions found in real-world environments. Hence, the results reported by those authors may be too optimistic for such scenarios.

In the area of signal processing, many techniques have been introduced for improving signal clarity and reliability. State-of-the-art methods include wavelet transform denoising for multi-resolution filtering, empirical mode decomposition (EMD) and its variants (like EEMD and VMD) for adaptive decomposition of non-stationary signals, and Kalman filtering for real-time, model-based noise estimation [32–34]. Additionally, DL based approaches such as Denoising Autoencoder (DAE) and CNN have shown good performance in learning complex noise characteristics directly from data [35]. In NILM, these techniques form the backbone of preprocessing and load disaggregation pipelines. However, while these methods focus on removing noise, they do not provide metrics to assess whether a detected event truly represents a valid steady-state. ApEn presents a complementary solution: instead of removing noise before event detections, ApEn quantifies signal regularity, providing a statistical measure of complexity in the current waveform which can be used to remove noise after event detections. By computing ApEn on the current waveform, NILM systems can verify whether it is a stable and repeatable load pattern—offering an additional layer of validation to reduce false positives and improve event detection reliability [36,37]. Importantly, ApEn is also computationally efficient, requiring only a short window of data and simple distance

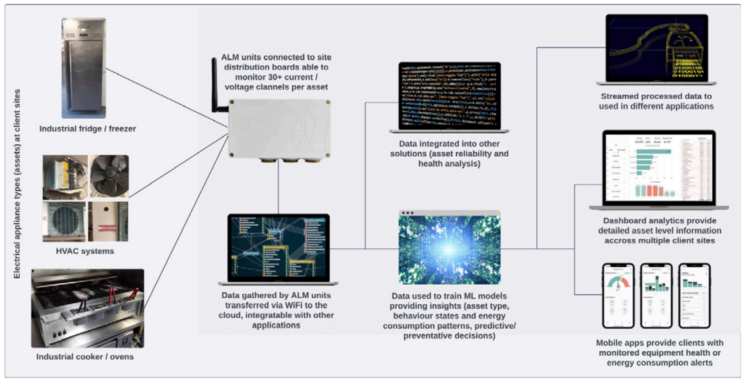

Fig. 1. FM digital platform supported by load monitoring units.

comparisons between patterns [9]. This makes it a practical and scalable addition to large-scale appliance classification pipelines, where maintaining low-latency and high-throughput processing is critical.

Commercial applications for appliance load identification which are integral to monitoring the operation of facilities such as warehouses, hospitals or commercial kitchens has not been extensively explored. Previous work in [38] installed NILM system in three commercial building types (office, barracks and vehicle maintenance) and evaluated its ability to disaggregate electric loads pertaining to a combination of servicing and specialised appliances (computing equipment, lighting, sockets, vehicle service bays, battery chargers, heating systems). However, the analysis failed to identify specific loads which authors attributed to the number and complexity of loads (including hundreds of small and similar loads) making identification ambiguous, difficulty in interpreting small changes in power consumption, and inability to identify continuously operating loads. In [39] authors proposed an NILM method based on a transformer neural network multi-label classification model for commercial load monitoring. Here a publicly available dataset [40] based on data collected from smart meters deployed in different buildings was used where data on electricity consumption was recorded per building, high energy loads, and total electricity consumption per floor with loads related to air conditioning, lighting, and elevators. Both these examples although relevant have again been limited to specific and specialised appliance types. A large dataset on commercial electronic appliances similar to those found across the FM sector was recently published in [41]. Here electricity consumption data has been collected from three restaurant kitchens during daily operations where 45 types of consumer and commercial electronic appliances were monitored. However, the individual appliance consumption data was collected at 0.2 Hz. This makes it unsuitable for V–I trajectories reconstruction which requires the instantaneous values of voltage and current to be able to plot the dynamic relationship between them.

Building on our initial work [42], we present and evaluate a feature extraction and classification pipeline to extract V-I trajectory features from correctly classified steady-state events. We use ApEn of the steady-state current and correlation between steady-state current and voltage to determine the validity of a detected steady-state event for V-I trajectory extraction. Using our proposed pipeline, we process a large scale FM dataset collected from commercial assets operating across several sites consisting of 86,268 samples of V-I trajectories based features from 29 appliance types. We validate the correctness of extracted features in the processed data by using it to compare well known and state-of-the-art ML algorithms in the task of appliance type classification. The next sections discusses in details our proposed pipeline.

3. Preliminaries

This section discusses the essential concepts and background information that our paper is built upon. In particular, the first subsection describes the ALM device which is the source of the data used for this study. Then, the second and third subsections explain the concepts of V–I trajectories and approximate entropy, which are core components of our feature extraction pipeline.

The UK FM industry is projected to be worth over $52 billion by 2027 [43]. A major FM company Cloudfm Group Ltd has developed an industry-leading Internet of Things (IoT) based ALM unit called PRISM® that is supported by a digital platform called Mindsett which provides an end-to-end solution combining AI and data analytics visualisation tools for appliance-level energy management and predictive maintenance.

3.1. Load monitoring data collection

The PRISM® ALM hardware unit connects directly to a site distribution board. A single unit consists of 36 channels that can be used to connect different appliances. For each appliance, the signal information that the ALM unit collects includes: (i) Apparent Energy (*VA*), (ii) Real Energy (*W*), (iii) Current of the Given Circuit (*irms*), (iv) Voltage Frequency (*vfreq*), (v) Voltage of the Power Supply (*vrms*), (vi) The fundamental harmonic of the voltage (*vh1*), and (vii) The harmonic data up to the 10th component (*Hi1–Hi10*, *Hq1–Hq10*). According to Cloudfm Group Ltd, through the use of this information along with the appliance type, a comparison can be made among different brands/models within an appliance type which has huge decision making benefits for customers. In addition, other metadata includes: (i) The appliance type (*thing_type*), (ii) The channel in which the appliance is connected (*channel*), (iii) The unique identifier for an appliance (*nid*). Fig. 1 illustrates the overview of how the ALM unit is used as part of the company's FM supporting digital platform.

3.2. V–I trajectories

In 2007, Lam et al. [23] demonstrates that unique appliance signatures can be created by examining a curve made up of one phase current and voltage of the appliance in a steady-state event. The unique signature is often based on the shapes of the V–I Trajectories and as shown in Fig. 2, different appliance types exhibit different trajectory shapes. Earlier work describes the trajectory shapes using hand-crafted features (e.g. area of the curve, curvature of mean line). Recent developments treat the trajectory as images to be used with Convolutional Neural Networks (CNN) [13–15,17,18]. As mentioned, extraction of V–I trajectories from a raw signal is dependent on the detection of steady-state events in that signal. Multiple approaches exists for determining steady-state events [13,44]. The most common is to detect a period in between a sharp increase and decrease in apparent energy. Fig. 3 shows an example of a detected steady-state event and the resulting V–I trajectory.

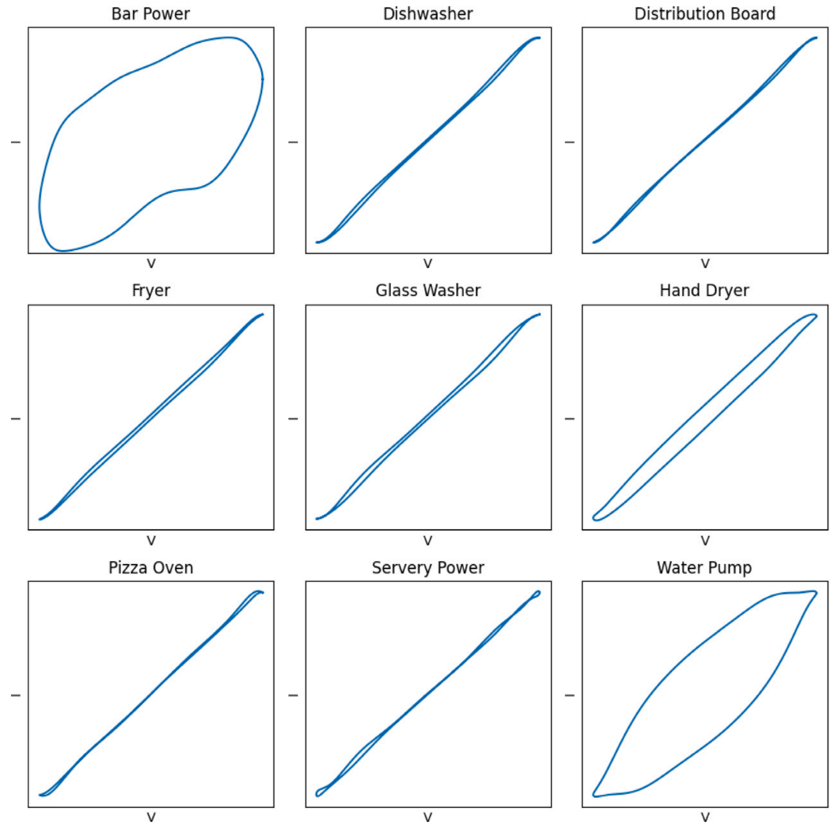

Fig. 2. Samples V–I Trajectories of Appliance Types.

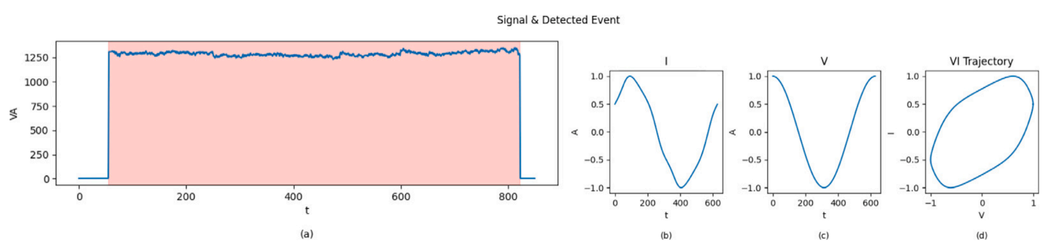

Fig. 3. Detected Event and Resulting VI Trajectory. (a) The shaded pink area signify the detected steady-state event using apparent energy (VA). (b) The Extracted Current Waveform. (c) The Extracted Voltage Waveform. (d) The Extracted V–I Trajectory.

3.3. Approximate entropy

$$\mathrm{ApEn}(m, r, N)(u) = \phi^{m}(r) - \phi^{m+1}(r) \tag{1}$$

$$\begin{aligned}
\text{where }\ \phi^{m}(r) &= \frac{1}{n}\sum_{i=1}^{n} log C_i^m(r) \\
C_i^m(r) &= \frac{\text{number of } j \text{ such that } d\left[\mathbf{x}(i), \mathbf{x}(j)\right] \le r}{n} \\
d\left[\mathbf{x}(i), \mathbf{x}(j)\right] &= \max_k \left(|\mathbf{x}(i)_k - \mathbf{x}(j)_k|\right) \\
\mathbf{x}(i) &= [u(i), u(i+1), \dots, u(i+m-1)] \\
n &= N - m + 1 \\
&\text{for } 1 \le k \le m;\ 1 \le i, j \le n;\ m \le N; \\
&r \in \mathbb{R}^+;\ m \in \mathbb{Z}^+
\end{aligned}$$

Approximate Entropy is a technique used for measuring the rate of regularity of a time series [45]. It was first introduced by Pincus back in 1991 [46]. Given a time series u, a sliding window of size m, and a filtering level r, the approximate entropy of u is defined as Eq. (1). It measures the logarithmic likelihood that patterns in a time series will remain similar when extended by an additional data point [46]. Fig. 4 shows the ApEn values of a time series with repeating pattern and a time series that has less repeating pattern. Fig. 4a shows a sine wave with four repeating cycles created through $y = sin(x)$ where $x = \frac{360}{8\pi}t; 0 < t < 8\pi$. Fig. 4b is created by randomly shuffling the sequence in Fig. 4a to create a sequence with few repeating patterns. A time series with repeating patterns will result in an ApEn close to zero. In contrast, a time series with few repeating pattern will have high value of ApEn. Hence, ApEn provides a useful mechanism for measuring noise in a signal.

4. Commercial use case and challenges

To illustrate the scale of the industry, Cloudfm Group Ltd serves FM clients ranging from restaurants, warehouses, 24/7 outlets, commercial properties, manufacturing plants and international clients. The number of appliances (also termed assets) at these multi-site establishments can range from 10 to 14,000. For instance a client in the amusement and recreation sector, produces data from 13,755 appliances (over 150 appliance types) operating on a 24 h day cycle across 40 sites in the UK.

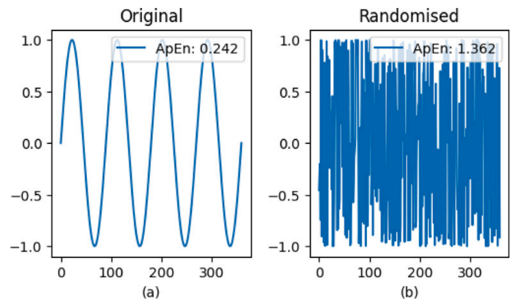

Fig. 4. Approximate Entropy calculation of $y = sin(x)$ (a) and its randomised sequence (b).

Depending on the number of appliances to be monitored per site, the installation of multiple ALM units are required, with each supporting up to 36 channels for monitoring different appliances. For another client, a casual dining restaurant chain, 482 appliance types operate across 101 spaces across 6 sites supported by 27 ALM units. Fig. 5 shows the logged entries for assets that includes meta information such as appliance category, type, phase information, circuit and distribution board they are connected to at a given site. Also shown are the ALM unit ids linked to a given distribution board and the assets connected to each channel of the ALM unit.

The process of installing ALM units starts when the FM mobilisation team sends installation surveyors to visit a given site to assess the number of ALM units that would be required, based on the number of distribution boards, the number of appliances that need to be monitored and their connectivity status via Wi-Fi coverage in the vicinity of each distribution board. Based on this initial assessment the distribution boards and assets each board is powering are logged into a shared data and workflows repository. Fig. 5 shows the logged entries for assets that includes meta information such as appliance category, type, phase information, circuit and distribution board they are connected to at a given site. The list of assets connected to each distribution board are normally pre-marked on the board by site operators. This is currently the only check for verifying that an asset of a particular type is connected to a board. During the on-site installation of ALM units, installation engineers visit the site to install the ALM units. Here they connect an ALM unit's channels to a distribution board and check if the ALM unit is transmitting signals. They then recheck the shared data repository to determine if the channel number listed there corresponds to the correct asset name based on those which were marked on the distribution board. Following installation of the ALMs which usually takes place during non-trading hours, the mobilisation team liaises with the hardware and data science teams to check the quality and correctness of received data (frequency of data points, phase mapping, etc.).

Some commercial appliances relay on a three-phase power supply where the current and voltage from each phase needs to be correctly mapped to the ALM channels to monitor the asset load and performance. Here mistakes can arise in how the phases are mapped. If this mapping is incorrect the installation engineer will have to revisit the site to correct it. During the mapping of phases, the placement of the Kct-16 current transformers used to monitor an asset's current, power and energy may be positioned incorrectly to collect signal information resulting in incorrect data being transmitted. The asset lists on the distribution board can also be outdated depending on when the last update was done by the site. There can be shifts between the plugs, and assets can be replaced with different appliances than those that are listed resulting in incorrect signals. These signal anomalies have to be manually identified by the hardware and data analytics teams by examining patterns in the signal harmonics for specific asset data streams flagged as needing further inspection.

For large or multi-site businesses with 100 s of connected assets requiring multiple ALMs to be deployed, the combination of these potential errors and manual inspection signal anomalies can prolong the installation process as there is no framework for consistently and sustainably verifying the correctness of each asset's type. Automatically verifying that load signal information being collected corresponds to a correct, incorrect or unidentifiable asset type can enable the site engineer to quickly recheck asset lists and their connections to specific ALM unit channels saving cost and time.

5. Proposed feature extraction and classification pipeline

This section details our proposed feature extraction and classification pipeline that has been applied on a real-world dataset provided by Cloudfm Group Ltd. Hence, the first part of this section discusses the original dataset and the cleaning procedure performed on the dataset to prepare for feature extraction. The second part details the feature extraction and ML pipeline.

5.1. Original dataset

Access to raw signals data collected from ALM units were provided which contained high frequency signals gathered from 391 appliances across 60 types over the span of six months from seven site locations. Locations comprised of various sites for a restaurant chain. The distribution of appliance type is highly imbalanced since a single location does not have all appliance types installed. The raw signals are grouped by the date and are separated into 1 week chunks. The appliance type label is stored in a different file which can be associated with the raw signal data by their *channel* in combination with their *nid*. The information in the label file, aside from the *channel* and *nid*, are entered manually by site engineers.

Data cleaning procedures were done on the label file and the files containing weekly raw signal data. For the label file, columns were stripped of leading and trailing spaces and rows with missing *thing_type* were removed. This resulted in a reduction in the number of appliances, their unique types, and locations. For the files containing weekly raw signal data, columns with null values were dropped. Furthermore, information in the label file were then joined with the raw signal data based on their *channel* and *nid*. The raw signal data was the regrouped by the appliance *channel* and *nid*. In particular, the weekly data was separated into different files where each file contained the raw signal data of one appliance on a single day. This was done to isolate each appliance's signals for more convenient processing in later stages of our preprocessing pipeline. In addition, the metadata information of each file was recorded as entries in a separate csv file.

5.2. Denoised V-I trajectories extraction approach

The feature extraction approach can be divided into two steps: (i) Extracting V–I Trajectories and (ii) Extracting shape features from V–I Trajectories. In the first step, for each of the per-appliance signal data files, outliers were removed from the apparent energy (VA) signal based on moving median replacement. More specifically, the signal is first divided into partitions. Then, a sliding window is used to replace points outside of $(Q1 - 1.5IQR, Q3 + 1.5IQR)$ with the median of the window. These steps are repeated for partitions of sizes 12.5%, 25%, 50%, and 100% of the signal length. Then, the filtered VA signal is fed into the steady-state detection algorithm outlined in [44]. Since the detection was done on a full day of data, multiple steady-state events could be detected for each signal data file. Here analyses was restricted to the first ten detected events due to computational resource constraints. For each of the events, the average of 2 cycles of voltage and current at the start and end of the event were extracted and marked as V_a, V_b, I_a,

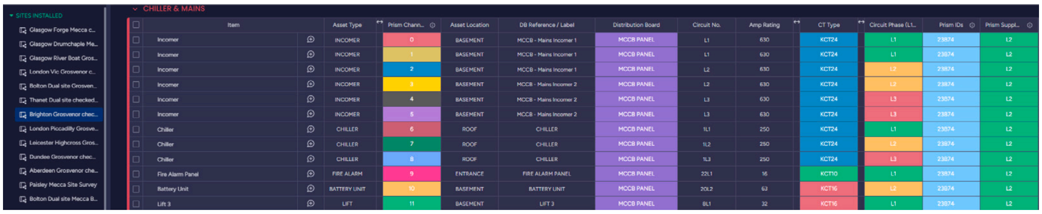

Fig. 5. FM digital platform supported by load monitoring units.

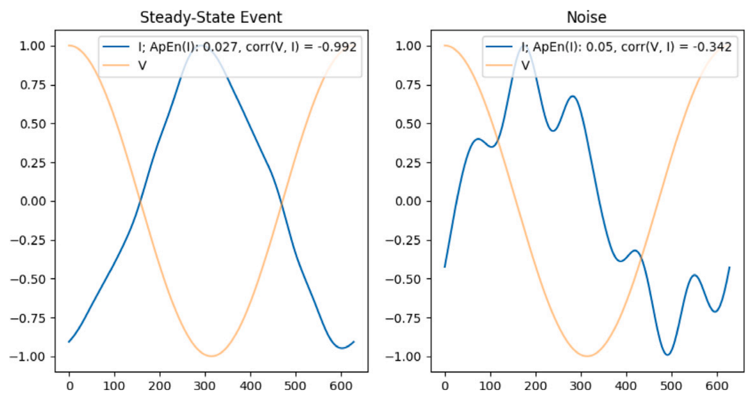

Fig. 6. Characteristics of correctly and incorrectly detected events.

and I_b. Then, Eq. (2) was used to calculate the steady-state voltage and current.

$$V = (V_a + V_b)/2;\ I = I_b - I_a \tag{2}$$

$$IQR = Q3 - Q1 \tag{3}$$

$$\text{ApEn}(m, r, N)(I) > 0.03 \vee |\text{corr}(V, I)| < 0.5 \tag{4}$$

where $Q1 =$ The first percentile of the signal.

$Q3 =$ The third percentile of the signal.

$$\text{corr}(V, I) = \frac{\sum_{i=1}^{N}(V_i - \bar{V})(I_i - \bar{I})}{\sqrt{\sum_{i=1}^{N}(V_i - \bar{V})^2}\sqrt{\sum_{i=1}^{N}(I_i - \bar{I})^2}}$$

$$m = 2;\ r = 0.2$$

As mentioned in Section 2, the steady-state detection algorithms can sometimes select the wrong data points as start and end points. We observed that these misclassified events result in steady-state current that is highly unstructured. To measure this, we calculate its approximate entropy (ApEn(I)) and its correlation with the steady-state voltage (corr(V, I)). Our experiment showed that misclassified events tend to have the characteristics shown in Eq. (4). This aligns with the work of Yan and Gao [9] in which a similar ApEn(I) threshold was reported where for a healthy bearing, sampled at 26 kHz with 1000 data points, the ApEn(I) value was approximately less than 0.038. Fig. 6 shows a comparison between steady-state current and voltage in correctly classified and misclassified events. Based on this, we discarded any events that matched the condition in Eq. (4). Finally, shape features were extracted from plots created by the remaining steady-state voltage and current pairs. We discuss the details of this extraction process in the next section.

5.3. Rationale for approximate entropy and fixed threshold

We motivate our use of ApEn as a noise-robust statistical measure because of two factors. First, its computational efficiency is particularly suitable for large-scale FM applications. Its lightweight computation — based only on local distance comparisons within a short window — ensures scalability and low-latency processing for both streaming and batch-mode feature extraction [9]. Second, ApEn's ability to quantify complexity without relying on any assumed signal distribution or frequency model makes it broadly applicable across different electrical environments and device types [46]. This is unlike conventional denoising or filtering techniques such as wavelet transform, empirical mode decomposition, or Kalman filtering, which require prior assumptions about the signal's spectral content or stationarity [32,33,45].

Building on this justification, we also opted to use a fixed ApEn threshold (0.03 in this study) rather than an adaptive one for three key reasons. First, empirical testing as detailed in Section 6.6 shows that this threshold effectively separated valid steady-state events from misclassified noisy segments across a wide range of datasets, appliance types and operating conditions, consistent with the findings of Yan and Gao [9]. Second, adaptive thresholding approaches typically require iterative recalibration or dynamic windowing, which introduces additional computational overhead and instability during batch processing of high-volume data streams [34,35]. Third, the non-uniform and high-variance sampling characteristics of real-world signals can cause adaptive thresholds to fluctuate excessively, potentially discarding valid but temporally irregular events [38,39]. The fixed threshold therefore provides a balance between generalisation and computational simplicity—maintaining interpretability, reproducibility, and consistent denoising performance across diverse datasets.

The effectiveness of using ApEn in this way aligns with broader findings in the energy-efficiency literature, where robust noise handling and feature selection have been shown to enhance the performance and scalability of ML-based monitoring systems. For example, Sheela et al. [10] demonstrated that incorporating noise-aware preprocessing significantly improves prediction accuracy in smart building energy management systems using ML and IoT frameworks. Similarly, our use of ApEn and fixed thresholding ensures that only statistically stable and physically meaningful steady-state events are retained for V–I trajectory feature extraction, thereby improving classification reliability and reproducibility in real-world FM deployments.

5.4. V–I trajectories shape features extraction

A total of 10 features were extracted from the V–I Trajectories based on those identified in [44]: (i) Current Span (*itc*), (ii) Area (*ar*), (iii)

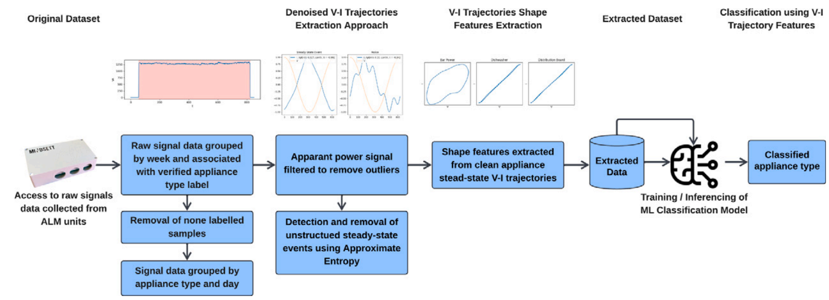

Fig. 7. Feature extraction and classification pipeline for asset verification.

Area with loop direction (*lpa*), (iv) Asymmetry (*asy*), (v) Curvature of mean line (M), (vi) Self-intersection (*sc*), (vii) Peak of middle segment (*mi*), (viii) Shape of middle segment (*sh*), (ix) Area of left and right segments (*alr*), and (x) Variation of instantaneous admittance (D). Aside from *itc*, the other features were calculated using normalised V–I trajectories. Furthermore, for some features the calculation method adopted by Wang et al. [44] rely on the points in the trajectory to be sorted. Hence, prior to the calculation of those features, we sorted the points in the trajectory based on their distance and direction from each other. In particular, starting from the point with the maximum steady-state voltage (*vmax*), the next point in the trajectory was the closest point in the same direction within a 40 degree angle.

5.5. Extracted dataset

Our proposed pipeline generated a dataset of 103,120 samples with 43 appliance types. However, due to the imbalanced nature of the dataset, some appliance types made up less than 0.1% of the total samples (e.g. Blind Supply, Cooling Tunnel Condenser, External Illuminate Signage, etc.). We opted to remove those appliance types from the dataset. In addition, appliances labelled as *Appliances*, *Mains*, and *Sockets* did not reflect an underlying appliance type. For those labelled as *Appliances*, we opted to split these into two based on their *thing_name* which are *Servery General* and *Post Mix Coke Machine*. Appliances labelled as *Mains* and *Sockets* were removed from the dataset. The final dataset consisted of 86,268 samples with 29 appliance types. Table 1 shows the distribution of the 29 classes which represent electronic appliances ranging from commercial chillers and pizza ovens to consumer devices such as coffee machines, indoor lighting and fridges. These represent a fraction of the various appliance types found across FM clients business users.

5.6. Classification using V-I trajectory features

The extracted dataset comprising of the 10 V–I trajectory features was used to classify the 29 appliance type classes. Here data was used to train and evaluate different supervised ML classification models after which inference could be performed on the trained model to predict the class probabilities of extracted trajectory features of unseen appliance signals. Data normalisation, in particular Min-Max normalisation was also performed on the extracted features. Fig. 7 illustrates the entire feature extraction and classification pipeline.

6. Experiments and results

To verify the quality of the extracted denoised features dataset, we opted to train and test different classifiers in correctly identifying appliance types from the extracted V–I trajectories. This section outlines our experimental setup by discussing the algorithms used, our train/test splits creation strategy, evaluation metrics, hyperparameter tuning procedure, and cross validation strategy. We then present the results of how the algorithms perform, comparisons with other state-of-the-art approaches and public datasets.

Table 1
Appliance type distributions of clean and noisy training set.

	Clean set	Training set with Noise	
		Noise samples	Noise ratio
AC	2205	884	0.401
Bar power	1069	117	0.11
Chiller	485	243	0.503
Clean power	919	424	0.462
Coffee machine	394	107	0.272
Dishwasher	12 046	156	0.013
Distribution board	6359	540	0.085
Fire alarm	104	104	1.0
Fridge	796	177	0.223
Fryer	15 115	120	0.008
Glass washer	5227	83	0.016
Grease trap	276	1	0.005
Hand dryer	9351	411	0.044
Hoist	111	25	0.231
Ice machine	918	136	0.149
Indoor lighting	745	745	1.0
Microwave	1036	221	0.214
Oven	1645	41	0.025
Overdoor heater	539	539	1.0
Pasta boiler	6432	12	0.002
Pizza oven	8456	329	0.039
Post mix coke machine	1003	14	0.014
Power	366	366	1.0
Servery general	654	43	0.066
Servery power	1316	205	0.156
Walk in freezer	3221	119	0.037
Walk in fridge	377	248	0.659
Water heater	2641	126	0.048
Water pump	2462	59	0.024

6.1. Algorithms for appliance type classification

The choice of ML classification algorithms are based on what have been proposed in the literature so that the results can be compared back to those sources. In particular, five algorithms are chosen: Decision Tree (DT), Random Forests (RF), k-Nearest Neighbours (kNN), Multi-Layer Perceptron (MLP) and Extreme Gradient Boosting (XGB) which has been shown to perform well against other classical ML classifiers applied to Non-Intrusive Load identification from V-I Trajectories [12].

6.2. Train/test splits

We split the extracted 86,268 samples into training (70%) and testing (30%). Stratified sampling was adopted to ensure that the two splits had the same appliance type distribution as the original dataset. We used the training set to perform hyperparameter tuning and cross validation. The test set was used to determine the final performance of the best performing hyperparameters on unseen data. Furthermore, to demonstrate the efficacy of our approach, a training set with noise was created to illustrate the performance differences between models trained with and without noisy samples. The number of samples and class distribution in this new set was equal to the clean training set.

Table 2
Results on test set.

	Training data	Acc	Precision	Recall	F1	mAP
XGB	Clean	0.906	0.884	0.841	0.860	0.871
	With noise	0.887	0.831	0.738	0.755	0.785
RF	Clean	0.903	0.890	0.841	0.863	0.863
	With noise	0.884	0.839	0.719	0.736	0.785
DT	Clean	0.861	0.808	0.786	0.796	0.740
	With noise	0.840	0.695	0.667	0.665	0.631
kNN	Clean	0.875	0.775	0.779	0.775	0.765
	With noise	0.866	0.690	0.708	0.697	0.686
ML	Clean	0.748	0.700	0.636	0.641	0.705
	With noise	0.748	0.608	0.545	0.550	0.623

Table 1 shows the noise ratio of each appliance type class in the training set with noise.

Conventional classification metrics were used to measure the performance of different models trained using data with and without noise removal. More specifically, Accuracy, Precision, Recall, F1-Score, and Mean Average Precision (mAP) and Average Precision (AP) where used. We primarily focus on AP and mAP as AP measures the area under the precision–recall curve for a specific appliance class, while mAP provides an overall measure by averaging AP across all appliance classes, offering a more comprehensive evaluation of classification performance.

6.3. Validity of the extracted dataset

Table 2 presents the performance of all classifiers on both clean and noisy test samples, revealing several key observations on the effectiveness of the extracted denoised dataset.

First, the mAP of all models trained with noise removed is larger than 60% while the highest performance achieved by *XGB* reaches 87%. This suggests that the extracted denoised dataset does consist of patterns that are useful for classifying different appliance types, confirming the applicability of V–I trajectory features as a useful representation of electrical appliances.

Second, the average mAP across models dropped by approximately 8.7% when noise was reintroduced, clearly demonstrating the benefit of the proposed denoising step in filtering out unstructured signal components. To further quantify this effect, model performance was analysed in terms of classification error reduction, defined as $1 - mAP$. This analysis highlights how much incorrect prediction each model eliminated after entropy-based denoising. The results show that *XGB*, *RF*, and *DT* reduced their classification errors by 40%, 36%, and 30%, respectively, indicating that ensemble and hierarchical models are particularly effective at leveraging the cleaner and more statistically regular V-I trajectory features. In contrast, *kNN* and *MLP* achieved smaller reductions of 25% and 22%, reflecting their continued sensitivity to residual feature irregularities and local noise. These findings align with prior research showing that tree-based models are inherently more robust to outliers, capture non-linear feature dependencies more effectively, and benefit from built-in regularisation mechanisms compared with distance-based or shallow neural approaches [19,21,44,47,48].

Third, while the denoising step substantially improves classification performance, it also introduces a trade-off. Overly aggressive filtering, as applied to remove unstructured current and voltage segments (Section 5.2), may risk discarding valid yet temporarily irregular steady-state events. Such balancing between predictive accuracy, computational efficiency, and potential data loss is a common consideration in AI-driven energy-efficient systems [11]. In this study, however, the observed performance gains and the resulting operational efficiency benefits far outweigh these risks.

Finally, the close alignment between cross-validation and test results further indicates strong generalisation capability of the models trained on the denoised dataset, reinforcing the stability and scalability of the proposed feature extraction and classification pipeline.

6.4. Algorithm performances

Table 3 shows the various classifier performances on test samples for each class in the dataset. These results provide two performance insights of the algorithms tested. First, tree based methods (*XGB*, *RF* and *DT*) perform the best compared to *kNN* and *MLP*. An examination of the minority class (*Fire Alarm*) shows that *XGB*, *RF* and *DT* all produce high AP for that class compared to the other two models. This suggests that the low performance of *kNN* and *MLP* is due to prediction bias towards the majority class. In the case of *kNN*, Table 3 shows clearly that the appliance types with the highest AP are the majority classes *Fryer* and *Hand Dryer* while the lowest AP class is the minority class *Fire Alarm*. Secondly, when examining Table 3, the class type with the lowest AP across all four models is *Servery General*. Deeper analysis of the misclassifications of *Servery General* as shown in Table 4 shows that they are mostly misclassified as *Distribution Board* and *Servery Power*.

6.5. Hyperparameter tuning and cross validation

We performed hyperparameter tuning using a Bayesian Optimisation approached introduced by Bergstra et al. [47]. The optimisation process was run 20 times for each algorithms resulting in 20 different tested configurations. For each hyperparameter configuration, we performed k-fold cross validation where k was set to 5. The best configuration was selected based on its average F1-Score across the different splits. Table 5 shows the performances of the best configuration of each model across the 5 splits. The values in parentheses represent the standard deviation which can be further used to determine the confidence interval by multiplying by 2. The best hyperparameter configuration was then used to train a model on the full training set and evaluated on the test set.

6.6. Thresholds sensitivity analysis

We analyse the sensitivity of the selected $\mathrm{ApEn}(I)$ and $\mathrm{corr}(V, I)$ by performing a full 10×9 grid search over $\mathrm{ApEn}(I) \in [0.01, 0.10]$ and $|\mathrm{corr}(V, I)| \in [0.1, 0.9]$ resulting in 90 combinations. For each pair, a subset of the training set with noise, and a *DT* classifier was evaluated under 5-fold cross-validation. The mean and standard deviation of F1 scores were recorded to assess robustness. The optimal thresholds ($\mathrm{ApEn}(I) = 0.03$, $|\mathrm{corr}(V, I)| = 0.5$) achieved the highest mean $F1 = 0.72(\pm 0.06)$, and lie within a broad plateau of stable performance (Fig. 8). This region represents a good balance between excluding noisy signals and retaining sufficient valid samples for reliable learning.

6.7. Comparative performance analysis with existing approaches

We further validate our proposed denoised V–I trajectory extraction approach by adding it as part of the data preprocessing stage for De Beats [19], Liu [21] and Zhao [49]. Their works use deep learning approaches that utilise V–I trajectories as image inputs into CNNs and Vision Transformer (ViT). De Beats et al. [19] converts V–I trajectories to grey scale pixelated images. These images are then fed to a novel

Table 3
Individual classes AP.

	AC	Bar power	Chiller	Clean power	Coffee machine	Dishwasher
XGB	0.862	0.810	0.922	0.917	0.883	0.920
RF	0.857	0.799	0.909	0.917	0.883	0.912
KNN	0.778	0.617	0.88	0.888	0.878	0.885
DT	0.67	0.612	0.79	0.902	0.755	0.852
MLP	0.665	0.584	0.845	0.846	0.797	0.771
	Distribution board	**Fire alarm**	**Fridge**	**Fryer**	**Glass washer**	**Grease trap**
XGB	0.813	0.983	0.862	0.944	0.903	0.848
RF	0.805	0.981	0.87	0.925	0.9	0.846
KNN	0.703	0.186	0.743	0.916	0.864	0.798
DT	0.617	0.896	0.654	0.904	0.802	0.691
MLP	0.619	0.75	0.64	0.89	0.777	0.633
	Hand dryer	**Hoist**	**Ice machine**	**Indoor lighting**	**Microwave**	**Oven**
XGB	0.946	0.840	0.808	0.910	0.914	0.898
RF	0.93	0.872	0.802	0.917	0.91	0.89
KNN	0.911	0.836	0.639	0.795	0.905	0.856
DT	0.901	0.607	0.571	0.78	0.896	0.76
MLP	0.919	0.576	0.5	0.659	0.865	0.717
	Overdoor heater	**Pasta boiler**	**Pizza oven**	**Post mix coke machine**	**Power**	**Servery general**
XGB	0.880	0.919	0.915	0.952	0.789	0.524
RF	0.875	0.916	0.912	0.909	0.774	0.523
KNN	0.854	0.903	0.894	0.891	0.579	0.233
DT	0.807	0.897	0.858	0.875	0.553	0.236
MLP	0.675	0.868	0.811	0.881	0.375	0.21
	Servery power	**Walk in freezer**	**Walk in fridge**	**Water heater**	**Water pump**	
XGB	0.646	0.946	0.865	0.909	0.918	
RF	0.624	0.907	0.86	0.895	0.91	
KNN	0.324	0.886	0.778	0.875	0.884	
DT	0.388	0.863	0.657	0.784	0.889	
MLP	0.284	0.876	0.744	0.81	0.874	

Table 4
Number of Servery General Misclassifications.

	AC	Bar power	Chiller	Clean power	Dishwasher	Distribution board	Fridge
XGB	1	1	–	–	4	37	6
RF	1	1	–	–	4	25	6
KNN	1	7	–	1	1	47	8
DT	–	7	4	–	5	45	8
MLP	–	3	–	–	1	3	–
	Glass washer	**Hand dryer**	**Hoist**	**Ice machine**	**Indoor lighting**	**Microwave**	**Power**
XGB	–	–	5	5	–	1	2
RF	–	–	3	3	–	1	1
KNN	–	–	1	6	–	1	4
DT	1	1	–	4	4	2	1
MLP	–	–	1	–	–	–	2
	Servery power	**Walk in fridge**	**Water pump**				
XGB	15	–	–				
RF	13	–	–				
KNN	32	–	–				
DT	20	2	–				
MLP	4	–	1				

Table 5
Cross validation results.

	DT	DT (Noise)	KNN	KNN (Noise)	MLP	MLP (Noise)
Acc	0.852 (±0.003)	0.817 (±0.044)	0.852 (±0.003)	0.817 (±0.044)	0.852 (±0.003)	0.817 (±0.044)
Precision	0.798 (±0.01)	0.707 (±0.094)	0.798 (±0.01)	0.707 (±0.094)	0.798 (±0.01)	0.707 (±0.094)
Recall	0.788 (±0.004)	0.685 (±0.065)	0.788 (±0.004)	0.685 (±0.065)	0.788 (±0.004)	0.685 (±0.065)
F1	0.792 (±0.007)	0.684 (±0.09)	0.792 (±0.007)	0.684 (±0.09)	0.792 (±0.007)	0.684 (±0.09)
	RF	**RF (Noise)**	**XGB**	**XGB (Noise)**	**Zhao etal**	**Zhao etal (Noise)**
Acc	0.852 (±0.003)	0.817 (±0.044)	0.901 (±0.002)	0.866 (±0.048)	0.622 (±0.028)	0.606 (±0.024)
Precision	0.798 (±0.01)	0.707 (±0.094)	0.878 (±0.006)	0.791 (±0.097)	0.801 (±0.011)	0.798 (±0.006)
Recall	0.788 (±0.004)	0.685 (±0.065)	0.841 (±0.003)	0.739 (±0.07)	0.472 (±0.045)	0.455 (±0.037)
F1	0.792 (±0.007)	0.684 (±0.09)	0.858 (±0.003)	0.749 (±0.1)	0.418 (±0.052)	0.356 (±0.043)
	Liu etal	**Liu etal (Noise)**	**De Beat etal**	**De Beat etal (Noise)**		
Acc	0.671 (±0.025)	0.665 (±0.028)	0.737 (±0.046)	0.734 (±0.054)		
Precision	0.835 (±0.01)	0.826 (±0.011)	0.815 (±0.024)	0.808 (±0.026)		
Recall	0.529 (±0.039)	0.521 (±0.041)	0.661 (±0.064)	0.662 (±0.082)		
F1	0.604 (±0.053)	0.598 (±0.069)	0.672 (±0.084)	0.67 (±0.105)		

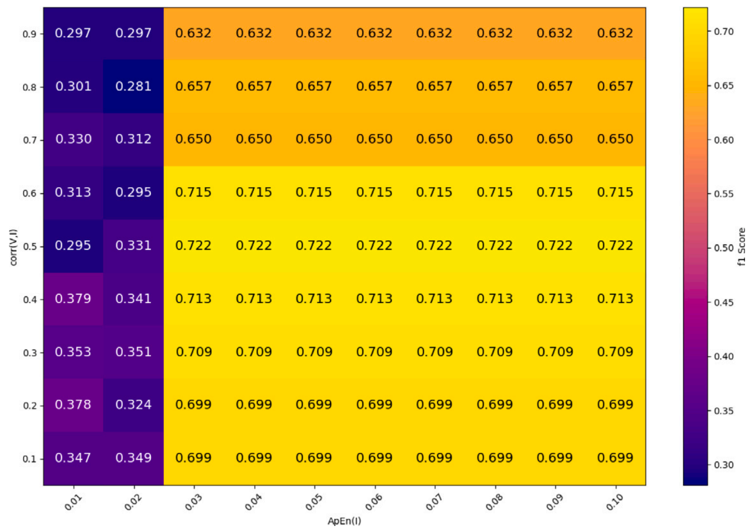

Fig. 8. F1 Score by ApEn(I) and corr(V,I)

Table 6
Performance on Benchmark Datasets.

Dataset	Method	F1	F1 (with Noise)
PLAID2014	RF	0.985 (±0.001)	0.799 (±0.099)
PLAID2017	RF	0.992 (±0.001)	0.813 (±0.229)
PLAID2018	XGB	0.988 (±0.001)	0.792 (±0.171)
COOLL	RF	0.998 (±0.001)	0.943 (±0.042)
WHITED	RF	0.99 (±0.001)	0.937 (±0.045)

CNN architecture to perform classifications [19]. Following a similar technique Zhao et al. [49] and Liu et al. [21] converts V–I trajectories to a coloured HSV image. Leveraging the power of transfer learning, an AlexNet model is then pre-trained [50] on ImageNet and fine tuned using the HSV images as input for appliance classifications. Zhao et al. uses this same HSV images as input into a ViT.

We trained the network of all three approaches on the same training and testing splits as outlined in Section 6.2. Table 5 shows that our trajectory extraction approach improves the F1 score of all three approaches. Performance is not as high compared to the ML models shown in Table 2, and we attribute this to two factors. First, for all three approaches, extensive hyperparameter searches were conducted to select the best performing configuration for their benchmarking datasets. Our experiment only selects the best proposed configurations from their works and applies them directly to our dataset. Second, due to these extensive hyperparameter optimisations, their models were able to achieve high performance with a relative small amount of training samples as compared with normal scenarios. CNNs and ViTs tend to require a much larger amount of data to achieve similar performances [48].

6.8. Comparative analysis with benchmark datasets

To further validate the proposed approach beyond the FM specific dataset used, it was evaluated on five available benchmark datasets researchers have widely used to evaluate NILM based approaches. Specifically, PLAID2014, PLAID2017, PLAID2018, COOLL, and WHITED were used to train and evaluate the ML classifiers selected in Section 6.1. Table 6 reports the F1 scores of top-performing algorithms from cross-validation on datasets with and without noise. The result shows that our approach achieves F1 scores of over 98% for each respective dataset. In comparison recent works on NILM approaches evaluated on these research datasets has achieved comparable mean F1 scores [12,22].

6.9. Number of samples per appliance type

We also performed experiments to determine the number of samples for each appliance type that yield the highest AP. To achieve this, we created new training sets where the number of samples of a specific appliance type is set as a certain percentage of the number of instances of that appliance type in the clean training set. We then trained DT on these new training sets and evaluated their performances on the test set. The AP of each appliance type at different percentages levels is shown in Fig. 9. The result provides two useful insights. First, Fig. 9 shows that the number of samples needed to achieve an AP of over 80% varies between different appliance types. This can be attributed to the V–I trajectory signature variation of some appliance types being drastically different between brands. Second, with exception of *Hand Dryer*, the addition of more samples leads to an increase in AP for all other appliance types. In fact, the number of samples is strongly correlated with the AP in the positive direction ($r = 0.371$).

7. Discussion

This research introduces and demonstrates a new denoised voltage-current feature extraction approach for improving modelling and classification of electronic appliance load signatures as part of a commercially viable NILM system for supporting asset verification in FM. Compared to existing noise reduction and steady-state event detection algorithms, the novel application of ApEn is able to provide a statistically effective and computational efficient method detecting load instability and removal from steady-state appliance signal waveforms to improve the accuracy of feature extraction for load classification. The proposed approach was evaluated using both an FM appliance load dataset comprising of 29 assets operating in real world commercial properties as well as five popular benchmark research datasets. Our approach was shown to consistently improve asset classification performances when applied across both traditional ML classifiers and a number of state-of-the-art DL approaches.

The FM appliance data used consisted of assets representing a variety of consumer and commercial appliance types beyond those found in existing available research datasets. We performed further analysis and comparisons of the dataset with three other FM datasets acquired using the ALM units described in section 3.1. These datasets were acquired from other client sites comprising of hireable commercial office spaces and multi site 12 and 24 h amusement/entertainment establishments. Across all these datasets we found that the recorded assets were reflective of similar types, makes and models found across similar

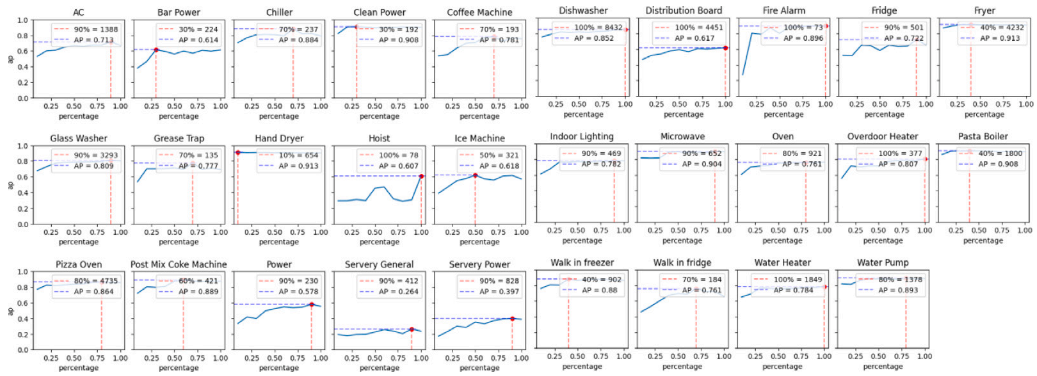

Fig. 9. AP of each appliance type at different number of samples.

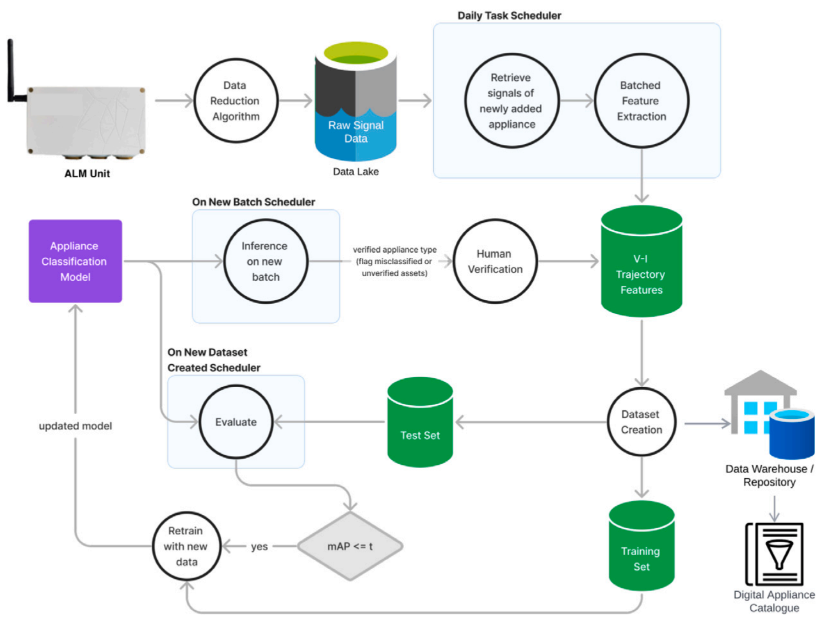

Fig. 10. Deployment mechanism for V–I trajectory feature extraction pipeline.

commercial sites for which our approach was achieving comparable performances with accuracies and F1 scores of between 88% and 97%. This shows the relevance of the commercial dataset selected to validate our approach and its potential generalisability when deployed.

7.1. FM deployment scenario

In this section, we outline a mechanism for deploying our proposed pipeline and for model training and evaluation as part of the digital FM platform shown in Fig. 3. This approach aligns with broader efforts in smart building energy management systems that integrate machine learning and IoT for enhanced energy efficiency and automation in various environments, including those with resource-constrained devices for data collection and control [10]. Given that the company's infrastructure and products are built on top of Microsoft Azure, we recommend the use of Microsoft Azure Machine Learning (Azure ML) for the deployment [51]. Fig. 10 shows the flowchart of the mechanism.

We suggest the feature extraction be performed in batches where each batch contains all the new appliance data collected through the ALM units over a single day. These might be newly installed ALM units or the reconfiguring of existing units based on connecting new application channels for monitoring. This can allow the processing of multiple appliances to be done in parallel. The raw data collected from the ALM units would be stored and maintained in a data lake for pre-processing. Once the V–I trajectory features have been extracted for the batch, inferences for appliance classification can then be performed on the extracted batched features. The predicted labels of each appliance in the batch can then be used to flag any incorrectly classified appliances or those that cannot be verified with a high enough confidence. These flagged appliances can then be checked by the hardware and data analytics teams.

To continually assess the performance of the trained classification model, an evaluation mechanism is needed. The creation of a new training and test set can be done automatically from the extracted V–I trajectory features database. Once the new dataset is created, the current model can be evaluated on the test set. If its *mAP* drops below a threshold *t*, then a new model is trained on the updated (verified) training set. This training and evaluation mechanism minimises the cost of model retraining while the model is operating with acceptable performance levels. Future work will investigate more efficient model fine tuning approaches such as semi-supervised learning techniques. These algorithms learn the distribution characteristics of data by simultaneously using smaller amounts of labelled data with larger available unlabelled data where labelled data points are scarce or expensive to obtain [25,52].

Data collected on the newly verified assets would be stored where the labelled V–I features in combination with the other signal (see Section 3.1) and appliance information can be used to distinguish and catalogue different electronic appliance types. FM clients can use this digital catalogue to compare and contrast Return on Investment (ROI), age/lifespan, power factor, and other energy related metrics between similar appliance makes and models in order to make informed purchasing or upgrade decisions considering, usage requirements, operational, maintenance costs and net zero sustainability targets. The collected asset data can also be integrated with smarter digital interventions for building infrastructures and their underpinning data systems. Computational simulation models can be used to visualise and comprehend the complex interplay between environmental variables and performance while using Large Language Models to interface with and explore real-time response strategies to different FM operational and energy consumption scenarios. Here, data and ML driven digital twins can model building infrastructures and connected assets to help simulate these scenarios and optimise performance against client requirements and constraints for improving operational and maintenance efficiencies while contributing to reducing CO2 emissions [53,54].

7.2. Impact on ALM installation process for FM

The proposed asset classification pipeline can have a profound impact on the existing ALM installation process at Cloudfm Group Ltd . For a typical client site with 100 assets, the current installation process can take between 2–3 h. This can be longer for larger sites and does not include manual checking of misclassified or unverifiable assets requiring return visits to client sites for implementing corrections. This can extend the installation and verification process over several days. The trained asset classification model can significantly reduce the time needed for these processes in two ways. Firstly, immediately after installation of an ALM unit and confirmation it is transmitting signals, the model can start determining the appliance type after the first steady-state event has been detected which varies between appliances starting from 10 min. In addition, each model's inference after determining a steady-state event takes approximately 300 ms. This means that for some appliances, the model can perform classifications while the engineer connects other appliances to the ALM unit or installs other ALM units on-site. In this scenario, engineers can confirm the predictions with low confidence scores without having to revisit the site. Secondly, for appliances whose steady-state takes a longer time to determine, the engineer can perform manual tagging on-site and the model's predictions can be cross-checked with the manual tags at a later time against any conflicting tags highlighted by the data science team. This reduces the number of appliances that need to be checked by the team.

The timely and correct installation of multiple ALMs units at client sites can start to have a positive impact on their operational efficiency and energy consumption. Data from the ALMs can be used to determine the digital fingerprint of its operation which can be used to track performance and raise an alarm if the asset starts to perform sub optimally to recommend a maintenance check by also comparing against similar asset types across several sites. Additionally the data acquired from ALM units can be used to identify operating inefficiencies and drive behaviour change in the interaction with operating staff. This can have a signifiant impact on reducing CO2 emissions of clients helping towards achieving the UK's Net-zero targets.

8. Conclusion

This paper presents a new approach for ensuring V–I Trajectories are extracted correctly from non-uniformly sampled high volume commercial appliance data. The proposed approach applies approximate entropy to identify and remove unstructured current and voltage signals. These could otherwise be misclassified as steady-state events leading to poorly extracted V–I Trajectory features for identifying asset load signatures. ApEn provides and efficient and scalable approach for measuring and identifying noise misclassified as steady-state events not previously applied for the processing of commercial operating appliance load signals for load classification and verification solutions.

Our denoised feature extraction approach is shown to significantly improve the classification of asset types when applied as a preprocessing step in training ML pipelines on extensive real world data reflecting the variety of consumer and commercially operating appliances and units found across industry sectors covered by FM. The validity of the approach is further proven by evaluating its performance gain on multiple ML classification models, widely used publicly accessible datasets and comparison with state-of-the-art ML techniques as part of other NILM solutions.

The proposed approach can be used to benefit the FM sector in the automatic identification or verification of connected appliances across multi-site properties and buildings to improve installation of ALM units. The timely and correct installation of ALM units will directly contribute to monitoring the operating states of assets enabling preventive maintenance improving economic returns to clients while also improving their operating efficiencies impacting the environment through reduced energy consumption.

In addition to improving operational accuracy and cost efficiency, enhanced appliance classification accuracy contributes to broader sustainability outcomes. Reliable identification and monitoring of assets enable more informed energy management decisions, reducing unnecessary energy use and associated emissions. Consistent with recent advances in sustainable machine learning applications [55–57], the proposed approach supports the long-term objectives of improved energy efficiency and reduced environmental impact.

This study was limited in its exploration of other relevant shape features as input to ML and DL algorithms. Future research directions will include building on the current feature extraction and classification pipeline to include more shape features as well as process and combine other recorded signal parameters such as voltage harmonics. Furthermore, future work will also focus on quantifying how often valid but irregular steady-state events occur and assessing their potential impact on appliances with inherently variable load behaviours such as batteries. Converting the extracted data into nodes and edges of a graph and applying Graph Neural Networks (GNN) for appliance classification will also be explored.

CRediT authorship contribution statement

Socretquuliqaa Lee: Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Conceptualization. **Faiyaz Doctor:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization. **Mohammad Hossein Anisi:** Writing – review & editing, Supervision, Investigation. **Shashank Goud:** Writing – review & editing, Resources, Methodology, Data curation, Conceptualization. **Xiao Wang:** Writing – review & editing, Resources, Data curation, Conceptualization.

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.

Acknowledgements

This work was supported by the University of Essex QR Impact Fund and in collaboration with Cloudfm Group Limited. For the purposes of open access, the authors have applied a Creative Commons Attribution (CC BY) License to any Author Accepted Manuscript (AAM) version arising from this submission.

Data availability

The authors have made the full feature extraction workflow openly available to support replication and further research at: https://github.com/Socret360/denoised-vi-pipeline.

References

[1] D.L. Capuano, International energy outlook 2018 (IEO2018), 2018.

[2] Y. Liu, W. Liu, Y. Shen, X. Zhao, S. Gao, Toward smart energy user: Real time non-intrusive load monitoring with simultaneous switching operations, Appl. Energy 287 (2021) 116616, http://dx.doi.org/10.1016/j.apenergy.2021.116616, URL https://linkinghub.elsevier.com/retrieve/pii/S0306261921001537.

[3] P.A. Schirmer, I. Mporas, Non-intrusive load monitoring: A review, IEEE Trans. Smart Grid 14 (1) (2023) 769–784, http://dx.doi.org/10.1109/TSG.2022.3189598, URL https://ieeexplore.ieee.org/document/9820770/.

[4] S. Dash, N. Sahoo, Electric energy disaggregation via non-intrusive load monitoring: A state-of-the-art systematic review, Electr. Power Syst. Res. 213 (2022) 108673, http://dx.doi.org/10.1016/j.epsr.2022.108673, URL https://linkinghub.elsevier.com/retrieve/pii/S0378779622007398.

[5] K. Carrie Armel, A. Gupta, G. Shrimali, A. Albert, Is disaggregation the holy grail of energy efficiency? The case of electricity, Energy Policy 52 (2013) 213–234, http://dx.doi.org/10.1016/j.enpol.2012.08.062, URL https://linkinghub.elsevier.com/retrieve/pii/S0301421512007446.

[6] A. Zoha, A. Gluhak, M. Imran, S. Rajasegarar, Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey, Sensors 12 (12) (2012) 16838–16866, http://dx.doi.org/10.3390/s121216838, URL http://www.mdpi.com/1424-8220/12/12/16838.

[7] J. Lu, R. Zhao, B. Liu, Z. Yu, J. Zhang, Z. Xu, An overview of non-intrusive load monitoring based on V-I trajectory signature, Energies 16 (2) (2023) 939, http://dx.doi.org/10.3390/en16020939, URL https://www.mdpi.com/1996-1073/16/2/939.

[8] N.A.-H.K. Hussein, S. Khalaf, K.D. Shah, Advanced composite materials for sustainable construction: Innovations in civil engineering applications, KHWARIZMIA 2024 (2024) 8–16, http://dx.doi.org/10.70470/KHWARIZMIA/2024/003, URL https://peninsula-press.ae/Journals/index.php/KHWARIZMIA/article/view/88.

[9] R. Yan, R.X. Gao, Approximate entropy as a diagnostic tool for machine health monitoring, Mech. Syst. Signal Process. 21 (2) (2007) 824–839, http://dx.doi.org/10.1016/j.ymssp.2006.02.009, URL https://linkinghub.elsevier.com/retrieve/pii/S0888327006000458.

[10] M. Sheela, S. Gopalakrishnan, I. Begum, J.J. Hephzipah, M. Gopianand, D. Harika, Enhancing energy efficiency with smart building energy management system using machine learning and IOT, Babylon. J. Mach. Learn. 2024 (2024) 80–88, http://dx.doi.org/10.58496/BJML/2024/008, URL https://mesopotamian.press/journals/index.php/BJML/article/view/423.

[11] M.C. Ghanem, S. Salloum, Integrating AI-driven deep learning for energy-efficient smart buildings in internet of thing-based industry 4.0, Babylon. J. Internet Things 2025 (2025) 121–130, http://dx.doi.org/10.58496/BJIoT/2025/007, URL https://mesopotamian.press/journals/index.php/BJIoT/article/view/804.

[12] B.M. Mulinari, L.D.S. Nolasco, E. Oroski, A.E. Lazzaretti, R.R. Linhares, D.P.B. Renaux, Feature extraction of V–I trajectory using 2-D Fourier series for electrical load classification, IEEE Sensors J. 22 (18) (2022) 17988–17996, http://dx.doi.org/10.1109/JSEN.2022.3194999, URL https://ieeexplore.ieee.org/document/9850433/.

[13] T. Hassan, F. Javed, N. Arshad, An empirical investigation of V-I trajectory based load signatures for non-intrusive load monitoring, IEEE Trans. Smart Grid 5 (2) (2014) 870–878, http://dx.doi.org/10.1109/TSG.2013.2271282, URL http://ieeexplore.ieee.org/document/6575197/.

[14] B.M. Mulinari, R.R. Linhares, D.P. De Campos, C.H. Da Costa, H.C. Ancelmo, A.E. Lazzaretti, E. Oroski, C.R.E. Lima, D.P.B. Renaux, F. Pottker, A new set of steady-state and transient features for power signature analysis based on V-I trajectory, in: 2019 IEEE PES Innovative Smart Grid Technologies Conference - Latin America (ISGT Latin America), IEEE, Gramado, Brazil, 2019, pp. 1–6, http://dx.doi.org/10.1109/ISGT-LA.2019.8895360, URL https://ieeexplore.ieee.org/document/8895360/.

[15] J. Gao, E.C. Kara, S. Giri, M. Berges, A feasibility study of automated plug-load identification from high-frequency measurements, in: 2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP), IEEE, Orlando, FL, USA, 2015, pp. 220–224, http://dx.doi.org/10.1109/GlobalSIP.2015.7418189, URL http://ieeexplore.ieee.org/document/7418189/.

[16] N. Iksan, J. Sembiring, N. Haryanto, S.H. Supangkat, Appliances identification method of non-intrusive load monitoring based on load signature of V-I trajectory, in: 2015 International Conference on Information Technology Systems and Innovation (ICITSI), IEEE, Bandung, Bali, Indonesia, 2015, pp. 1–6, http://dx.doi.org/10.1109/ICITSI.2015.7437744, URL http://ieeexplore.ieee.org/document/7437744/.

[17] L. De Baets, C. Develder, T. Dhaene, D. Deschrijver, Automated classification of appliances using elliptical fourier descriptors, in: 2017 IEEE International Conference on Smart Grid Communications (SmartGridComm), IEEE, Dresden, Germany, 2017, pp. 153–158, http://dx.doi.org/10.1109/SmartGridComm.2017.8340669, URL http://ieeexplore.ieee.org/document/8340669/.

[18] L. De Baets, T. Dhaene, D. Deschrijver, C. Develder, M. Berges, VI-based appliance classification using aggregated power consumption data, in: 2018 IEEE International Conference on Smart Computing (SMARTCOMP), IEEE, Taormina, 2018, pp. 179–186, http://dx.doi.org/10.1109/SMARTCOMP.2018.00089, URL https://ieeexplore.ieee.org/document/8421347/.

[19] L. De Baets, J. Ruyssinck, C. Develder, T. Dhaene, D. Deschrijver, Appliance classification using VI trajectories and convolutional neural networks, Energy Build. 158 (2018) 32–36, http://dx.doi.org/10.1016/j.enbuild.2017.09.087, URL https://linkinghub.elsevier.com/retrieve/pii/S0378778817312690.

[20] D. Baptista, S. Mostafa, L. Pereira, L. Sousa, F. Morgado-Dias, Implementation strategy of convolution neural networks on field programmable gate arrays for appliance classification using the voltage and current (V-I) trajectory, Energies 11 (9) (2018) 2460, http://dx.doi.org/10.3390/en11092460, URL http://www.mdpi.com/1996-1073/11/9/2460.

[21] Y. Liu, X. Wang, W. You, Non-intrusive load monitoring by voltage–current trajectory enabled transfer learning, IEEE Trans. Smart Grid 10 (5) (2019) 5609–5619, http://dx.doi.org/10.1109/TSG.2018.2888581, URL https://ieeexplore.ieee.org/document/8580416/.

[22] T. Chen, Y. Yuan, J. Gao, S. Guo, P. Yang, Non-intrusive load monitoring based on time-enhanced multidimensional feature visualization, Sci. Rep. 15 (1) (2025) 4800, http://dx.doi.org/10.1038/s41598-025-89191-x, URL https://www.nature.com/articles/s41598-025-89191-x.

[23] H. Lam, G. Fung, W. Lee, A novel method to construct taxonomy electrical appliances based on load signaturesof, IEEE Trans. Consum. Electron. 53 (2) (2007) 653–660, http://dx.doi.org/10.1109/TCE.2007.381742, URL http://ieeexplore.ieee.org/document/4266955/.

[24] L. De Baets, C. Develder, T. Dhaene, D. Deschrijver, Detection of unidentified appliances in non-intrusive load monitoring using siamese neural networks, Int. J. Electr. Power Energy Syst. 104 (2019) 645–653, http://dx.doi.org/10.1016/j.ijepes.2018.07.026, URL https://linkinghub.elsevier.com/retrieve/pii/S0142061517331253.

[25] T. Tang, K. Li, C. Su, Z. Liu, Semi-supervised learning with flexible threshold for non-intrusive load monitoring, Heliyon 10 (14) (2024) e34457, http://dx.doi.org/10.1016/j.heliyon.2024.e34457, URL https://linkinghub.elsevier.com/retrieve/pii/S2405844024104884.

[26] J. Kolter, M. Johnson, REDD: A public data set for energy disaggregation research, Artificial Intelligence 25 (2011).

[27] J. Gao, S. Giri, E.C. Kara, M. Bergés, PLAID: a public dataset of high-resoultion electrical appliance measurements for load identification research: demo abstract, in: Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings, ACM, Memphis Tennessee, 2014, pp. 198–199, http://dx.doi.org/10.1145/2674061.2675032, URL https://dl.acm.org/doi/10.1145/2674061.2675032.

[28] L.D. Baets, M. Berges, T. Dhaene, C. Develder, Jingkun Gao, D. Deschrijver, PLAID 2017, figshare, 2020, http://dx.doi.org/10.6084/M9.FIGSHARE.11605215.V1.

[29] M. Kahl, A. Haq, T. Kriechbaumer, H.-a. Jacobsen, WHITED - a worldwide household and industry transient energy data set, 2016.

[30] M.N. Meziane, T. Picon, P. Ravier, G. Lamarque, J.-C. Le Bunetel, Y. Raingeaud, A measurement system for creating datasets of on/off-controlled electrical loads, in: 2016 IEEE 16th International Conference on Environment and Electrical Engineering (EEEIC), IEEE, Florence, Italy, 2016, pp. 1–5, http://dx.doi.org/10.1109/EEEIC.2016.7555847, URL http://ieeexplore.ieee.org/document/7555847/.

[31] T. Kriechbaumer, H.-A. Jacobsen, BLOND, a building-level office environment dataset of typical electrical appliances, Sci. Data 5 (1) (2018-03-27) 180048, http://dx.doi.org/10.1038/sdata.2018.48, URL https://www.nature.com/articles/sdata201848.

[32] C.I. Salis, A.E. Malissovas, P.A. Bizopoulos, A.T. Tzallas, P.A. Angelidis, D.G. Tsalikakis, Denoising simulated EEG signals: A comparative study of EMD, wavelet transform and Kalman filter, in: 13th IEEE International Conference on BioInformatics and BioEngineering, IEEE, Chania, Greece, 2013, pp. 1–4, http://dx.doi.org/10.1109/BIBE.2013.6701613, URL http://ieeexplore.ieee.org/document/6701613/.

[33] K.-M. Chang, Arrhythmia ECG noise reduction by ensemble empirical mode decomposition, Sensors 10 (6) (2010) 6063–6080, http://dx.doi.org/10.3390/s100606063, URL https://www.mdpi.com/1424-8220/10/6/6063.

[34] Y. Jia, H. Pei, J. Liang, Y. Zhou, Y. Yang, Y. Cui, M. Xiang, Preprocessing and denoising techniques for electrocardiography and magnetocardiography: A review, Bioengineering 11 (11) (2024) 1109, http://dx.doi.org/10.3390/bioengineering11111109, URL https://www.mdpi.com/2306-5354/11/11/1109.

[35] S. Chatterjee, R.S. Thakur, R.N. Yadav, L. Gupta, D.K. Raghuvanshi, Review of noise removal techniques in ECG signals, IET Signal Process. 14 (9) (2020) 569–590, http://dx.doi.org/10.1049/iet-spr.2020.0104, URL https://onlinelibrary.wiley.com/doi/10.1049/iet-spr.2020.0104.

[36] Y.X. Su, G.P. Liu, X.Q. Shen, EEMD-ApEn applied in power quality detection and classification, Appl. Mech. Mater. 423–426 (2013) 2443–2447, http://dx.doi.org/10.4028/www.scientific.net/AMM.423-426.2443, URL https://www.scientific.net/AMM.423-426.2443.

[37] S. Aguayo-Tapia, G. Avalos-Almazan, J.D.J. Rangel-Magdaleno, Entropy-based methods for motor fault detection: A review, Entropy 26 (4) (2024) 299, http://dx.doi.org/10.3390/e26040299, URL https://www.mdpi.com/1099-4300/26/4/299.

[38] A. Meier, D. Cautley, Practical limits to the use of non-intrusive load monitoring in commercial buildings, Energy Build. 251 (2021) 111308, http://dx.doi.org/10.1016/j.enbuild.2021.111308, URL https://linkinghub.elsevier.com/retrieve/pii/S0378778821005922.

[39] M. Zhou, S. Shao, X. Wang, Z. Zhu, F. Hu, Deep learning-based non-intrusive commercial load monitoring, Sensors 22 (14) (2022) 5250, http://dx.doi.org/10.3390/s22145250, URL https://www.mdpi.com/1424-8220/22/14/5250.

[40] N. Batra, O. Parson, M. Berges, A. Singh, A. Rogers, A comparison of non-intrusive load monitoring methods for commercial and residential buildings, 2014, arXiv:1408.6595 [cs]. URL http://arxiv.org/abs/1408.6595.

[41] L. Pereira, V. Aguiar, F. Vasconcelos, R. Martins, T. Garcês, H. Morais, Fikelectricity: A electricity consumption dataset from three restaurant kitchens in Portugal, Sci. Data 10 (1) (2023) 779, http://dx.doi.org/10.1038/s41597-023-02698-8, URL https://www.nature.com/articles/s41597-023-02698-8.

[42] S. Lee, F. Doctor, M.H. Anisi, S. Goud, X. Wang, S. Ruthven, AI driven streamlining of appliance load monitoring in facilities management, in: 2024 19th Annual System of Systems Engineering Conference (SoSE), IEEE, Tacoma, WA, USA, 2024, pp. 130–133, http://dx.doi.org/10.1109/SOSE62659.2024.10620932, URL https://ieeexplore.ieee.org/document/10620932/.

[43] S. Baskar, UK Facilities Management Market Report 2024 – 2027. URL https://baachurain.com/register/uk-facilities-management-market-2024-2027-report/.

[44] A.L. Wang, B.X. Chen, C.G. Wang, D. Hua, Non-intrusive load monitoring algorithm based on features of V–I trajectory, Electr. Power Syst. Res. 157 (2018) 134–144, http://dx.doi.org/10.1016/j.epsr.2017.12.012, URL https://linkinghub.elsevier.com/retrieve/pii/S0378779617304911.

[45] A. Delgado-Bonal, A. Marshak, Approximate entropy and sample entropy: A comprehensive tutorial, Entropy 21 (6) (2019) 541, http://dx.doi.org/10.3390/e21060541, URL https://www.mdpi.com/1099-4300/21/6/541.

[46] S.M. Pincus, Approximate entropy as a measure of system complexity., Proc. Natl. Acad. Sci. 88 (6) (1991) 2297–2301, http://dx.doi.org/10.1073/pnas.88.6.2297, URL https://pnas.org/doi/full/10.1073/pnas.88.6.2297.

[47] J. Bergstra, R. Bardenet, Y. Bengio, B. Kégl, Algorithms for hyper-parameter optimization, in: J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, K. Weinberger (Eds.), in: Advances in Neural Information Processing Systems, vol. 24, Curran Associates, Inc., 2011, URL https://proceedings.neurips.cc/paper_files/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf.

[48] S. Takahashi, Y. Sakaguchi, N. Kouno, K. Takasawa, K. Ishizu, Y. Akagi, R. Aoyama, N. Teraya, A. Bolatkan, N. Shinkai, H. Machino, K. Kobayashi, K. Asada, M. Komatsu, S. Kaneko, M. Sugiyama, R. Hamamoto, Comparison of vision transformers and convolutional neural networks in medical image analysis: A systematic review, J. Med. Syst. 48 (1) (2024) 84, http://dx.doi.org/10.1007/s10916-024-02105-8, URL https://link.springer.com/10.1007/s10916-024-02105-8.

[49] Q. Zhao, W. Liu, K. Li, Y. Wei, Y. Han, Unknown appliances detection for non-intrusive load monitoring based on vision transformer with an additional detection head, Heliyon 10 (9) (2024) e30666, http://dx.doi.org/10.1016/j.heliyon.2024.e30666, URL https://linkinghub.elsevier.com/retrieve/pii/S2405844024066970.

[50] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: F. Pereira, C. Burges, L. Bottou, K. Weinberger (Eds.), in: Advances in Neural Information Processing Systems, vol. 25, Curran Associates, Inc., 2012, URL https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.

[51] Azure Machine Learning - ML as a Service | Microsoft Azure. URL https://azure.microsoft.com/en-us/products/machine-learning.

[52] F. Zhao, J. Zhou, Y. Ren, H. Liu, Y. Wang, W. Xia, Mean teacher for semi-supervised non-intrusive load monitoring based on V-I trajectory, in: 2023 8th Asia Conference on Power and Electrical Engineering (ACPEE), IEEE, Tianjin, China, 2023, pp. 1623–1627, http://dx.doi.org/10.1109/ACPEE56931.2023.10135645, URL https://ieeexplore.ieee.org/document/10135645/.

[53] F. Khalifa, M. Marzouk, Integrated blockchain and digital twin framework for sustainable building energy management, J. Ind. Inf. Integr. 43 (2025) 100747, http://dx.doi.org/10.1016/j.jii.2024.100747, URL https://linkinghub.elsevier.com/retrieve/pii/S2452414X24001900.

[54] I. Petri, Y. Rezgui, A. Ghoroghi, A. Alzahrani, Digital twins for performance management in the built environment, J. Ind. Inf. Integr. 33 (2023) 100445, http://dx.doi.org/10.1016/j.jii.2023.100445, URL https://linkinghub.elsevier.com/retrieve/pii/S2452414X23000183.

[55] S. Tao, H. Liu, C. Sun, H. Ji, G. Ji, Z. Han, R. Gao, J. Ma, R. Ma, Y. Chen, S. Fu, Y. Wang, Y. Sun, Y. Rong, X. Zhang, G. Zhou, H. Sun, Collaborative and privacy-preserving retired battery sorting for profitable direct recycling via federated machine learning, Nat. Commun. 14 (1) (2023) 8032, http://dx.doi.org/10.1038/s41467-023-43883-y, URL https://www.nature.com/articles/s41467-023-43883-y.

[56] S. Tao, R. Ma, Z. Zhao, G. Ma, L. Su, H. Chang, Y. Chen, H. Liu, Z. Liang, T. Cao, H. Ji, Z. Han, M. Lu, H. Yang, Z. Wen, J. Yao, R. Yu, G. Wei, Y. Li, X. Zhang, T. Xu, G. Zhou, Generative learning assisted state-of-health estimation for sustainable battery recycling with random retirement conditions, Nat. Commun. 15 (1) (2024) 10154, http://dx.doi.org/10.1038/s41467-024-54454-0, URL https://www.nature.com/articles/s41467-024-54454-0.

[57] Sustainable Development Goal 7: Affordable and Clean Energy | United Nations in China. URL https://china.un.org/en/sdgs/7.