Computers and Electrical Engineering 123 (2025) 110163

Contents lists available at ScienceDirect

Computers and Electrical Engineering

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

Unveiling energy usage patterns in industrial kitchens: From detection to clustering of appliance usage

Ricardo Martins [a], Hugo Morais [a,b], Lucas Pereira [a,c,*]

[a] Instituto Superior Técnico - IST, Universidade de Lisboa, Lisboa, 1049-001, Portugal
[b] INESC-ID – Instituto de Engenharia de Sistemas e Computadores-nvestigacão e, Desenvolvimento, 1049-001, Portugal
[c] Interactive Technologies Institute, LARSyS, Rua da Manutenção 71, Building F, S05, Lisboa, 1900-500, Portugal

ARTICLE INFO

Dataset link: https://doi.org/10.17605/OSF.IO/K3G8N, https://doi.org/10.17605/OSF.IO/SPWD9

Keywords:
Consumption patterns
Data-driven analysis
Energy efficiency
Industrial kitchen devices
Smart grid
Time series clustering

ABSTRACT

Industrial Kitchens (IKs) are characterized by high energy consumption, yet they remain largely overlooked in energy research. Understanding how electricity is used in IKs is crucial for identifying opportunities for energy optimization and improving sustainability in this sector. This paper presents a data-driven methodology for analyzing appliance consumption by automatically detecting and classifying appliance activations. The approach combines automatic activity detection with unsupervised clustering to reveal usage patterns. Evaluated on data from nine IK appliances, the methodology achieves outstanding performance, with average balanced accuracy and F1-scores exceeding 0.98. The unsupervised classification identifies distinct cycle modes for each appliance, with the optimal number of clusters varying across appliances. Load fluctuation patterns are found to be the most significant feature, with appliances like the ice machine exhibiting unique consumption behaviors compared to similar appliances like refrigerators. In contrast, appliances such as the salamander draw power consistently, regardless of activity duration. These findings not only contribute to a better understanding of energy use in IKs but also lay the groundwork for future research on demand response strategies and energy efficiency improvements in small-scale commercial kitchens.

1. Introduction

The kitchen is the visible cultural manifestation of the technology human beings employ to store, prepare, and eat food [1]. The transition from the family kitchen and the breakthrough of the Industrial Kitchens (IKs), widely known as restaurants, dates back to the eighteenth century in Paris or the thirteenth century in China [2]. With the popularization of such establishments, restaurants recently emerged as one of the most energy-intensive players in the commercial sector, using about five to seven times more energy per square meter than other commercial buildings [3].

Furthermore, it is important to emphasize that energy demand is poised to surge in the foreseeable future due to the transformative shifts occurring within the restaurant industry, driven by the burgeoning food-delivery ecosystem where orders placed using platforms represent additional tables in the dining area. According to [4], food delivery is already a global market worth more than $150 billion, having tripled since 2017. Ultimately, considering the projected expansion of the food and beverage services market to reach $4651.03 billion by 2027, with a Compound Annual Growth Rate (CAGR) of 5.4% [5], a compelling opportunity emerges to enhance the operational efficiency of IKs within food and beverage establishments.

* Corresponding author at: Instituto Superior Técnico - IST, Universidade de Lisboa, Lisboa, 1049-001, Portugal.
E-mail address: lucas.pereira@tecnico.ulisboa.pt (L. Pereira).

https://doi.org/10.1016/j.compeleceng.2025.110163
Received 22 May 2024; Received in revised form 10 December 2024; Accepted 4 February 2025
Available online 19 February 2025

This enhancement can be achieved by implementing strategies that promote better coordination between using IKs' appliances and integrating small production from renewable energy sources and storage technologies or even promoting carbon offsetting [6] and the participation of restaurants in electricity markets [7]. Such improvements in electricity consumption will not only help to reduce the carbon footprint but will also strongly contribute to the energy transition targets and reduce the waste of electricity and water [8,9], and in the end, reduce the costs to the owners of the restaurants [10,11].

Understanding the energy consumption profiles of appliances is, therefore, a critical task, as this will allow the development of prescriptive models to provide forecasts, feedback, and recommendations, promoting sustainable behaviors and opening room to the study of the potential participation of IKs in demand-response programs [12].

Despite the high preponderance and consumption in IKs, very little research has been carried out so far to understand individual appliance consumption. To the best of the author's knowledge, the few published works focus on forecasting the demand for restaurant appliances considering hour [13] and day-ahead scenarios [13,14].

Instead, most studies on IKs focus on benchmarking, aiming at finding proper normalization factors [15–17]. Whereas other studies focus on understanding how distributed energy resources such as Solar PhotoVoltaic (PV) and Battery Energy Storage System (BESS) can be incorporated into the IK operation [18–20]. While these studies are important because they define common metrics to compare different establishments and contribute to better energy management in the sector, they do not rely on analyzing the behavior of individual appliances.

In this context, the original contributions of this work to ongoing the body of knowledge are twofold:

1. The proposal of a data-driven methodology to examine and unveil patterns in the electricity consumption of industrial kitchen appliances. The methodology directly classifies the appliance mode of operation, leveraging an expert heuristic to automatically detect the working cycles from the consumption time series and the combination of a data-driven feature extraction method, leveraging the *Catch22* feature set [21], and k-means clustering for unsupervised classification.
   Focusing on individual device cycles rather than aggregate consumption is relevant since it facilitates a detailed comprehension of the appliance modes of operation that would not be available otherwise. Additionally, the data-driven approach for feature extraction liberates users from engineering their own features, which would require a prior understanding of the appliance's working principles.
2. An in-depth analysis of electricity consumption of nine devices monitored from one IK during its quotidian operation. Besides illustrating the effectiveness and usefulness of the proposed methodology, this analysis has revealed some unexpected behaviors from certain appliances that share the same working principles and unveiled energy efficiency opportunities for others. In both cases, such conclusions are only possible due to the possibility of scrutinizing the operation cycles of the appliances at stake. This is, to the best of the authors' knowledge, the first work that attempts to profile IK appliances.

This paper is organized into five main sections. An overview of existing works in data-driven analysis of appliance electric demand is provided in Section 2. The proposed methods are thoroughly described in Section 3, and a case study designed to illustrate the application of the proposed methodology to real IK devices is specified in Section 4. The results of the case study are then comprehensively presented and discussed in Section 5. Finally, the main conclusions, limitations, and future work directions are presented in Section 6.

2. Related works

In the residential sector, many authors have addressed and continue to address the topic of excessive energy consumption by developing several methods and techniques that leverage smart meter data. This includes numerous works in areas such as forecasting, clustering, classification, and optimization to enable potential applications of smart meter data analytics, including user feedback, anomaly detection, and participation in demand-response programs [22,23]. In this regard, most related works focus on applying clustering methods to appliance-level consumption data. For example, in [24–27], clustering techniques were used to group the consumption traces of different appliances. These clusters were later used as inputs to the development of forecasting or classification models for residential appliances. Likewise, several authors have employed clustering techniques to categorize the consumption profiles of several appliances [28–31]. The majority of these works aimed at assessing the Demand-Response (DR) potential of household appliances.

On the other hand, a few authors have explored correlation-based approaches to gain insights into application consumption data. For example, in [32], the authors proposed a methodology based on auto-correlation and probability distributions to model the intensity of usage of residential appliances. Appliance intensity of usage and correlation with total demand was also explored in [33] to assess how appliance behavior patterns can help support and optimize DR programs. Finally, in [34], the authors employed Sequentially Distributing Auto Regression to detect changes or outliers in the trend of individual appliance profiles.

A brief summary of the reviewed papers is provided in Table 1. It can be observed that K-means clustering is the most commonly employed technique in related works. However, these studies, except for [31], typically use the entire time series as input for clustering. Our work takes a different approach in two key aspects: first, we focus on appliance activation cycles, merging relevant periods to streamline analysis, improve computational efficiency, and enhance interpretability. Second, unlike [31], which uses a fixed threshold (10% of the peak demand) to identify activation cycles, we propose an adaptive method that dynamically adjusts to the specific consumption profile of each appliance.

Additionally, apart from [30], most works rely solely on power as the input feature. While our work also transforms the input signal, like [30], it avoids black-box models such as Auto-Encoders. Instead, we utilize a carefully curated set of features selected from extensive literature designed for general utility. Finally, apart from two existing works on forecasting IK appliance demand, this study is, to the best of our knowledge, the first to comprehensively explore the usage patterns of IK appliances.

Table 1
Brief summary of the reviewed papers.

Ref.	Summary
[24]	The k-medoids clustering algorithm is employed to find representative appliance consumption groups. These groups are used as inputs for forecasting peak demand. Several appliances were studied, including a refrigerator, an oven, a hair dryer, an air conditioner (AC), and a water heater.
[25]	The Clustering Large Application (CLARA) algorithm was used to identify consumption clusters in household appliances, including ACs, furnaces, refrigerators, microwaves, and clothes dryers. The clusters were used as inputs to forecast appliance demand in several periods (e.g., daily and weekly).
[26]	K-means was used to identify the different working states of household appliances. These states are then used as inputs to conditional hidden semi-Markov Models (cHMM), which are later used to develop short-term demand forecasting algorithms for each individual appliance. Five types of residential appliances were considered: ACs, Refrigerators, pool pumps, EVs, and water heaters.
[27]	The k-means++ is used to cluster consumption profiles into different appliance categories. These clusters were then used as inputs to a Graph Convolutional Neural Network (GCN) to classify the appliance type information. The proposed method was tested with five types of loads: microwave, refrigerator, TV, laptop, and water heater.
[28]	K-means clustering is performed to categorize AC unit usage patterns. The clustered load profiles, combined with load control strategies, are then used to estimate peak load reductions from ACs DR.
[29]	K-means clustering is applied to the forecasted profiles of a large number of refrigerator units. One of the created clusters is then selected to be used as the target of the DR event based on its general characteristics and alignment with the purpose of the DR event.
[30]	Uses autoencoders to automatically extract data representation of the appliance data. Then, the K-means algorithm is applied to the latent features to cluster operation cycles with similar power signatures. The methodology is tested using residential washing machines and dishwashers.
[31]	Applies K-means clustering to the appliance activations, assuming that an appliance is active when the consumption is higher than 10% of its peak power. For clustering, the appliance activations are represented by eight human-engineered features, including mean, minimum, and maximum power values. This approach was tested using residential dishwashers, clothes washers, and dryers.
[32]	Applies Auto-Correlation Function (ACFs) for detecting the different appliance use patterns. Then, a Probability Distribution Analysis based on the ACF results and the Kullback–Leibler divergence is used to select the prevailing consumption patterns. The method was tested using four residential appliances: a kettle, a washing machine, a dishwasher, and a microwave.
[33]	Applies Pearson correlation on appliance consumption and explores the frequency of usage to identify their potential for DR. Several household appliances were studied, including microwaves, cooktops, and ovens.
[34]	Applies ACFs and Partial Auto-Correlation (PACFs) to identify trends and outliers in appliance consumption. The base assumption is that the presence of outliers occurs due to some appliance anomaly. This method was applied to household appliances, such as refrigerators, dishwashers, and furnaces.

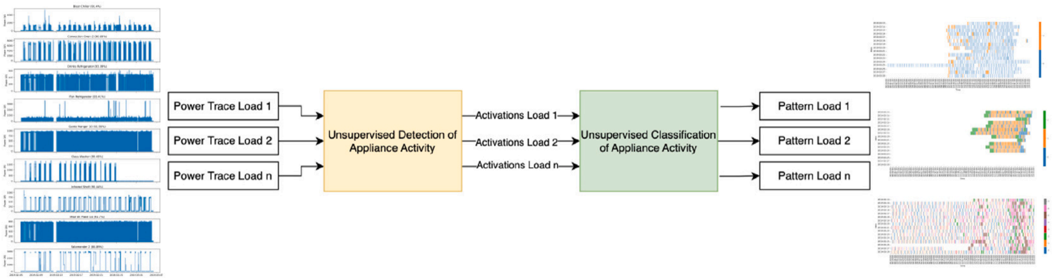

Fig. 1. Overview of the proposed methodology.

3. Methods

The proposed methodology is depicted in Fig. 1. The first step automatically identifies the appliance activations from the individual consumption time series (Section 3.1). Then, the appliance activities are classified in an unsupervised fashion using a data-driven feature extraction method and k-means clustering (Section 3.2).

3.1. Unsupervised detection of appliance activity

The operational states of appliances are identified by detecting and merging the appliance transitions in the power time series. This sub-section first addresses the definition of an appliance activity to reduce the ambiguity of ground-truth data acquisition (i.e., labeling) and automatic activity detection. Finally, it presents the proposed algorithm for automatic activity detection.

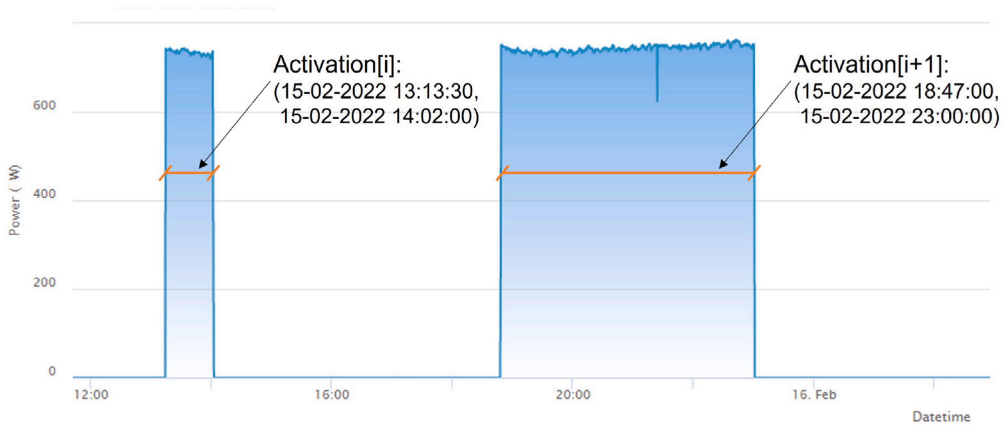

Fig. 2. Time series of the power consumption with identification of two activations for a salamander.

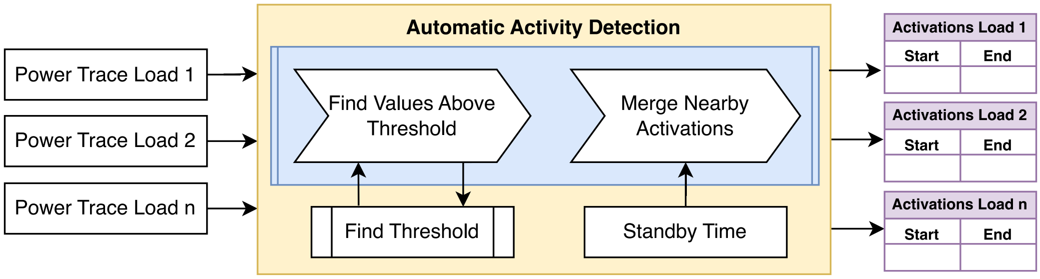

Fig. 3. Schematic representation of the workflow for automatically detecting appliance activities.

3.1.1. Definition of appliance activity

The activation period of an appliance corresponds to the time span since the appliance starts consuming energy until it returns to the same inactive state, where this inactive state could be a standby mode where the power consumption is negligible or a total switch-off period. These activations are bounded by events, which are associated with the change in operational states of appliances, and they are reflected on the time series as sudden changes in the average value of aggregate electricity consumption [35].

To illustrate the above, Fig. 2 depicts two consecutive activities in the same active power time series. Activation i is bounded by events at 15-02-2022 13:13:30 (ON) and 15-02-2022 14:02:00 (OFF), whereas the second activation $(i+1)$ is bounded by events 15-02-2022 18:47:00 (ON) and 15-02-2022 23:00:00 (OFF).

Appliance activation cycles depend on both the device's working mode and on human intervention. For example, heating lamps and infrared shelves' activations are directly controlled by human action, despite the possibility of implementing time controllers for automatic switching; in this case, it would still be humanly controlled. On the other hand, the refrigerator cycle is fully controlled by the integrated thermostats that regulate the cycles to maintain the setup temperature. Nevertheless, for such appliances, human behavior, such as opening and closing the door, highly influences the usage cycle. In this case, some of the appliances' activations are indirectly influenced by human action. The agent that promotes an activation period is out of the scope of this work because it would be required to have external information to confirm some of the hypotheses.

3.1.2. Automatic activity detection

The proposed algorithm for automatic activation detection falls into the category of heuristic-based approaches, and the main workflow is depicted in Fig. 3. The workflow starts with the input as a time series with the active power consumption and a user-defined standby time. This parameter controls the maximum time the appliance can be OFF between consecutive ON states, resulting in a single activation. This allows for mitigating the effect that instantaneous power shortage would have in producing several activations. The algorithm's output is a table with the information on the timestamps with the start and end time of each activation. The algorithm used to detect the activations is described in Algorithm 1.

This algorithm is implemented for use with the consumption data for a single appliance. Therefore, every event found belongs to this appliance and can be classified into turn-ON or turn-OFF, representing the beginning and end of the activation, respectively.

In short, the critical step of the proposed algorithm is to identify the steady state stages and observations that are lower than an evaluated threshold, marking those observations as the appliance is considered OFF and the remaining observation as ON. The standby time is then used to merge consecutive OFF-ON transitions that belong to the same appliance activation.

Algorithm 1: Activation detection

```
activations ← [];
for A in appliances do
    tseries ← get_appliance_consumption(A);
    threshold ← find_threshold(tseries);
    max_standby ← read_standby_parameter(A);
    A_status ← [ones()];
    for t, idx in tseries do
        if t < threshold then
            A_status[idx] ← 0;
        end
    end
    A_merged ← [];
    for idx in len(A_status(==1))-1 do
        if standby(idx, idx+1) < max_standby then
            A_merged ← merge_activations(idx, idx + 1);
        end
    end
    activations[A] ← A_merged
end
return activations
```

Algorithm 2: Find threshold

```
n_classes ← 100;
if max(tseries) > 1000 then
    n_classes ← 500;
end
histogram ← make_histogram(tseries, n_classes);
for b in get_bins(histogram) do
    if contains_zero(b) then
        r_limit ← max(b);
    end
end
return r_limit/2;
```

The threshold is defined by Algorithm 2. The algorithm generates a histogram of all measurements and chooses the threshold half of the supreme limit of the class that contains zero. The underlying assumption behind this algorithm is that this value will represent the minimum power change that would trigger an appliance transition. This decision is particularly relevant when considering appliances that are not necessarily OFF most of the time, such as refrigerators and freezers. An illustration of the method is presented in Fig. 4, showing the selected thresholds for the Blast Chiller (Fig. 4(a)) and the Refrigerator 1 (Fig. 4(b)).

In the last step, the algorithm merges activations that are separated by a standby or OFF period that is below a certain time delta. Currently, this parameter is defined by the user via an empirical analysis of consumption time series, but this could be improved by sensitivity analysis in a small portion of the dataset.

3.2. Unsupervised classification of activations

Since no additional information exists other than the time series values, the classification had to be done following an unsupervised approach. The proposed approach relies on clustering to group the appliance activations and is summarized in Fig. 5.

The input for this procedure consists of all the activations for a given load L, encoded as a matrix with dimension $A \times S$, where A is the total number of activations and S is the number of samples of the longest activation. Each row of the matrix contains one activation, and the rows are ordered chronologically. Each column corresponds to one sample within each activation. The different steps of the methodology are described next.

3.2.1. Data-driven feature extraction

The feature extraction step transforms the time-dependent dataset into a static one by capturing the dynamical properties of time series concisely as interpretable feature vectors [21]. This step was done independently for each activation. To this end, a data-driven approach using the open-source library *catch22* [36], and Principal Component Analysis (PCA) [37] was employed.

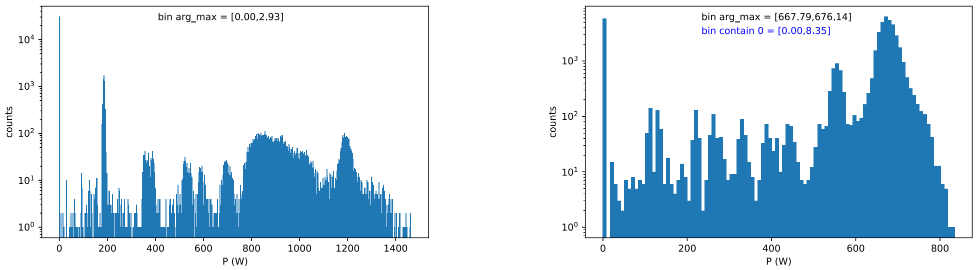

(a) Blast Chiller (b) Refrigerator 1

Fig. 4. Histogram of the instant active power measurements for two appliances.

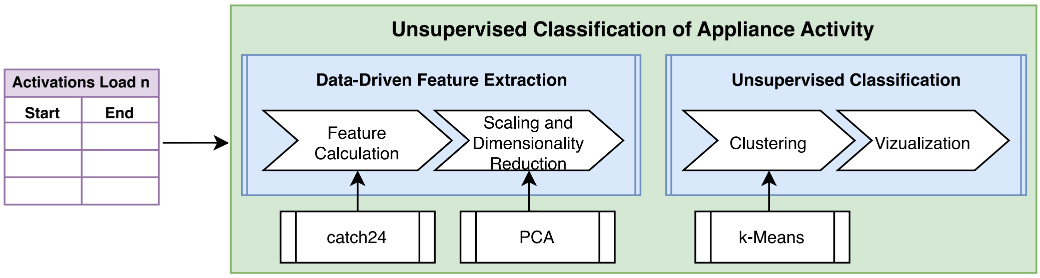

Fig. 5. Schematic representation of the workflow for unsupervised classification of appliance activities.

In the first step, *pycatch22* [21] was used to calculate 22 features for each activation. The *Catch22* feature set is effective for capturing appliance-specific consumption patterns due to its ability to represent diverse time-series properties. It identifies temporal variations through features like autocorrelation and periodicity, captures dynamic behaviors with entropy measures, and quantifies variability and extremes using statistical metrics such as variance and kurtosis. These capabilities make it well-suited to distinguishing between steady-use appliances (e.g., salamanders) and those with irregular or transient patterns (e.g., convection ovens). In addition to these features, given that the time series data is not normalized, the mean and variance were also considered, giving a total of 24 features (*catch24*).

In the second step, since it is not guaranteed that the features are mutually exclusive, dimensionality reduction using PCA was applied to the dataset of features. Given that the feature matrix contains attributes with different orders of magnitude between each feature, before implementing PCA, the data was normalized using a standard data transformation, i.e., subtracting the mean and dividing by the standard deviation. With the PCA strategy, the final set of features was selected based on the explained variance. All the principal components with explained variance higher than the one were selected as features and passed to the unsupervised classification procedure.

3.2.2. Unsupervised classification

For the unsupervised classification of the activities, this work relied on the well-known and widely used k-means algorithm [38]. The only required parameter is the number of clusters, k, used to partition the data. A distance function is also required to attribute each sample to the nearest k-clusters, which in this work was set to the Euclidean distance. The selection of the most appropriate number of clusters was performed by evaluating the silhouette coefficient [38]. This coefficient takes values from -1 to 1, where a value of 1 means that the clusters are well apart from each other, 0 means that the distance between clusters is not significant, and -1 indicates that the clusters are wrongly assigned.

The analysis of the best distribution of activations in clusters for each appliance obtained after the analysis of the silhouette coefficient was made through a graphical representation of the results. To this end, the visualizations consisted of the daily representation of the activations per appliance in a heatmap, having different colors assigned to different clusters. This representation allows for visually identifying routines of the use of the appliances and enables the interpretability of the grouping strategy.

4. Case study specification

The proposed methodology was evaluated using consumption data from nine appliances monitored from an IK during its normal operation. The present section first describes the appliance consumption data and the necessary preparation steps (Section 4.1).

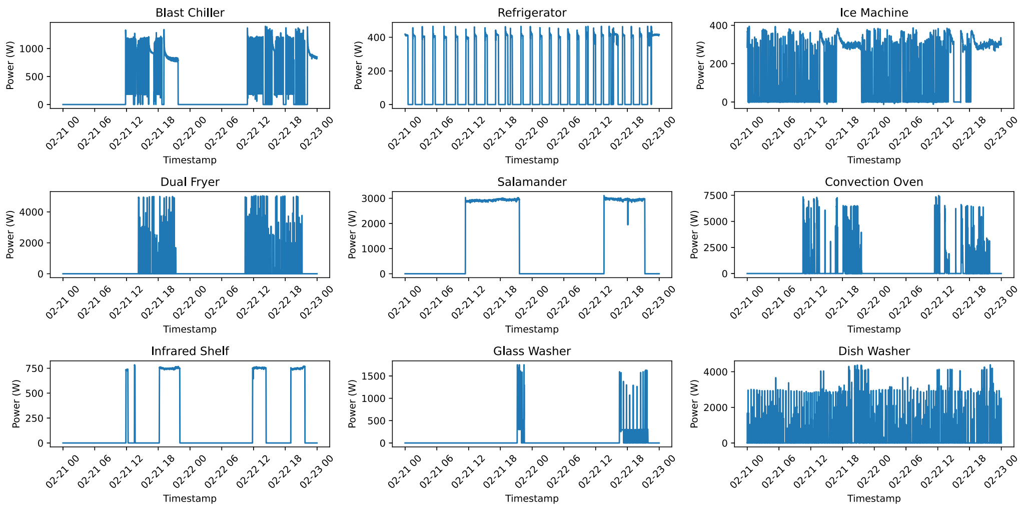

Fig. 6. Two days of active power for each appliance considered.

Then, the evaluation methodology is presented (Section 4.2), including details on the setups of the different components of the methodology and the different performance evaluation metrics.

4.1. Appliance consumption data

IKs are equipped with professional appliances, many of which have the same end as their domestic counterparts. In such cases, the main difference is the device's size, robustness, and intensity of use, which are naturally much higher in IKs. Overall, the studied appliances can be classified into five distinct groups: (1) the *cooling devices* are responsible for preserving the ingredients and prepared food. Some examples are refrigerators, freezers, and blast chillers; (2) the *food preparation* includes a wide variety of appliances such as induction plates, convection ovens, deep fryers, and salamanders; (3) *cleaning appliances* includes dishwashers and glasswashers; (4) appliances directed for *service* include infrared shelves and heating lamps; and (5) *HVAC* includes appliances such as the air conditioning and smoke extractors.

The consumption data for these appliances was taken from the FIKElectricity dataset [39], more precisely from IK 1. This dataset consists of aggregated and individual appliance electricity consumption data collected from three restaurant kitchens in Funchal, Portugal. The individual appliance consumption data were collected at the sampling rate of $\frac{1}{5}$ Hz. This particular IK was monitored for 26 consecutive days, from 06-02-2019 until 04-03-2019. This kitchen operated seven days a week for dinner, with occasional breakfast services. The nine appliances used are the following: blast chiller, refrigerator 1, ice machine, convection oven 1, dual dryer, salamander 1, infrared shelf, dishwasher, and glass washer. An overview of the appliance consumption traces of each appliance is provided in Fig. 6.

4.1.1. Data preprocessing

Before starting the application of the proposed methodology, a preliminary evaluation of the data showed consistent information for each appliance, namely the number of observations available was similar for all cases, and the range of values was within the expected values for the underlying physical quantities being measured.

The first days monitored in this kitchen showed data inconsistency due to the misplacement of some physical sensors and parameter adjustment of the acquisition system. For IK 1, there was also a period with no measurements due to an acquisition problem that occurred from 11:06 of 12-02-2019 until 01:16 of 13-02-2019, resulting in a period without observations for all sensors. By trimming all observations before 13-02-2019, it was possible to maintain a continuous time span of 17 days, from 13-02-2019 until 01-03-2019.

This period was also checked for missing observations. The salamander 1 and the dishwasher are the appliances with the larger number of missing observations, with 8% of missing data for the case of salamander 1. All the other appliances have missing values below 1.2% of the dataset. In this regard, time interpolation was used to fill the missing data. This method ensures continuity and a smooth transition from the left branch to the right branch in a linear model.

Outlier removal was also applied by removing the data points above the upper boundary, calculated by taking three standard deviations from the mean value. This technique ensures that all the data points are within a 99.7% confidence interval. The infrared shelf was the appliance with a larger number of outliers identified using this methodology, and the total number of observations modified is 1614, corresponding to 0.590% of the dataset.

Finally, the data was downsampled from $\frac{1}{5}$ Hz to $\frac{1}{30}$ Hz. Despite a possible loss of information, this allows decreased computation requirements, helps to reduce the existent noise in the original data, and still captures the load fluctuation at the appliance level since the appliances will be ON for much more than 30 s.

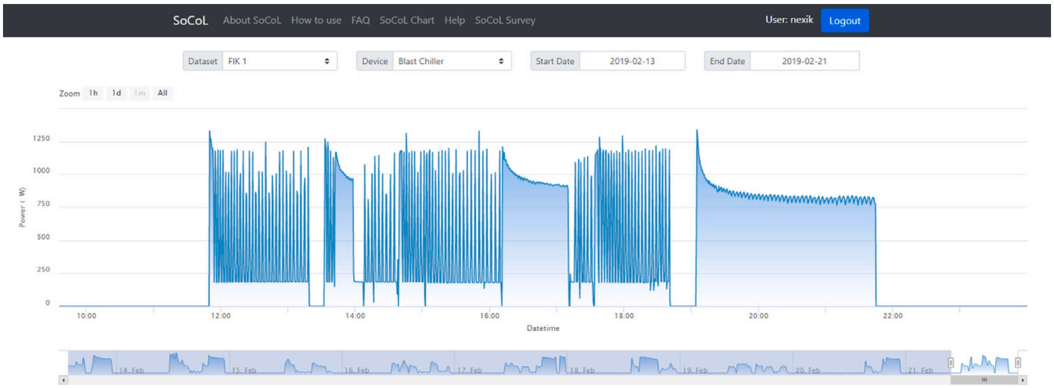

Fig. 7. Illustration of the SoCoL labeling platform.

Fig. 8. Labeling example of the dishwasher for a period of 24-02-2019.

4.1.2. Dataset labeling

The labeling of the data consisted of the process of generating a ground-truth dataset for evaluation of the algorithm proposed. This was a necessary task in this work since no labels are available in the original dataset [39].

The labeling process was processed in the Semi-automatic Collaborative labeling (SoCol) platform [40,41]. It was developed to be a collaborative labeling platform where several users can participate in the labeling process of energy datasets. Each registered user has access to the datasets that the administrator made available. The user can select one appliance at a time and the analysis period; the information is presented as a line graph with time in the x-axis and the power consumption in the y-axis as represented in Fig. 7.

The platform allows the user to mark events with a single click directly on the data point where an event starts or ends. A double click clears a wrong mark. To define appliance activations, it is necessary to create two marks, namely, the event that marks the beginning of such activation and another event for the end. The processing of such information to retrieve the activations has to be performed manually after the labeling process is concluded.

For a properly labeled dataset, to reduce bias and uncertainty, different experts should mark the same appliance, and a voting scheme should decide the final selection of the timestamps of the activation. In this case, that was not possible, and the voting scheme was performed only by the authors of this paper and the team members working in the nexIK research project.[1] The labeled dataset was created for a period of one week, from 22-02-2019 until 28-02-2019. One challenge identified during the labeling activity was maintaining labeling consistency through the process within one appliance and between different appliances. To better illustrate this, Fig. 8 shows the labels for the dishwasher on 24-02-2019. In this case, a double interpretation can be made in the cycle with three spikes between 12:45 and 13:05 and a similar pattern between 13:30 and 13:47.*Should they belong to the same activation period?* By capitalizing on the definition of activation in this work, these were considered three distinct activations, and the labels in the right group should be updated.

For the nine appliances selected for the ground-truth evaluation, a total of 2637 activations were registered. The descriptive statistics of this process are presented in Table 2. It presents the total number of activations per appliance and the time in minutes: total time of activations and min, max, and mean time per activation.

[1] nexIK project, https://nexik.tecnico.ulisboa.pt.

Table 2
Activations of the dataset based on the ground-truth evaluation from 2019-02-22 until 2019-02-28.

Appliance	Activities	Total time	Mean time	Min time	Max time
Blast Chiller	101	3552.5	30.36	1.5	390.0
Convection Oven 1	86	2826.0	32.86	1.5	276.0
Dishwasher	740	2521.5	3.40	1.5	31.5
Refrigerator 1	95	3429.5	36.1	6.5	813.5
Dual Fryer	527	1205.5	2.29	1.5	30.0
Glass Washer	121	757.0	6.26	2.0	47.5
Ice Machine	654	5770.5	8.82	0.5	550.0
Infrared Shelf	14	2666.5	190.46	1.5	551.5
Salamander 1	9	3601.0	400.11	4.0	636.5

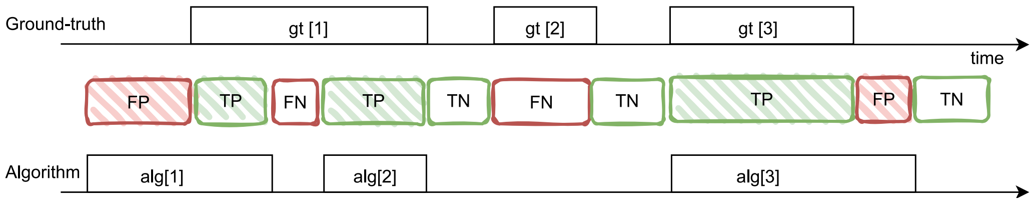

Fig. 9. Schematic representation of the elements of the confusion matrix: The activations are represented as blocks both for the ground-truth (top) and the results from the algorithm (bottom).

From the results, it is possible to observe that there are appliances with shorter cycles, such as the dishwasher, glass washer, dual fryer, and ice machine. These appliances have an average activation time below 9 min and many activations. On the other hand, some appliances have longer activation cycles, such as the blast chiller, the convection oven, and the refrigerator with long activations or extra long activations cycles, as the infrared shelf and the salamanders, with an average time per activation of approximately 35 min and higher than 3 h, respectively. These results show that in the case of the appliances that keep the food hot during service, they are usually turned ON when the service starts and turned OFF only at the end of service.

This labeled dataset will be used to evaluate the algorithm described next and analyzed in the next section.

4.2. *Performance evaluation methodology*

The assessment of the proposed methodology is conducted in two steps. First, the performance of the automatic appliance activity detection is evaluated against the manually labeled data. Then, the resulting appliance activities are assessed visually using time series plots that represent the resulting clusters. The quality of the clusters is identified using the silhouette score.

4.2.1. *Activity identification*

The performance of the proposed algorithm was evaluated using a set of metrics extracted from the confusion matrix considering the one week of labeled data. In this particular problem, as depicted in Fig. 9, the true values correspond to correctly identified time instants as ON or OFF stage, corresponding to True Positives (TPs) and True Negatives (TNs), respectively. The False Positives (FPs) are the time instants when the appliance is identified as being ON but is OFF, and the False Negatives (FNs) relates to the opposite situation.

To calculate the confusion matrix, it was first necessary to translate the timestamps of the ground-truth annotations and the algorithm activations into a matrix. This was done via one-hot encoding, which assigns zero to periods where appliances are OFF and one where appliances are ON [42]. The resulting matrix was then used to calculate the confusion matrix, which was done using the available functions of the *sklearn* package for *Python*. This methodology ensures a correct implementation for the above-mentioned interpretation of the four possible classifications of the tuple actual-predicted.

To quantify the performance of the proposed algorithm, different performance metrics were analyzed. One such metric is the Accuracy (ACC), which quantifies the number of correct observations as defined by Eq. (1). This metric is relevant in this context since it translates the model's capacity to correctly identify real activations while not triggering activations where they do not exist (i.e., keep a low number of FPs).

One important characteristic of appliance consumption data is that the number of activations is not homogeneous across appliances. For example, a refrigerator will be ON during several periods of the day. In contrast, heating lamps are used only about 10% of the time. As such, a performance metric such as ACC may be misleading, especially when considering appliances that are in the OFF mode most of the time [35]. A better approach in such cases is to consider the balanced Accuracy (bACC), which reflects the capacity to identify all real cases balanced by the wrong classifications, as defined in Eq. (2).

$$ACC = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

Table 3
List of monitored appliances with ground-truth labels and the evaluation results for the detection of activations (data from 22-02-2019 until 28-02-2019).

Appliance	E	TN	FP	FN	TP	ACC	bACC	R	P	F1-score
Blast Chiller	7105	13 043	12	21	7084	0.998	0.998	0.997	0.998	0.998
Convection Oven 1	5652	14 360	148	121	5531	0.987	0.984	0.979	0.974	0.976
Dish Washer	5043	15 086	31	146	4897	0.991	0.984	0.971	0.994	0.982
Refrigerator 1	6859	13 296	5	78	6781	0.996	0.994	0.989	0.999	0.994
Dual Fryer	2411	17 740	9	276	2135	0.986	0.943	**0.886**	0.996	**0.932**
Glass Washer	1514	18 642	4	2	1512	0.999	0.999	0.999	0.998	0.998
Ice Machine	11 541	8465	154	378	11 163	0.974	0.975	0.967	0.986	0.977
Infrared Shelf	5333	14 818	9	2	5331	0.999	0.999	0.998	0.999	**0.999**
Salamander 1	7202	12 955	3	1	7201	0.999	0.999	0.999	0.999	**0.999**
Sum	52 660	128 405	375	1025	41 635	–	–	–	–	–
Mean	–	–	–	–	–	0.992	0.986	0.976	0.994	0.985
StDev	–	–	–	–	–	0.009	0.019	0.036	0.008	0.020

$$bACC = \frac{1}{2} \cdot \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right) \tag{2}$$

Precision and Recall are other metrics often considered to evaluate scenarios with data imbalance. Precision and Recall are given by Eqs. (3) and (4), respectively. Precision focuses on the positive events, indicating the flagged activations that are real activations. In contrast, Recall analyses the missed opportunities, accounting for the real activations that were not correctly identified. Finally, the F1-score (given by Eq. (5)) is also used in the performance evaluation. This metric, which combines Precision and Recall via the harmonic mean, translates to the problem domain as the ability of a model to capture the appliance activations (Recall) and simultaneously be accurate in the cases it captures (Precision) [43].

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

$$Recall = \frac{TP}{TP + FN} \tag{4}$$

$$F_1 score = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{5}$$

4.2.2. *Activity classification*

All the activations identified for the nine appliances were analyzed to gather insights into their characteristics. The 24 features were reduced by PCA. In the set of final features, the number of components after dimensionality reduction was 5 for the convection oven 1, glass washer, ice machine, and infrared shelf; 6 for the blast chiller, dishwasher, dual fryer, and salamander 1; and 7 for the refrigerator 1.

The k-means algorithm requires starting the algorithm with a known number of clusters. For this problem, 2 to 8 groups were tested for each appliance, and the selection of the best number of partitions for each case can be made via the evaluation of the silhouette coefficient. The highest silhouette score indicates the best distribution of the activations per cluster, as this score favors the maximum distance of observations between clusters and the minimum distance within each group. Heatmaps were then used to visualize the cluster arrangements.

5. Results and discussion

The results were analyzed individually for the activity identification (Section 5.1) and unsupervised activity classification (Section 5.2). Additional graphs, data, and source code are available in an online repository https://doi.org/10.17605/OSF.IO/SPWD9.

5.1. *Activity identification*

The algorithm was applied directly to the prepossessed time series of the active power of each appliance with aggregated power consumption with a frequency of 2 observations per minute. The standby parameter was defined by visual inspection of the data, being set to two minutes (four samples) for the blast chiller, three minutes (six samples) for the convection oven 1 and the dishwasher, 30 s (1 sample) for the refrigerator 1, and 1 min (2 samples) for the remaining five devices. The obtained results are presented in Table 3.

The table shows that the proposed detection methodology performs exceptionally well across most appliances, with F1-scores and balanced accuracy consistently exceeding 0.98. Appliances like the Infrared Shelf and Salamander exhibit near-perfect detection performance (F1 = 0.999, bACC = 0.999), indicating highly accurate activation detection. Overall, the methodology demonstrates excellent consistency with a mean F1-score of 0.985 and balanced accuracy of 0.986, suggesting the model works reliably across

Table 4
Number of activations labeled, found by the algorithm from 22-02-2019 until 28-02-2019, and activations found in the full data.

Appliance	Labeled activations	Detected in labeled data	Detected in full data
Blast Chiller	117	112 (−5)	234
Convection Oven 1	86	89 (−3)	181
Dish Washer	740	715 (−25)	1689
Refrigerator 1	95	96 (+1)	230
Dual Fryer	527	506 (−21)	1076
Glass Washer	121	118 (−3)	485
Ice Machine	654	676 (+22)	1460
Infrared Shelf	14	17 (+3)	38
Salamander 1	9	10 (+1)	21
	2363	2339 (−24)	5414

various appliances in the industrial kitchen setting. This is further supported by the relatively low standard deviations (StDev), with a maximum deviation of 0.036 in F1-scores.

However, some appliances, particularly the Dual Fryer and Ice Machine, show slightly lower performance, with the Dual Fryer having an F1-score of 0.932 and balanced accuracy of 0.943. This suggests that the detection method faces challenges with appliances exhibiting more complex or fluctuating consumption patterns. Despite this, the overall precision and recall remain strong, and the True Positives far outweigh the misclassifications.

Given the superior performance of the proposed algorithms, the next step was to apply it to the remaining dataset samples for which no ground truth is available. This was performed such that a complete set of activations to test the unsupervised classification algorithm was available. Table 4 summarizes the number of activations found in three situations, the labeling process and the algorithm from 22 until 28-02-2019, and the total number of activations found in the dataset from 13-02-2019 until 01-03-2019. Note that per the definition in Section 3.1.1, one activation corresponds to the time span since one appliance starts consuming energy until it returns to the same inactive state.

As observed, the proposed algorithm was able to approximate the number of activations of each appliance correctly. Looking at the individual results, it is important to notice the dual fryer for which 21 activities were missed. This is related to the lower value for the Recall metric observed in Table 3. The same observation is true for the dishwasher. Another interesting observation is that of the ice machine, for which there are 22 false detections. This is caused by the high number of FPs observed in Table 3. Nevertheless, it is important to stress that these errors reflect only a very small percentage of the total activations. More precisely, −3.98% in the dual fryer, −3.37% in the dishwasher, and +3.36% in the ice machine.

In summary, the results demonstrate that the proposed event detection model performs well across the dataset, with only minor discrepancies between labeled and detected activations. The model shows good scalability, effectively detecting appliance activations in a larger dataset beyond the labeled data, highlighting its robustness. While some discrepancies suggest areas for improvement in detection accuracy, the model's ability to generalize and scale to different appliances and larger datasets is a positive outcome, indicating its potential for broader applications in real-world scenarios.

5.2. Activity classification

The activations identified when considering the entire dataset were analyzed to gain insights into their working characteristics. The silhouette scores obtained for each appliance and the number of clusters are depicted in Fig. 10. From these results, several observations can be drawn.

First, all the silhouette scores are > 0 independently of the appliance. However, only in the case of the dual-fryer is there a clear separation with two clusters (score = 0.7). Interestingly, grouping the data into two clusters seems to be the best option in several appliances despite having less pronounced silhouette scores. This is the case with the blast chiller, refrigerator, and infrared self (scores of 0.54, 0.53, and 0.58, respectively). Concerning the ice machine, the selected number of clusters is five. However, four clusters would also be a good option since the silhouette scores are very close (0.56). The dishwasher is the appliance that generates the highest number of clusters, which can be either seven or eight since the score is practically the same (0.54). Surprisingly, despite being an appliance of the same category, the number of clusters of the glass washer is only three (score = 0.50). Finally, the convection oven is the only appliance that does not have a silhouette score greater than 0.5 for any possible cluster arrangements. Still, in this case, the selection would be for the three clusters (0.414 vs. 0.401). The best number of clusters for each appliance and the respective silhouette scores are summarized in Table 5.

To better understand these results, it is interesting to verify when each activation takes place and if the activations identified as belonging to the same cluster are repeated every day at the same time. This is a reasonable interpretation as industrial kitchens have routines that should be very similar every day, which should be reflected in the use of appliances. To this end, the daily consumption of each individual appliance will be analyzed using heatmaps, where the activations are colored based on the best cluster distribution selected by the highest silhouette coefficient. Each row corresponds to one day, and the columns are the hours of the day. The results are presented by appliance category, i.e., cooling, food preparation, service, and cleaning, to facilitate the analysis.

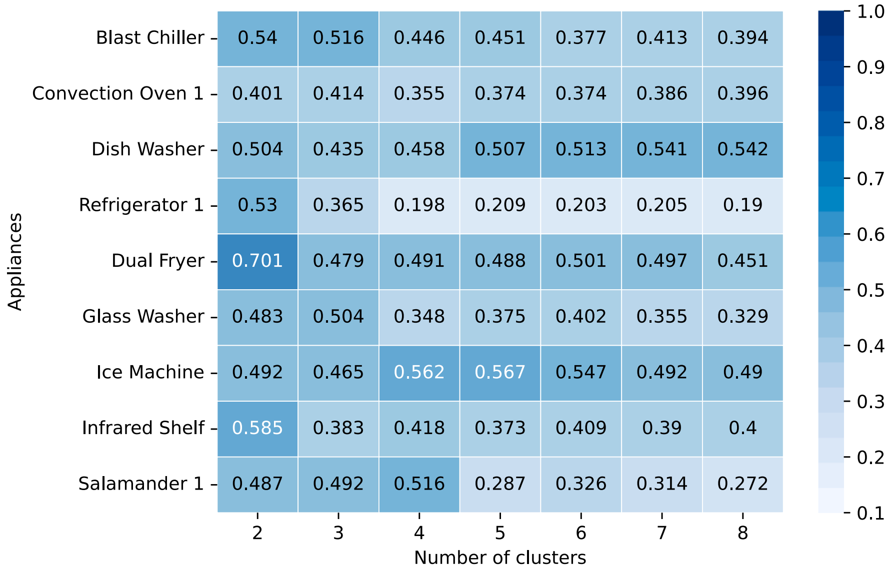

Fig. 10. Average silhouette coefficient obtained by the clustering algorithm using k-means with 2 to 8 clusters and feature extraction using *catch24*.

Table 5
Summary of the best number of clusters selected for each device, and the respective silhouette score.

Appliance	Clusters	Score
Blast Chiller	2	0.54
Convection Oven 1	3	0.414
Dish Washer	8	0.542
Refrigerator 1	2	0.53
Dual Fryer	2	0.701
Glass Washer	2	0.504
Ice Machine	5	0.567
Infrared Shelf	2	0.585
Salamander 1	4	0.516

As a remark, it should be noted that the clustering algorithm has no information about the time instant any of the activations started since all the information about time instants was lost in the process before the feature extraction step. Therefore, any time correlation the following analysis can retrieve is related to the appliance activations' cycles.

5.2.1. Cooling devices

There are three devices in the cooling category: blast chiller, refrigerator, and ice machine.

Blast Chiller. A blast chiller is a piece of equipment that is used to quickly lower the temperature of food. Typically, one such device can bring the temperature of foods down from 70 (°C) to 5 (°C) or less in about 90 min.

As per Fig. 11, it can be seen that this appliance operates from morning to evening, meaning that the cooling of foods is performed very often in this restaurant. The only exception for this was the 26 of February when the chiller was left ON also during the night (most likely by oblivion). Regarding the clusters, the first one (cluster 0) groups all the smaller activations together. These activations only happen when the device is turned ON and serve essentially to keep the device at a certain temperature. The other cluster, on the other hand, indicates the periods when the blast chiller is working. Interestingly, the activations under this cluster tend to last at least 90 min, taking longer on some occasions. The fact that longer activations occur has several interpretations. One would be placing food at different initial temperatures. Another one would be placing different quantities of food.

Refrigerator 1. For the refrigerator, the highest silhouette occurs for two clusters. The distribution of the clusters is shown in Fig. 12.

This figure shows that the number of elements in each cluster is very different, with more than 90% of the activation in cluster 1. The dominating cluster contains all the activations that show a normal operation, whereas cluster 0 contains considerably short activations. Interestingly, there is no separation between the very short activations and the longer ones, which leads to also looking at the results for three clusters (Fig. 13).

It can be observed that when considering three clusters, a new cluster is created that contains mostly the activations with an average duration (cluster 2). While this clustering seems to make more sense, one possible explanation for the lower silhouette score

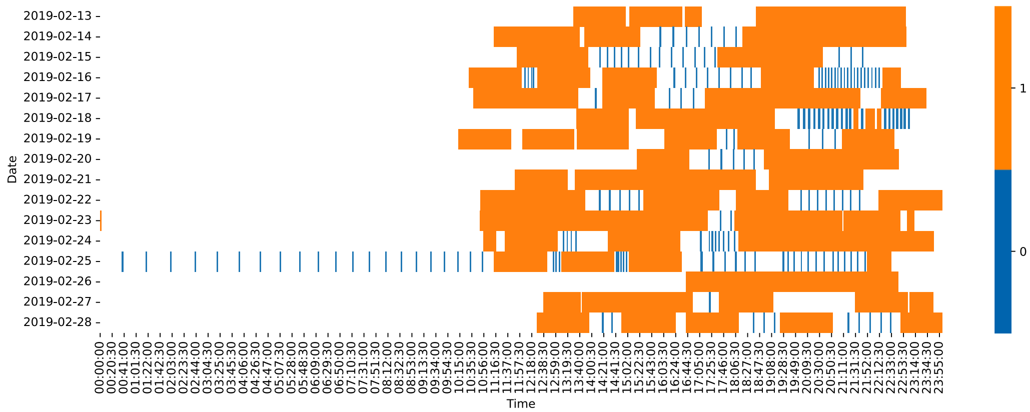

Fig. 11. Clusters of activations for the blast chiller in a date × time matrix.

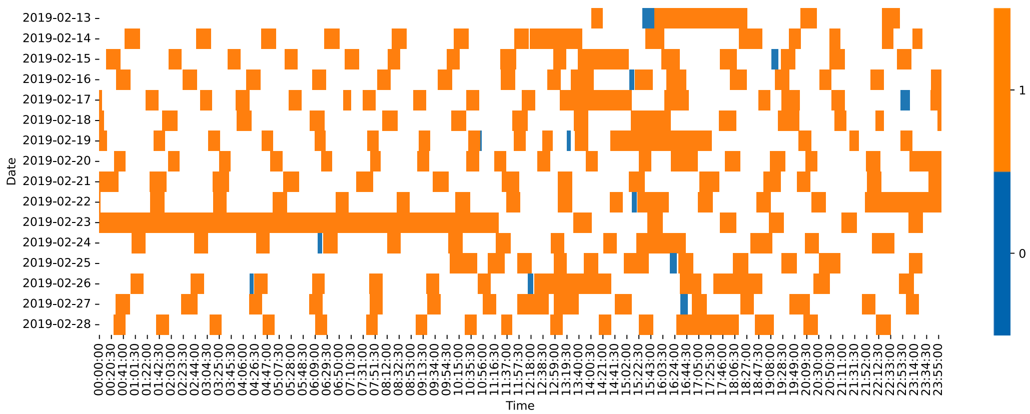

Fig. 12. Clusters of activations for the refrigerator in a date × time matrix.

is the presence of one activation in cluster 1 that is much longer than the others (February 22 and 23), which causes an increase in the intra-cluster distance of this cluster and necessarily a decrease in the silhouette score.

Ice Machine. For this appliance, the clustering of activations and the analysis of the silhouette coefficient suggest five partitions with a silhouette score of 0.567. The results are depicted in Fig. 14.

Clusters 3 and 4 are frequent during service time and show long periods of continuous motor operation. Cluster 4 has only three occurrences, two of which are characterized by long duration. This suggests that these can be merged with cluster 2, which is what happens when considering only 4 clusters.

The other three clusters are mostly visible during the out-of-service periods and are therefore associated with keeping the temperature of the systems at the pre-established values. Interestingly, the clustering can separate this into three moments. The first moment (cluster 3) shows a slightly higher consumption and longer duration. In the other two moments, differences in duration are not easily observed. Instead, the clustering is related to the power levels, with more power in cluster 1 than in cluster 0. Interestingly, this is a very different behavior from the refrigerator, which always operates at the same power levels.

5.2.2. Food preparation

Three devices can be found in the food preparation category: dual fryer, salamander, and convection oven.

Dual Fryer. An industrial fryer is an appliance used to heat large volumes of cooking oil for deep frying large batches of food. It is very commonly found in restaurants. The results of the best cluster distribution for this appliance are shown in Fig. 15.

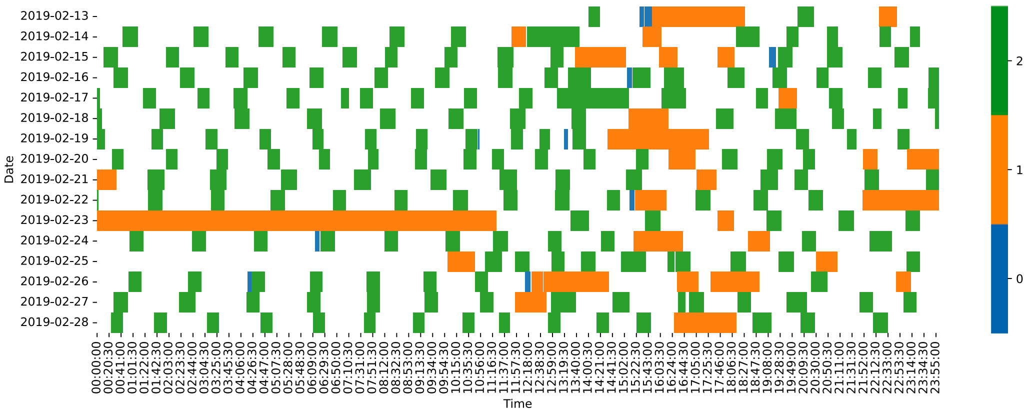

Fig. 13. Clusters of activations for the refrigerator in a date × time matrix. Results of the second-best silhouette score.

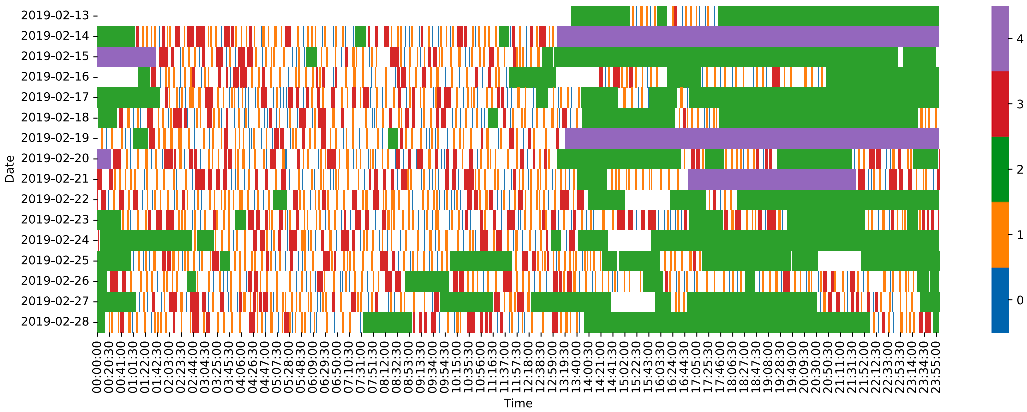

Fig. 14. Clusters of activations for the ice machine in a date × time matrix.

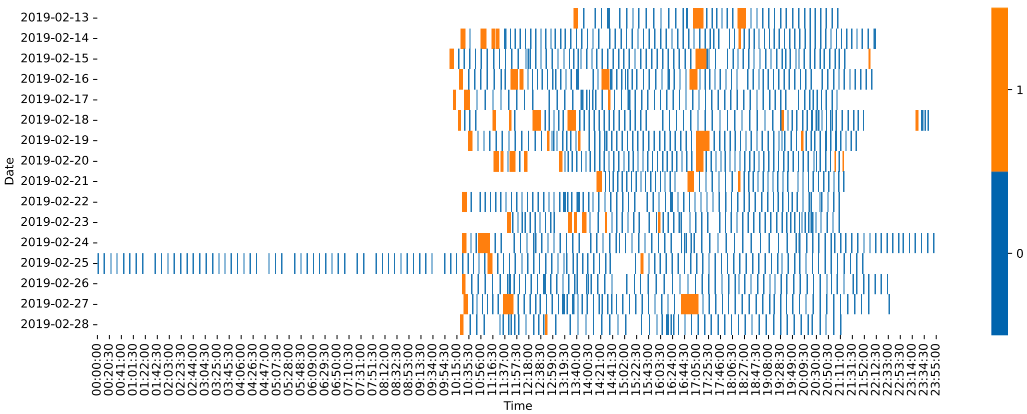

Fig. 15. Clusters of activations for the dual fryer in a date × time matrix.

This appliance is only used during service time, except for 25-02-2019, where the data suggest the appliance was not turned OFF. The activations belonging to Cluster 0 are short and recurrent, corresponding to cycles that maintain the oil at a certain temperature. The orange activations, Cluster 1, correspond to more intensive periods of work, where the oil was heated for a larger temperature difference, from cold to hot, or when the flux of food preparation increased; for example, when a large number of guests order french-fries simultaneously the thermal inertia of the oil is not sufficient to keep up the temperature and a more intensive activation has to start to match the setup temperature.

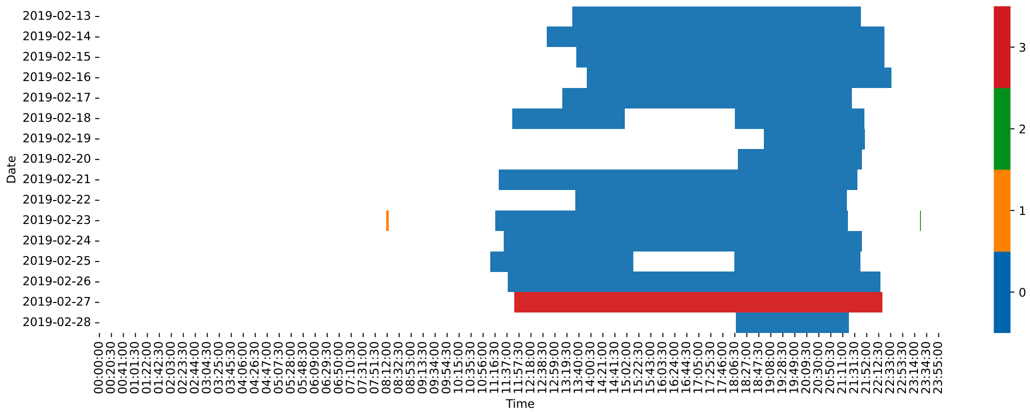

Fig. 16. Clusters of activations for the salamander in a date × time matrix.

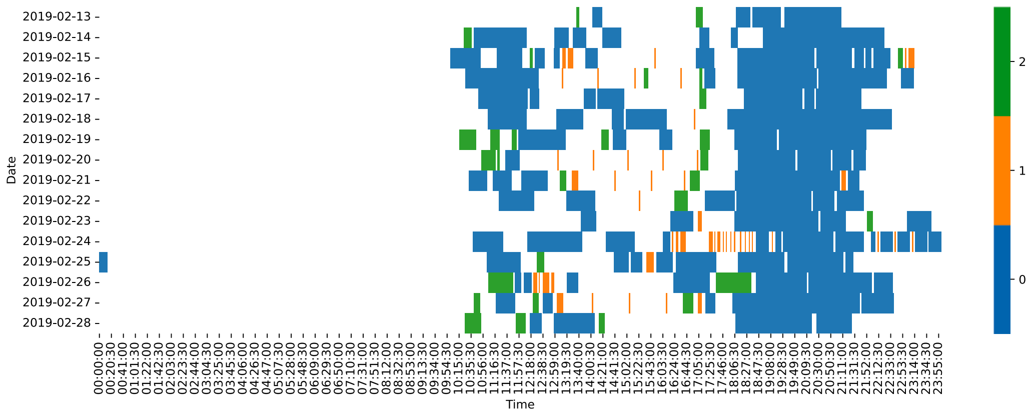

Fig. 17. Clusters of activations for the Convection Oven in a date × time matrix.

Salamander. Salamanders are specialized kitchen appliances dedicated to broiling foods at higher temperatures than a traditional oven's broil setting. A salamander uses infrared or radiant heat to push heat evenly onto food. The results of the best cluster distribution for this appliance are shown in Fig. 16.

This appliance is dominated by cluster 0, which contains 17 of the 20 activations. The smaller clusters correspond to very short activations in one specific day, whereas there is one cluster with one single activation on the 27th, which occurs due to a change in the power, possibly due to the downsampling to 30 s.

Ultimately, the clustering shows that the usage of this appliance follows a very strict routine, which is somewhat expected since it has a very simple mode of operation, which is to keep a resistor to a certain temperature. It is interesting to notice that even though there are some days where the usage is split into two, or one single activation, the clustering does not take this into consideration since more weight is given to power drawn than the energy consumed.

Convection Oven. For the convection oven, the highest silhouette occurs for three partitions. The distribution of the clusters is shown in Fig. 17.

This appliance is used more intensively during service time, often in the evening. However, it can be seen that it is also used in the morning for preparation for the evening service. Cluster 0 shows the activations that correspond to the periods of more intense work. Cluster 1 contains the smallest activations, which seem to occur when food is not being prepared but the device is still open. Finally, cluster 2 is interesting since it occurs a few moments before cluster 1. The power consumption in this cluster is slightly higher than that of cluster 0, which may indicate the need to reach a certain temperature. Interestingly, when using only two clusters, the activations in cluster 2 merge with cluster zero, indicating that the duration is the main feature here.

5.2.3. Service appliances

This service category includes only the infrared shelf, which keeps the food warm before being sent to the table.

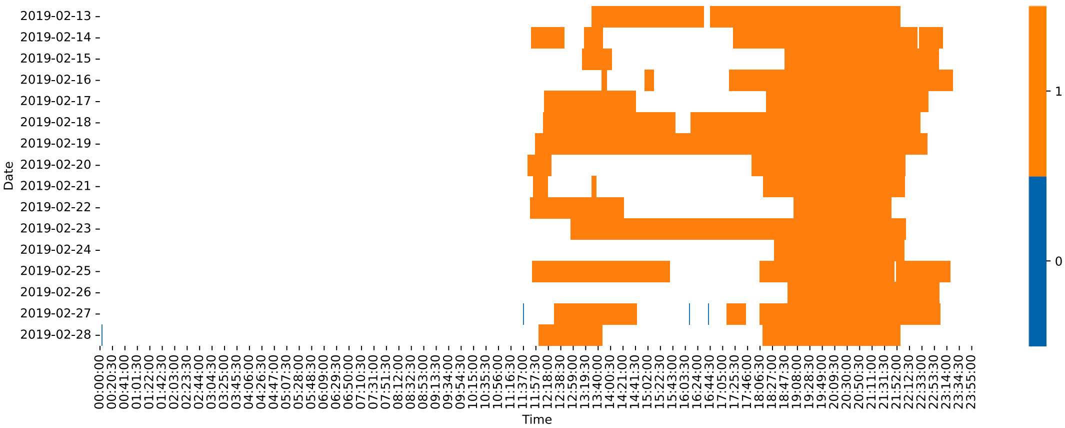

Fig. 18. Clusters of activations for the infrared shelf in a date × time matrix.

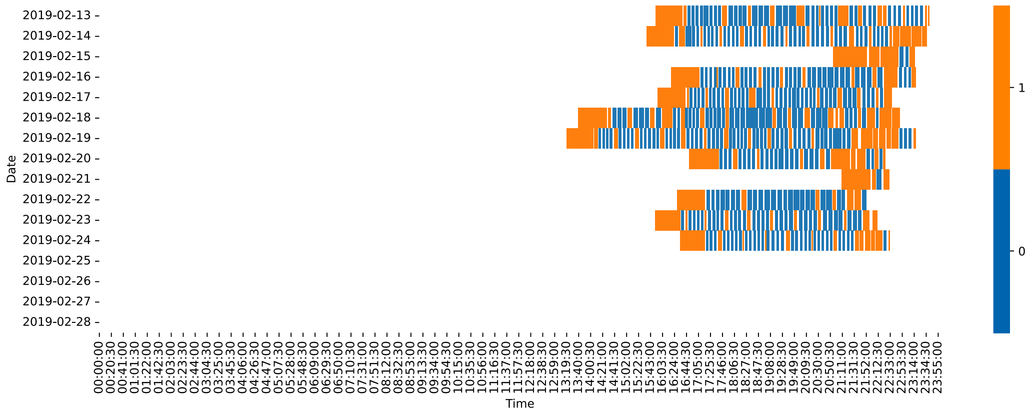

Fig. 19. Clusters of activations for the Glass Washer in a date * time matrix.

Infrared Shelf. The analysis of the silhouette coefficient suggests the organization of the activations into two clusters, with a silhouette value of 0.585. The results are depicted in Fig. 18.

This appliance has a behavior that is very similar to the salamander, which is somewhat expected since both appliances are mainly resistive loads. Like with the salamander, the duration of the activation has an important role in the grouping of observations. Cluster 0 contains the four activations with the smallest duration, which can be associated with anomalies in the smart meter reading or a simple human-made verification to check if the appliance was, for example, connected to power.

5.2.4. Cleaning appliances

There are two appliances in the cleaning category: glass washer and dishwasher.

Glass Washer. For this appliance, the best number of clusters is two. The results are depicted in Fig. 19.

This appliance is only used during service time. The first cleaning of the day is always the most intensive, identified in Cluster 1. This group of activations is also popular at the end of service. The occurrences of cluster 0 are frequent during the day. Since these activities do not represent glass-washing activities, they represent an opportunity to save energy by turning the glass washer off when not in use.

Dish Washer. The dishwasher is an appliance from the same group as the glass washer. However, the obtained results (depicted in Fig. 20) are considerably different.

The first major difference is that this appliance has activities 24 h per day, not only during preparation, service, and cleaning time, as expected. The clustering framework suggests grouping the activation into eight distinct partitions with a silhouette coefficient of 0.542.

Cluster numbers 2, 5, and 6 are abundant during service hours. Except for cluster 6, the others appear outside of the period the kitchen is open. All the other clusters are frequent throughout the day. This suggests that this appliance is activated during the day, even if not for dishwashing. After careful inspection of the appliance datasheet, it was found that this behavior occurs for a self-cleaning operation and to keep the hot water in a reservoir at a pre-established temperature. Like with the glass washer, since no actual dishwashing is happening during this period, it would be possible to turn the device off to save energy.

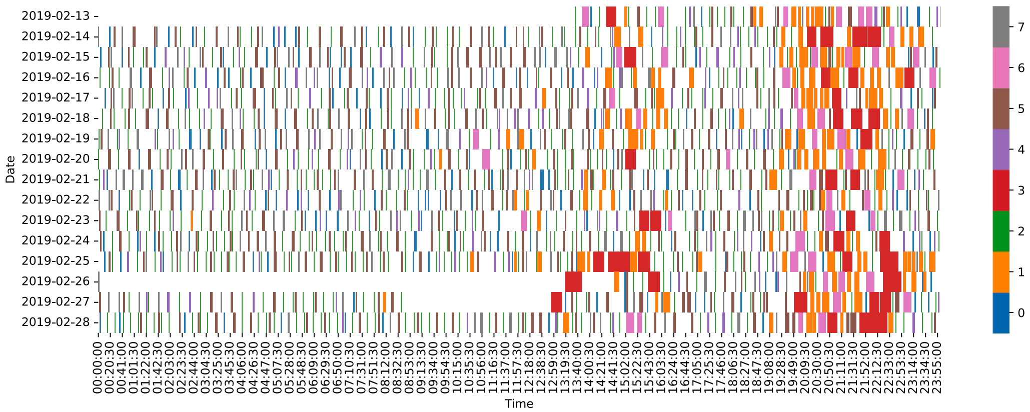

Fig. 20. Clusters of activations for the Dish Washer in a date * time matrix.

6. Conclusion

In this work, a data-driven methodology was proposed to analyze individual appliance consumption data in IKs. The proposed methodology leverages the concept of appliance activity to unveil different patterns in energy consumption. To this end, the first step of the methodology is to identify the activations in the individual appliance consumption measurements using a heuristic-based method. In the second step, the detected activities are classified via unsupervised learning. This was done via transforming the appliance activities' time series into static representations by feature extraction, leveraging the *catch22* feature set for time series data. The clustering was done using a k-means algorithm considering 2 to 8 clusters per appliance.

The proposed methodology was applied to the power consumption signals of nine industrial kitchen appliances taken from the FIKElectricity dataset. In particular, the consumption of the following appliances was considered: blast chiller, refrigerator, ice machine, convection oven, dual dryer, salamander, infrared shelf, dishwasher, and glass washer.

Concerning automatic activity detection, when compared to the manually labeled activities, the proposed solution showed excellent results in terms of balanced accuracy and F1-score, both above .98 on average. In this regard, the least performing appliance was the dual fryer, with 0.93 in the F1-Score (0.94 in the balanced accuracy). In contrast, the most performing appliances were the infrared shelf and the salamander, with 0.999 F1-score and balanced accuracy scores. Ultimately, the very good ability to correctly identify individual appliances clearly indicates that the method can be applied with a very high confidence level to other datasets even when no labeled data is available. Despite most existing works focusing on household appliances, our results surpass benchmarks from other studies that use similar methodologies, which typically report F1-scores and accuracies ranging from 0.85 to 0.95 [40,44]. This demonstrates the effectiveness and robustness of our approach in detecting appliance activities in the context of industrial kitchens, where the operational complexity and variability of appliance behaviors are much higher.

Concerning the unsupervised classification, each appliance showed a different optimum number of clusters, and the methodology proposed can identify different cycle modes within the same appliance. Overall, the feature selection method shows some important dependence on the activation duration when defining the clusters. Still, the load fluctuation pattern seems to be the most crucial feature. Ultimately, the load fluctuation patterns reflect the operational cycles of appliances, which are often characterized by periods of high and low power consumption. These fluctuations typically correspond to the on–off cycles, heating or cooling phases, or varying power needs during different cooking stages.

This is especially visible in the ice machine, which despite having a similar working principle to a refrigerator, has a different way of drawing power from the grid. More specifically, while both rely on refrigeration cycles to regulate temperature, they differ in their specific functions and operations. The refrigerator typically maintains a consistent, lower temperature to store food, while the ice machine goes through intermittent cycles of freezing water to create ice, which leads to more pronounced load fluctuations. The ice machine's power consumption pattern, therefore, shows more significant peaks corresponding to the freezing process, whereas the refrigerator has more stable consumption with periodic increases when the compressor starts to work. In contrast, the salamander draws power exactly the same way, independently of the duration of the activity. This occurs because this appliance typically maintains a consistent temperature for cooking or warming, and its power consumption is not influenced by the length of time it is in use.

6.1. Research implications and potential applications

Overall, this methodology gives detailed information on the energy use in IKs. For example, it allowed the detection of possible outliers in the activations of the infrared self. Moreover, it also unveiled energy efficiency opportunities in some appliances, such as the dishwasher and blast chiller, which should be turned OFF when not being used since they consume a fair amount of energy when in standby mode.

The ability to detect, classify, and categorize individual appliance activities can also be used to develop off-peak and peak-demand forecasting methods, which are crucial when planning the deployment and control of Distributed Energy Resources (DERs) such as PV and BESSs. Furthermore, understanding appliance activities is also critical to developing algorithms to assess and forecast the potential of certain appliances to participate in DR programs specifically designed for small-scale consumers [7]. The possibility of participating in such programs is very relevant for restaurants that have no means, either technical or financial, to install DERs [6]. Both these aspects have been widely studied in the residential sector, as highlighted in the literature [25,26,45,46], and should be adapted to the IKs sector.

Finally, having access to appliance activation information also enables the quantification of the inter-dependences between appliances. This is particularly relevant when considering different aspects of IK operation, such as menu composition, substitutability of appliances, and the simultaneity of requests, as a means not only to reduce electricity consumption but also to address peak demand concerns [10,47].

6.2. Limitations and future work

While the objectives of this work have been reached, some aspects should be improved in future iterations.

First, despite having used a real-world dataset, the number of days is quite short, which may affect the clustering results. Hence, in future work, it would be interesting to replicate the proposed methodology with a longer dataset if one becomes available. Furthermore, it was observed that, on some occasions, the method is very sensitive to variations in the consumption data. Hence, it would be relevant if, in future work, a data cleaning procedure could be applied to the detected appliance activations.

Additionally, despite the merits of the data-driven feature extraction methodology, this results in poor explainability since it is hard to understand the relation between the features and the consumption data. Hence, in future work, it would be interesting to develop human-engineered features and compare them with data-driven features to understand eventual patterns that help explain the obtained results.

Another limitation is the dependence on the level of granularity in data collection. The methodology assumes a high sampling resolution, and its effectiveness with lower-resolution datasets has not been tested. Addressing this in future work will help determine the robustness of the approach across diverse datasets.

At last, despite the potential of IKs to contribute to the energy transition, the research in this field is still in an early stage, and numerous challenges remain unexplored [12]. Among those challenges is the lack of publicly available datasets to develop and validate new and existing approaches. In this regard, an interesting future research direction would be to extend this methodology to generate synthetic datasets. To state more precisely, by keeping track of the sequence of cluster activations over time, it should be possible to develop a probabilistic generative model (e.g., Hidden Markov Models (HMMs) or Bayesian Networks (BNs)) to synthesize appliance consumption traces. Likewise, another important research direction would be to leverage the proposed methodology to study the interactions across appliances and, with that, unveil the dynamics of electricity consumption in IKs. Such an analysis could be done, for instance, by quantifying the synchrony between activations of different appliances using techniques such as Time-Lagged Cross-Correlations (TLCCs) and Windowed Time-Lagged Cross-Correlations (WTLCCs).

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

Lucas Pereira received funding from the Portuguese Foundation for Science and Technology (FCT) under grants CEECIND/01179/2017 and UIDB/50009/2020. Hugo Morais received funding from FCT under grant UIDB/50021/2020. This research was funded by FCT under the Exploring the Human-Water-Energy Nexus in Industrial Kitchens (nexIK) project (EXPL/CCI-COM/1234/2021).

Data availability

The data used in this work is publicly available in the following repositories: https://doi.org/10.17605/OSF.IO/K3G8N; https://doi.org/10.17605/OSF.IO/SPWD9.

References

[1] van den Eijnde J. The human touch in kitchen technology: How technology changes our relationship with food in the rational and ritual kitchen. APRIA J 2020;2(1):68–84.
[2] Kiefer NM. Economics and the origin of the restaurant. Cornell Hotel Restaur Adm Q 2002;43(4):58–64.
[3] EnergyStar. Energy star for small business: Restaurants. 2017, URL https://www.energystar.gov/buildings/facility-owners-and-managers/small-biz/restaurants.
[4] Ahuja K, Chandra V, Lord V, Peens C. Ordering in: The rapid evolution of food delivery. Tech. rep., McKinsey & Company; 2021.
[5] Food and beverage services global market report 2023. Tech. rep. 6193681, The Business Research Company; 2023, p. 400.
[6] Legrand W, Sloan P, Simons-Kaufmann C, Fleischer S. A review of restaurant sustainable indicators. In: S. Chen J, editor. Advances in hospitality and leisure. Advances in hospitality and leisure, vol. 6, Emerald Group Publishing Limited; 2010, p. 167–83. http://dx.doi.org/10.1108/S1745-3542(2010)0000006013.
[7] Pedro A, Krutnik M, Yadack VM, Pereira L, Morais H. Opportunities and challenges for small-scale flexibility in European electricity markets. Util Policy 2023;80:101477. http://dx.doi.org/10.1016/j.jup.2022.101477.
[8] Renewable energy in the water, energy and food nexus. Tech. rep., International Renewable Energy Agency (IRENA); 2015, p. 128.
[9] Stephan RM, Mohtar RH, Daher B, Embid Irujo A, Hillers A, Ganter JC, Karlberg L, Martin L, Nairizi S, Rodriguez DJ, Sarni W. Water–energy–food nexus: A platform for implementing the sustainable development goals. Water Int 2018;43(3):472–9. http://dx.doi.org/10.1080/02508060.2018.1446581.
[10] Mudie S, Vadhati M. Low energy catering strategy: insights from a novel carbon-energy calculator. Energy Procedia 2017;123:212–9. http://dx.doi.org/10.1016/j.egypro.2017.07.244, URL http://www.sciencedirect.com/science/article/pii/S1876610217328151.
[11] Young R. Electrification of commercial kitchens. 2020.
[12] Oliveira A, Ribeiro M, Rodrigues R, de Morais G, Morais H, Pereira L. On the role of industrial kitchens in sustainable energy systems: The nexIK vision. In: CIRED 2023 - The 27th international conference and exhibition on electricity distribution. Rome, Italy; 2023.
[13] Amantegui J, Morais H, Pereira L. Benchmark of electricity consumption forecasting methodologies applied to industrial kitchens. Buildings 2022;12(12):2231. http://dx.doi.org/10.3390/buildings12122231.
[14] Andrade V, Morais H, Pereira L. Industrial kitchen appliance consumption forecasting: Hour-ahead and day-ahead perspectives with post-processing improvements. Comput Electr Eng 2024;116:109145. http://dx.doi.org/10.1016/j.compeleceng.2024.109145.
[15] Hedrick R, Smith V, Field K. Restaurant energy use benchmarking guideline. Tech. rep. NREL/SR-5500-50547, 1019165, NREL; 2011.
[16] Hearnshaw SA, Essah EA, Grandison A, Felgate R. Energy reduction and benchmarking in commercial kitchens. In: Technologies for sustainable built environments. Reading, UK; 2012, p. 8.
[17] Mudie S. Energy benchmarking in UK commercial kitchens. Build Serv Eng Res Technol 2016;37(2):205–19.
[18] Young W. Applying solar energy to food trucks. In: Proceedings of the SOLAR 2017 conference. Denver: International Solar Energy Society; 2017, p. 1–8. http://dx.doi.org/10.18086/solar.2017.05.01, URL http://proceedings.ises.org/citation?doi=solar.2017.05.01.
[19] Hashmi MU, Cavaleiro J, Pereira L, Bušić A. Sizing and profitability of energy storage for prosumers in madeira, Portugal. In: 2020 IEEE power & energy society innovative smart grid technologies conference. ISGT 2020, Washington, DC, USA: IEEE; 2020.
[20] Pereira L, Cavaleiro J, Morais H. Understanding the role of solar PV and battery energy storage in a snack bar: A case study in madeira island. In: 2022 IEEE 20th international conference on industrial informatics. INDIN, Lemgo, Germany: IEEE; 2023.
[21] Lubba CH, Sethi SS, Knaute P, Schultz SR, Fulcher BD, Jones NS. Catch22: Canonical time-series characteristics. Data Min Knowl Discov 2019;33(6):1821–52.
[22] Yildiz B, Bilbao JI, Dore J, Sproul AB. Recent advances in the analysis of residential electricity consumption and applications of smart meter data. Appl Energy 2017;208:402–27. http://dx.doi.org/10.1016/j.apenergy.2017.10.014.
[23] Völker B, Reinhardt A, Faustine A, Pereira L. Watt's up at home? smart meter data analytics from a consumer-centric perspective. Energies 2021-01;14(3):719. http://dx.doi.org/10.3390/en14030719.
[24] Haq EU, Lyu X, Jia Y, Hua M, Ahmad F. Forecasting household electric appliances consumption and peak demand based on hybrid machine learning approach. Energy Rep 2020;6:1099–105. http://dx.doi.org/10.1016/j.egyr.2020.11.071.
[25] Abera FZ, Khedkar V. Machine learning approach electric appliance consumption and peak demand forecasting of residential customers using smart meter data. Wirel Pers Commun 2020;111(1):65–82. http://dx.doi.org/10.1007/s11277-019-06845-6.
[26] Ji Y, Buechler E, Rajagopal R. Data-driven load modeling and forecasting of residential appliances. IEEE Trans Smart Grid 2020;11(3):2652–61. http://dx.doi.org/10.1109/TSG.2019.2959770.
[27] Gao M, Pan S, Chen S, Li Y, Pan N, Pan D, Shen X. Identification method of electrical load for electrical appliances based on K-Means ++ and GCN. IEEE Access 2021;9:27026–37. http://dx.doi.org/10.1109/ACCESS.2021.3057722.
[28] Malik A, Haghdadi N, MacGill I, Ravishankar J. Appliance level data analysis of summer demand reduction potential from residential air conditioner control. Appl Energy 2019;235:776–85. http://dx.doi.org/10.1016/j.apenergy.2018.11.010.
[29] Faria P, Jozi A, Vale Z. Forecasting refrigerators consumption to support their aggregated participation in demand response. In: 2020 IEEE international conference on environment and electrical engineering and 2020 IEEE industrial and commercial power systems europe. EEEIC / I&CPS Europe, Madrid, Spain: IEEE; 2020, p. 1–6. http://dx.doi.org/10.1109/EEEIC/ICPSEurope49358.2020.9160527.
[30] Castangia M, Barletta N, Camarda C, Quer S, Macii E, Patti E. Clustering appliance operation modes with unsupervised deep learning techniques. IEEE Trans Ind Informatics 2023;19(7):8196–204. http://dx.doi.org/10.1109/TII.2022.3217495.
[31] Marcello F, Pilloni V. Energy consumption profiling of appliances inside smart buildings based on K-means clustering. In: 2023 IEEE international mediterranean conference on communications and networking. MeditCom, 2023, p. 240–5. http://dx.doi.org/10.1109/MeditCom58224.2023.10266646.
[32] Góis J, Pereira L. A novel methodology for identifying appliance usage patterns in buildings based on auto-correlation and probability distribution analysis. Energy Build 2022;256:111618. http://dx.doi.org/10.1016/j.enbuild.2021.111618.
[33] Cruz C, Tostado-Véliz M, Palomar E, Bravo I. Pattern-driven behaviour for demand-side management: An analysis of appliance use. Energy Build 2024;308:113988. http://dx.doi.org/10.1016/j.enbuild.2024.113988.
[34] Priyadarshini I, Alkhayyat A, Gehlot A, Kumar R. Time series analysis and anomaly detection for trustworthy smart homes. Comput Electr Eng 2022;102:108193. http://dx.doi.org/10.1016/j.compeleceng.2022.108193.
[35] Pereira L, Nunes N. An empirical exploration of performance metrics for event detection algorithms in non-intrusive load monitoring. Sustain Cities Soc 2020-11-01;62:102399. http://dx.doi.org/10.1016/j.scs.2020.102399.
[36] Pycatch22. 2022, http://dx.doi.org/10.5281/zenodo.6673597, URL https://github.com/DynamicsAndNeuralSystems/pycatch22.
[37] Maćkiewicz A, Ratajczak W. Principal components analysis (PCA). Comput Geosci 1993;19(3):303–42. http://dx.doi.org/10.1016/0098-3004(93)90090-R.
[38] Kaufman L, Rousseeuw PJ. Finding groups in data: An introduction to cluster analysis. John Wiley & Sons; 2009.
[39] Pereira L, Aguiar V, Vasconcelos F, Martins R, Garcês T, Morais H. FIKElectricity: A Electricity Consumption Dataset from Three Restaurant Kitchens in Portugal. Scientific Data 2023;10(1):779. http://dx.doi.org/10.1038/s41597-023-02698-8.
[40] Pereira L, Ribeiro M, Nunes N. Engineering and deploying a hardware and software platform to collect and label non-intrusive load monitoring datasets. In: 2017 sustainable internet and ICT for sustainability. SustainIT, Funchal, Portugal: IEEE / IFIP; 2017, p. 1–9. http://dx.doi.org/10.23919/SustainIT.2017.8379791.
[41] SoCol: Semi-automatic collaborative labelling. 2023, online, URL https://ml.socol.online/. [Accessed 25 October 2022].

[42] Okada S, Ohzeki M, Taguchi S. Efficient partition of integer optimization problems with one-hot encoding. Sci Rep 2019;9(1):13036. http://dx.doi.org/10.1038/s41598-019-49539-6.

[43] Powers DMW. Evaluation: From precision, recall and f-measure to roc., informedness, markedness & correlation. J Mach Learn Technol 2011;2(1):37–63.

[44] Völker B, Scholl PM, Becker B. Semi-automatic generation and labeling of training data for non-intrusive load monitoring. In: Proceedings of the tenth ACM international conference on future energy systems. E-energy '19, New York, NY, USA: ACM; 2019, p. 17–23. http://dx.doi.org/10.1145/3307772.3328295.

[45] Ji Y, Rajagopal R. Demand and flexibility of residential appliances: An empirical analysis. In: 2017 IEEE global conference on signal and information processing. GlobalSIP, 2017-11, p. 1020–4. http://dx.doi.org/10.1109/GlobalSIP.2017.8309115.

[46] Afzalan M, Jazizadeh F. Residential loads flexibility potential for demand response using energy consumption patterns and user segments. Appl Energy 2019;254:113693. http://dx.doi.org/10.1016/j.apenergy.2019.113693, URL http://www.sciencedirect.com/science/article/pii/S0306261919313807.

[47] Nonaka T, Shimmura T, Fujii N, Mizuyama H. Energy consumption in the food service industry: A conceptual model of energy management considering service properties. In: Advances in production management systems: Innovative production management towards sustainable growth. IFIP advances in information and communication technology, Cham: Springer International Publishing; 2015, p. 605–11. http://dx.doi.org/10.1007/978-3-319-22759-7_69.