

Article

Shapelets to Classify Energy Demand Time Series

Marco G. Pinheiro ^{1,2} , Sara C. Madeira ^{3,*}  and Alexandre P. Francisco ^{1,4} 

¹ Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisbon, Portugal; marco.pinheiro@tecnico.ulisboa.pt (M.G.P.); aplf@tecnico.ulisboa.pt (A.P.F.)

² EDP, Av. 24 de Julho 12, 1249-300 Lisbon, Portugal

³ LASIGE, Faculdade de Ciências, Universidade de Lisboa, 1749-016 Lisbon, Portugal

⁴ INESC-ID Lisboa, Rua Alves Redol 9, 1000-029 Lisbon, Portugal

* Correspondence: sacmadeira@fc.ul.pt

Abstract: Data are an important asset that the electric power industry have available today to support management decisions, excel in operational efficiency, and be more competitive. The advent of smart grids has increased power grid sensorization and so, too, the data availability. However, the inability to recognize the value of data beyond the siloed application in which data are collected is seen as a barrier. Power load time series are one of the most important types of data collected by utilities, because of the inherent information in them (e.g., power load time series comprehend human behavior, economic momentum, and other trends). The area of time series analysis in the energy domain is attracting considerable interest because of growing available data as more sensorization is deployed in power grids. This study considers the shapelet technique to create interpretable classifiers for four use cases. The study systematically applied the shapelet technique to data from different hierarchical power levels (national, primary power substations, and secondary power substations). The study has experimentally shown shapelets as a technique that embraces the interpretability and accuracy of the learning models, the ability to extract interpretable patterns and knowledge, and the ability to recognize and monetize the value of the data, important subjects to reinforce the importance of data-driven services within the energy sector.

Keywords: energy; energy digitalization; electricity demand; machine learning; time series classification; shapelets



Citation: Pinheiro, M.G.; Madeira, S.C.; Francisco, A.P. Shapelets to Classify Energy Demand Time Series. *Energies* **2022**, *15*, 2960. <https://doi.org/10.3390/en15082960>

Academic Editors: Christina N. Papadimitriou, Andrea Michiorri and Carsten Hoyer-Klick

Received: 14 March 2022

Accepted: 13 April 2022

Published: 18 April 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The electric power industry has been subject to constant changes. Utility companies have been aware of the threats and opportunities that arise as a result of this change. Climate action calls for the electric power industry to participate in the energy transition, once it is an important player in integrating more variable renewable energy sources and guaranteeing energy distribution for new needs, such as charging electrical vehicles, green hydrogen production, industrial heating electrification, and other opportunities for electrification all supported by renewable energy. The growing popularity of “behind the meter” on-site generation and storage, the new digital retail competitors that serve customers with bundle solutions towards energy-as-a-service business models, increased search for grid flexibility, and intense public and regulatory scrutiny are just some strengths that reinforce the need for more digitalization in the energy sector [1–4].

Data are an important asset that utilities have available today to support management decisions, excel in operational efficiency, and be more competitive. Moreover, data and technology to extract value from these data, generally addressed by artificial intelligence and machine learning, are no longer just a technological enabler, but rather an integrated part of the acceleration of the energy transition. The complex number of data sources along the value chain, led by higher levels of grid sensorization, has resulted in data streams whose value has not been fully explored.

Data-driven services within the energy sector pose challenges across regulatory, socio-economic, and organizational (RSEO) aspects. The review [5] states that the value of data from various sources must be clearly understood in order to overcome identified organizational barriers related to the lack of data compatibility between different sources, the complexity of the data, and the inability to recognize the value of the data in addition to the siloed application in which data are collected.

Smarts grid deployment has led the industry sector to recognize the inherent value of data besides the billing and settlement functions: asset monitoring, behavior profiling, customer classification, load curve classification, and many other uses to be considered using the same dataset. The study of pattern and knowledge extraction might therefore help to understand human activities as energy consumers, raising the general understanding about energy demand, and thus helping not only today's power grid operations and decisions, but also define policies for tomorrow's operations [6]. By virtue of their function, smart grids gather data mostly in the form of time series, that is, observations of the same signal over time.

In the last decade, new series-based algorithms have been developed and new studies have analyzed the contribution of these algorithms to extract value and insights from series across various sectors. This includes sequential data such as sequences of numeric values, text, audio, and even image. If the sequence is time-stamped, the sequence is generally named a time series. Time series forecasting is an major area with extensive literature that infers or estimates further steps in the series [7]. The clustering of time series seeks to discover temporal patterns with an unsupervised approach [8]. On the other hand, with a supervised approach, time series classification considers discriminatory features dependent on the ordering to organize time series into predefined labels or classes [9]. A set of data mining techniques for time series has also been developed: symbolic representation [10], motifs [11], discords [11], shapelets [11], time series chains [12], snippets [13], semantic segmentation [14] and so on.

When considering the energy-specific domain in the application of those techniques, shapelets is one that was not extensively applied. From the literature review, shapelets were used for non-intrusive load monitoring (NILM) [15], discovering customer weekend load patterns [16], classification of district heating substations [17], evaluation of voltage stability [18–20], and clustering power curves [21] with a modified version of shapelets to work as an unsupervised technique. Although with a few applications, it is evident that shapelets have not yet extensively assessed time series with a power load at the national level, as well as at primary and secondary substations. This paper aims to fill the gap in the literature on the energy domain, exploring the shapelet technique with four different goals with the same dataset. The dataset is related to the power load measured quarter-hourly at three levels: national, primary, and secondary substations. The data and the method are fully described in Section 2. Section 3 provides four use cases: (i) which pattern identifies weekend load curves from business days, (ii) which load pattern identifies Mondays from the rest of the business days, (iii) a classifier capable of identifying the load dynamics due to maneuvers across the grid, and (iv) a classifier capable of identifying the type of energy consumption from just the daily load curve. The study assesses whether shapelets are a technique capable of responding to the four challenges, and thus reinforces the value of smart grid data in addition to the siloed application in which data are gathered.

The most important added benefit provided by this study is the demonstration of the value and information that can be extracted as interpretable patterns from one of the most meaningful types of data collected in the energy sector, the load time series. Using appropriate machine learning techniques—in this case, shapelets—it is possible to extract value, which reinforces the importance of multi-source data-driven services within the energy sector and across its value chain. Section 4 draws conclusions and discusses the results.

Related Work

Time series do not have explicit features. For instance, each data point is a value, but the pattern or trend depends on the data points close to it. Whereas most classifiers consider discriminatory features non-dependent on the ordering values, which can be poisoned by even low levels of noise and distortions, the shapelets are local features, and thus they capture patterns as subsequences inside a time series able to discern different classes.

Shapelets were introduced by [22] as a new primitive representation of a time series that is highly representative of a class. At that time, they introduced shapelets as interpretable, more accurate, and significantly faster than state-of-the-art classifiers. Instead of looking for the global shape of time series and calculating the distance between them and a class representative, the shapelets technique only compares a local subsection of the shape that is particularly class discriminating. Shapelets have been continuously improved. The enormous quantity of shapelet candidates makes brute-force (exhaustive) shapelet discovery very slow for large datasets. Early abandon of Euclidean distance calculation combined with early entropy pruning were ideas introduced in the original paper [22]. Additional articles emphasize the reuse of computations and the pruning of the search space [23], while the projection of time series to the SAX representation was also elaborated by [24]. Furthermore, [25] proposed a novel fast method that avoids measuring the prediction accuracy of similar candidates through an online clustering/pruning technique. Although the shapelets concept was originally considered a supervised technique, it was exploited and extended that shapelets can be learned from unlabeled time series and used for unsupervised clustering [21]. A recent study proposes dynamic shapelets in which the differing representative of a long time series is considered in different time slices, as well as the evolution pattern of shapelets [26]. The evolution of shapelets is modeled as a graph, which represents how a time series evolves in terms of shape and interpretable patterns throughout time. To validate, the authors conducted experiments based on five time series datasets from different domains.

Shapelets were used in a few use cases in the energy domain with relative success. Recognizing the potential of shapelets as an interpretable technique for studying power load curves, the literature is not as extensive as one would like. Non-intrusive load monitoring discerns the individual electrical appliances of a residential or commercial building by disaggregating the accumulated energy consumption data without improving the individual sensorization of each appliance. The introduction of shapelets to this problem was proposed by [15], which starts with shapelets discovering from the current signatures of the recorded device instances present in a labeled database. The 60 Hz recorded samples consider the first few seconds of a device operation to search the shapelet instead of the entire time frame. Classification of district heating substations using shapelets as feature extraction was introduced by [17]. In this study, the classifier models are improved after the use of augmented features extraction from the shapelet transformation. The transformation is not described in detail, but generally consists of the difference between the original time series and the shapelets previously discovered from them. Another author has applied the shapelets technique in the assessment of voltage stability [18–20]. With the development of smart grids, phasor measurement units are massively available, and so are the data, making it possible to use data mining techniques such as shapelets. In this case, the author has to deal with classes that are imbalanced and incrementally updated, and with the online application of the model. Another study, which approximates to ours, analyzes the effectiveness of the shapelet algorithm in classifying various weekend consumption patterns extracted from real-life data [16]. The author suggests the potential of shapelets to determine which customers use more electricity at weekends, a period in which consumption usually reduces, which saves during peak hours, or who responds to demand response events. However, the assessment considered only the weekend pattern discovery using 15 minutes of the consumption data of 33 buildings with different noise levels. None of the studies applied the technique to power load time series at different levels

of aggregation (national, primary, and secondary substations) considering a systematic method and evaluation.

2. Materials and Methods

This section overviews the energy load time series, and introduces the methodology and key concepts used in the research.

2.1. Data Overview

A dataset of energy demand was used to perform the study of time series classification using shapelets. The dataset contains measurements that denote the national electricity demand (Mainland Portugal) between 2006 and the first quarter of 2014. These data also include measurements of electrical substations and secondary substations that represent two different levels in the hierarchy of the national power system. The energy consumption is dependent on human behavior and activities, and those reflect on the patterns of energy time series.

The time series resolution of the dataset is 15 min, i.e., a measurement for each 15 min. The longest time series—the national demand—are 8 years long: 288,768 points. The available substation measurements are 3 years long.

At the national level, the electricity demand follows a pattern with annual, weekly, and daily seasonality. Figure 1 shows the demand for a week from Monday to Sunday, in winter and summer.

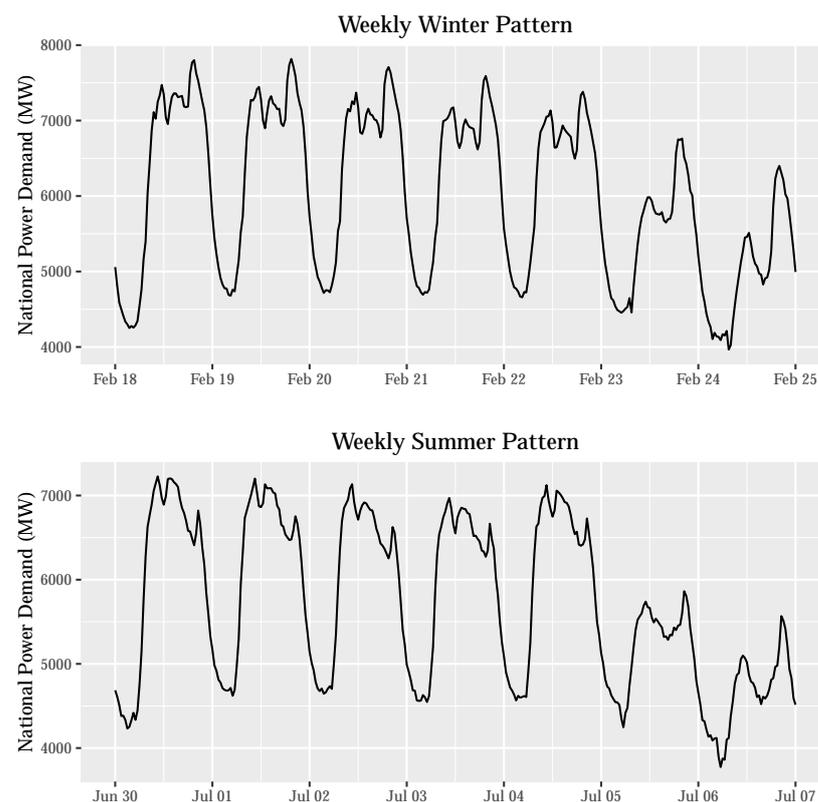


Figure 1. Portuguese demand from Monday to Sunday, in February 2008 and July 2008.

At the substation level, the electricity demand also follows a pattern with annual, weekly, and daily seasonality. However, due to the dynamics of the distribution grid and the possibility of shifting the power load from one substation to another according to the management of the grid operator, there exist high load peaks and periods where the substation is shut down and the measure values are zero, as shown in Figure 2.

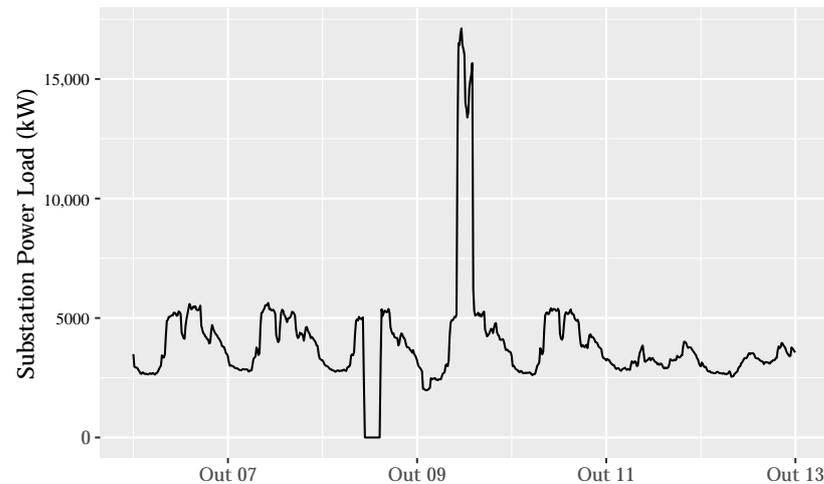


Figure 2. The substation load from Monday to Sunday with a shut-down on day 8 and a high load peak due to a load shift on day 9.

At the secondary substation level, seasonalities are generally kept. However, as the secondary substations feed a high-demand building, venue or industry, or even a residential neighborhood, the load curve is more unpredictable when compared with higher levels (power substations and national levels) and patterns might be found in accordance with the specific usage of electricity at that point of energy delivery. Generally, unpredictability and noise increases as we move down the level hierarchy (national, primary and secondary substations).

Additionally, a calendar dataset was available with the following information: (i) day of the week, (ii) public holidays of Portugal, (iii) local statutory holidays for Portugal's municipalities, (iv) strike days, and its local disrupted services when applicable.

2.2. Definitions

A simple symbolic sequence is an ordered list of symbols of a given alphabet. The dataset used in the use cases is a sequence of real values. So, we are actually interested in classifying time series. Time series can be univariate or multivariate. We considered only the simplest.

Definition 1. A time series T is a sequence of real values typically ordered in ascending order by timestamp. For example, $T = \langle (t_1, v_1), (t_2, v_2), \dots, (t_l, v_l) \rangle$ is a simple time series of length l that records data points from time t_1 to t_l . A subsequence S of time series T is a sampling length $m \leq l$ of contiguous positions from T , that is, $S = \langle (t_p, v_p), \dots, (t_{p+m-1}, v_{p+m-1}) \rangle$ for $1 \leq p \leq p+m-1 \leq l$.

Definition 2. The distance from a subsequence to a time series, $\mathcal{D}(S, T)$, returns a non-negative value which is the Euclidean distance (or other distance function) between S and its best matching location somewhere in T , that is, where the distance is minimum.

Definition 3. A dataset D of length n is a set of time series T_i and its class label c_i . Formally, $D = \langle T_1, c_1 \rangle, \dots, \langle T_n, c_n \rangle$ and $c_1, \dots, c_n \in C$, the set of possible labels. $\#C_i$ is the number of time series in class C_i .

Definition 4. The entropy \mathcal{E} of the dataset D is defined as $\mathcal{E}(D) = -\sum_{i=1}^{\#C} p_i \log_2(p_i)$ where $p_i = \frac{\#C_i}{n}$ is the probability of the class.

Definition 5. A split is a tuple $\langle S, \theta \rangle$ of a subsequence S and a distance threshold θ (or separation gap) that separates the dataset into two smaller datasets, D_L and D_R . When a split separates the dataset with the maximum information gain, the subsequence S is denominated as a shapelet.

Definition 6. The information gain \mathcal{G} of a split strategy that divides D into two subsets D_L and D_R of length n_L and n_R is given by $\mathcal{G} = \mathcal{E}(D) - \frac{n_L}{n} \mathcal{E}(D_L) - \frac{n_R}{n} \mathcal{E}(D_R)$.

Classifying with a shapelet S and its corresponding split gap θ produces a binary decision node on whether a time series belongs to a certain class or not. The shapelets are embodied on the nodes of a decision tree to create a universal classifier. At each step of decision tree induction, the shapelet and the corresponding split gap are computed on the training subset considered in that step [22].

2.3. Method

The four use cases follow a systematic method. First, a scenario is set that includes the question to be answered. The question is usually related to whether an interpretable pattern is discoverable by shapelets, which explains how two or more classes are discriminated. Second, the appropriate dataset is chosen with the respective data wrangling to set up the time series and its labels. Usually, the time series are transformed into daily load curves and chosen to balance the classes. For example, for the weekends use case, the daily national load curves are randomly selected within each class (weekend and business day). Daily curves that present outliers or missing data are kept out of the training and testing datasets. Third, the method applies the steps of the shapelet algorithm: (i) a pool of candidates is created from the time series inputs and minimum and maximum shapelet-length parameters, (ii) the best performing candidates are ranked using the information gain criterion over the target (or other prediction quality metrics like the Kruskal–Wallis or Mood’s median [27], or F-Stats [28]), (iii) the chosen best performing candidate is used as a tree node to create a tree-based model interactively, (iv) the time series input are split using the model built so far and a new iteration starts on each of the new leaves, (v) the algorithm stops when pruning parameters are met or a whole leaf time series corresponds to a class. Finally, it is checked if the shapelets were able to answer the initial question, and the performance of the resulting decision tree classifier is evaluated against the testing dataset through the accuracy metric, $\frac{\sum_k TP_k}{N}$, where TP_k are the objects correctly classified by the true label k , and N is the number of objects.

3. Results

This section will describe four use cases based on energy load time series of different power levels. Four particular classification problems were studied in order to assess decision trees using shapelets applied to energy data.

3.1. Weekends

The first classification problem is related to the lower consumption that occurs on weekends. In this problem, the goal is to discover whether a 24 h time series corresponds to a weekend or a business day. As described above, the national load follows a very distinctive trend over the weekend, but even these are different over the year.

The input data are a set of time series of 96 points each (24 h with a quarter-hour resolution), extracted from the 8 year national load. Each time series has a binary classification that identifies whether it is from a business day (label A) or a weekend (label B)—Table 1. Public holidays were separated.

Figure 3 shows us the training dataset. Note that demand throughout the daylight period is smaller on weekends, and the morning ascending occurs later and demand peaks appear at different times in the day. On the other hand, on business days, there is a demand break around lunch. We ask whether the discovery algorithm will find the best shapelet(s) to distinguish those time series.

Table 1. The number of time series in the datasets for a business day (label A) and a weekend (label B).

Dataset	Label A	Label B	Total
Train	40	40	80
Test	2014	772	2786

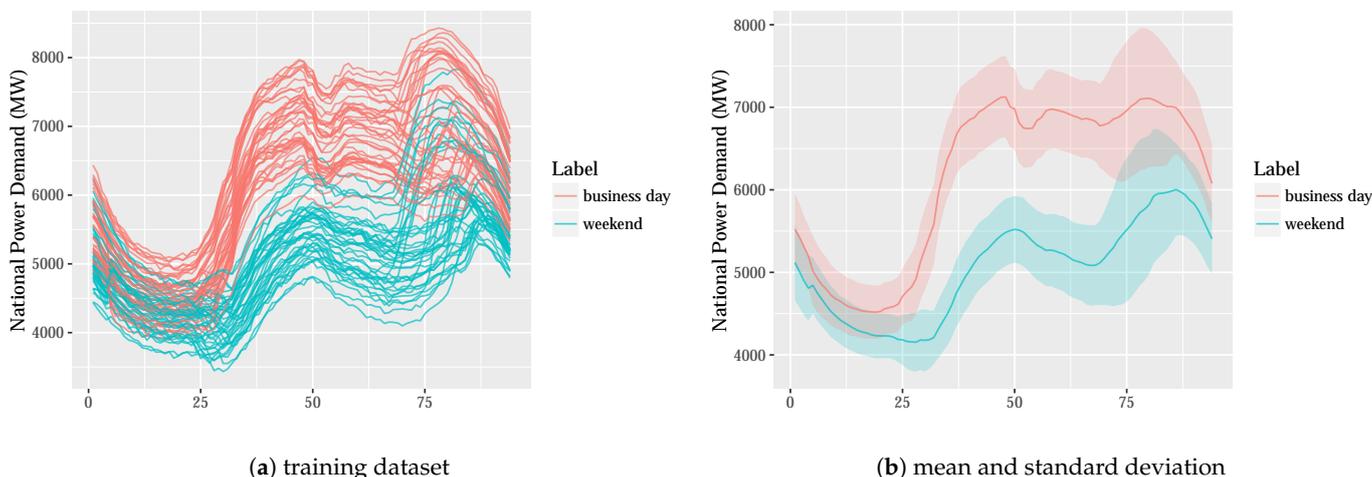


Figure 3. The complete training dataset with 40 time series of each class.

For this scenario, the shapelet discovery algorithm is parameterized with a minimum and maximum shapelet length of 20 and 35 positions. This corresponds to shapelets between 5 and 8.75 h long, a length that we have assumed to be enough to find the shape of weekend consumption.

Figure 4 shows the decision tree obtained after the training phase. One shapelet proved to be enough to decide whether a time series T is from a weekend or not. This results in an information gain of value 1.0 in the root node. Thus, during classification, all normalized subsequences of an arbitrary time series T are compared with the shapelet I . If just one normalized subsequence exists whose distance to the shapelet is less than 1.1617, then the time series is classified as label A, a business day.

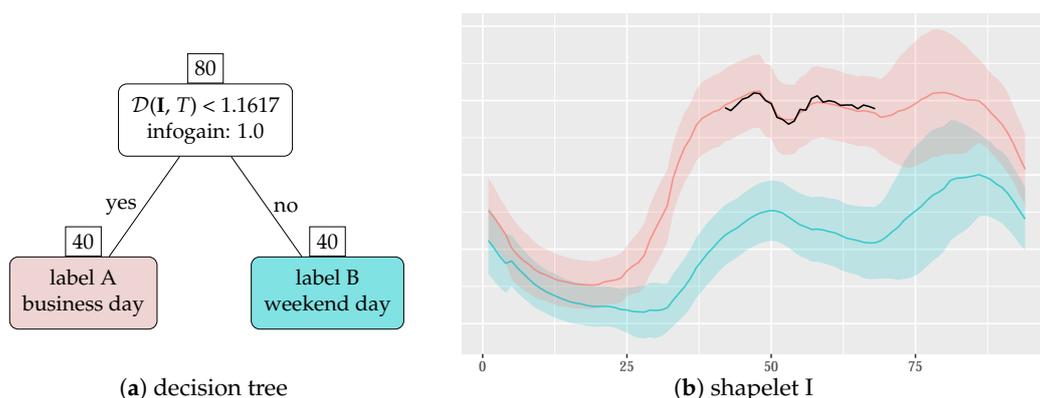


Figure 4. The decision tree classifier for the weekend problem. During the classification of an arbitrary time series T , its normalized subsequences are compared to the shapelet I (the black line). It is classified as label A if one of its normalized subsequences is *similar* to the shapelet. Note that the shapelet I was positioned according to where it was found, just to match the *similarities*. The shapelet might have a different scale and does not correspond to the coordinate axis.

The resulting shapelet is 27 points long, and it is positioned in the middle of the time series, i.e., in the middle of the day. As one can see, the shapelet has grasped the demand

break towards lunch—a trend which does not exist in the weekend time series. Note also that amplitude differences among the consumption time series, due to the effect of temperature or other human behavior, have not affected the accuracy of the decision tree.

Finally, the decision tree was evaluated with the testing dataset and it was found to perform with an accuracy of 96.77%.

3.2. Early Monday Morning

The aim of the second classification problem is to determine whether the following result can be found using the shapelets technique. There are several publications about short-term energy forecasting that use the day of the week as an explanatory variable following this idea: Mondays are different from the rest of the business days. They detach Monday into an individual category followed by the weekend category and the rest of the days of week as another category. The rationale behind this approach is the existence of differences between the consumption during Monday's first hours when compared with the other days, because Monday dawn comes from the end of the weekend which presents a different curve pattern as seen in the previous section. Is the shapelet algorithm able to discern the pattern presented in the Monday load curve from other business days and weekends?

The input data are similar to the previous scenario—time series of national demand and 24 hours long—but, this time, the target variable has three labels: Monday (label A), weekend (label B), and other business days (label C)—Table 2. Public holidays were separated from the dataset.

Table 2. The number of time series on the datasets from the category Monday (label A), weekends (label B), and other business days (label C).

Dataset	Label A	Label B	Label C	Total
Train	100	100	100	300
Test	309	712	1545	2566

From the training dataset—Figure 5b—the weekend consumption pattern is distinguishable from the other days, as we saw in the previous section. However, the Monday pattern is very similar, except during the early morning. Monday's mean demand follows a different pattern when compared with the other business days: the very first hour follows the same pattern, but the time series dives further until the first daylight hours.

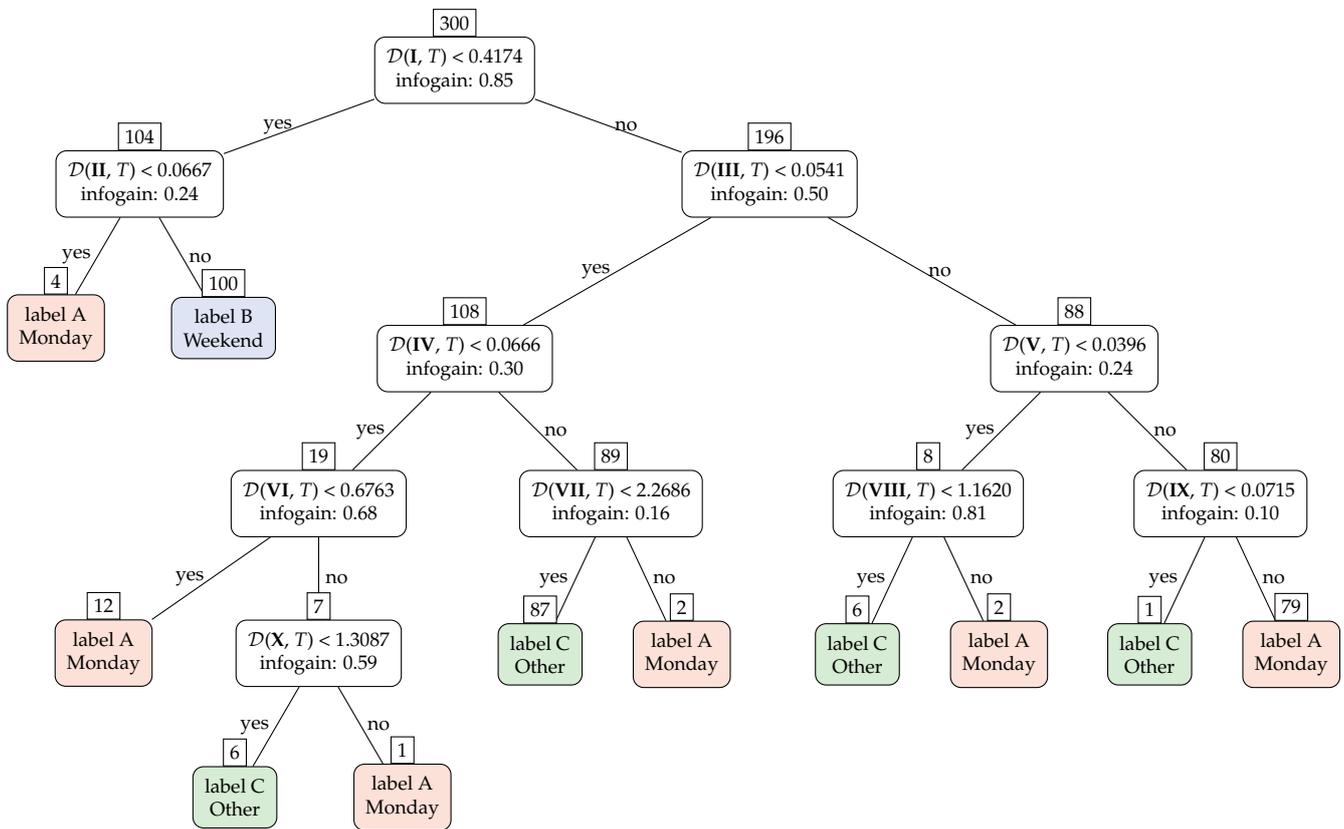
The resulting non-pruned decision tree—Figure 5a—has two important nodes that roughly split the training dataset into the three labels: the ones that use shapelets I and III. Actually, these are the nodes with the highest information gain multiplied by the number of time series to split (the number that appears above the node). In Figure 5c, the weight of the shapelet's line reflects its importance.

Therefore, shapelet I splits the weekend time series from the dataset (the left most branch of the decision tree). Shapelets III and V, which occur in the early morning, capture a pattern that favors class C (the opposite side of the most right branch of the decision tree). Meanwhile, the shapelet II, which also occurs in the early morning, makes a final split between the weekend time series and the Monday time series.

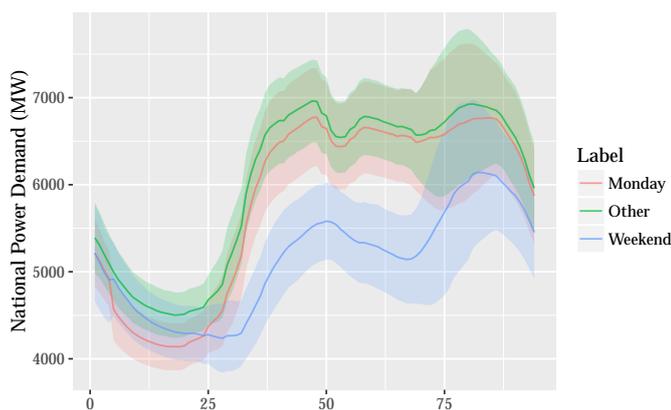
The non-pruned decision tree uses other similar shapelets (from VI to X), but they are not easily interpretable. However, even if those shapelets were ignored, i.e., the decision tree was pruned, the model would get right 278 times of 300 items (94%) from the training dataset.

Finally, the testing dataset was used to assess the accuracy of this decision tree. Despite its unpruned nodes and potential overfitting, the decision tree has a high accuracy of 88.74%. Evidently, there is a rationale to split Monday from the rest of business days when forecasting energy time series: the Monday demand is different from the other days and the shapelets discovery algorithm was able to find the pattern.

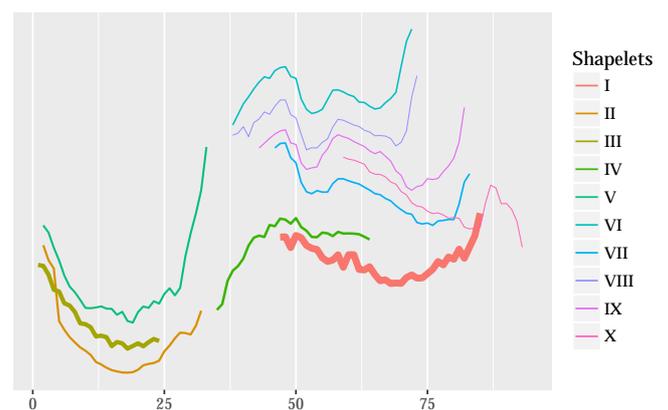
Furthermore, there are publications that break the business day, not into two categories but into three categories: (i) Monday, (ii) Friday, and (iii) Tuesday, Wednesday, and Thursday. We have tested this scenario, but the algorithm was not able to find a good shapelet that could accurately split time series. Indeed, the Friday time series shape is similar to Tuesday's, Wednesday's, and Thursday's. In that scenario, the same method resulted in a low accuracy of 58%.



(a) decision tree



(b) mean and standard deviation of training dataset



(c) shapelets

Figure 5. The decision tree classifier for the early morning pattern on Monday morning. (a) represents the non-pruned decision tree to classify time series as Monday (label A), weekend (label B) or other business day (label C) using the shapelets exhibited in (c). The line weight of the shapelets is proportional to its importance, i.e., the information gain times the number of time series splits in each node. (b) shows the mean and standard deviation of the training dataset with 100 time series of each class.

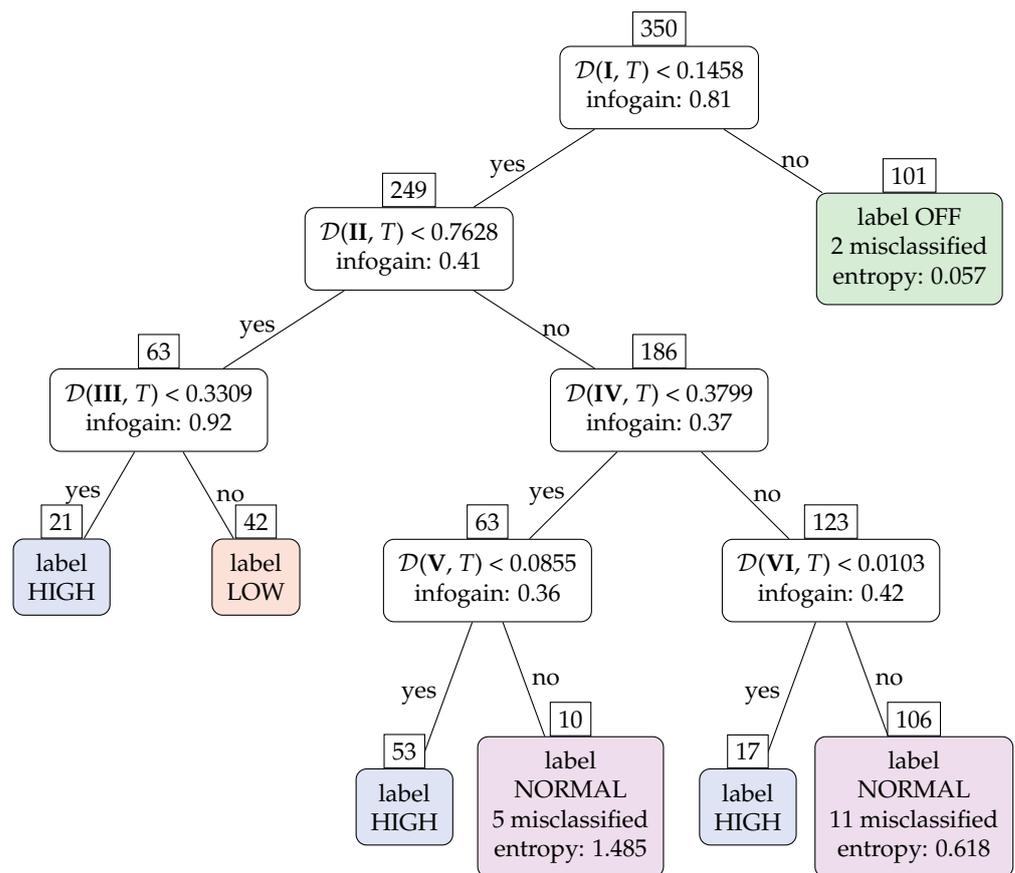
3.3. Load Dynamics in Substations

Power substations have the function of transforming energy from very high voltage (VHV) to high voltage (HV), or from high voltage to medium voltage (MV) and each usually feeds thousands of clients. As a redundant system, the power grid can manage the shutdown of a particular substation, whether due to scheduled maintenance or failure. In this situation, the load of a substation is shifted to another and thus holds out the energy flow without any significant outage. Obviously, the load time series reflects these maneuvers. While one substation reflects a very high load peak, the observations of the other reveal a stationary no-load state, see Figure 2.

Rebuilding the historical state of this dynamic grid at a particular moment in the past can be a large computational problem due to the enormous amount of assets that change states. However, this information is important to energy forecasting tasks because of the high impact on loads due to load shifts.

In this third classification problem, we have the goal of classifying whether a substation is in one of four states looking only for the load series. Input data are a set of time series extracted from three years of historical data through a 96-point sliding window (24 h). Each time series has a label that identifies the state of the substation at the 96th point—Table 3.

Figure 6 shows the decision tree obtained after the training phase using the first dataset. Shapelets II and IV discovered the load shifts that occurred when the grid system maneuvers the substation.



(a) decision tree

Figure 6. Cont.

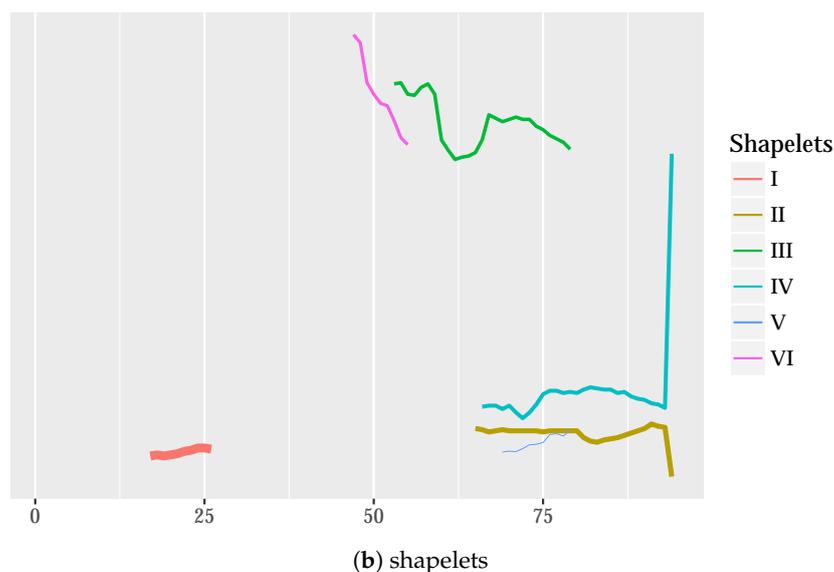


Figure 6. The decision tree classifies the state of specific substations. (a) represents the classifier to reason the state of specific substations among four possible classes—OFF, LOW, NORMAL, or HIGH—using the shapelets exhibited in (b).

Table 3. The number of time series in the datasets for the states OFF, LOW, NORMAL, and HIGH. The first dataset has an approximate distribution along the labels when comparing the training and testing dataset. The second includes all available time series.

Dataset		Label OFF	Label LOW	Label NORMAL	Label HIGH	Total
1	Train	100	50	100	100	350
	Test	400	37	400	400	1237
2	Train	150	50	150	150	500
	Test	7893	37	121,139	787	129,856

The decision tree used on the testing dataset was the non-pruned. In conclusion, the decision tree calculated with the first dataset has a high accuracy of 90%, and 83% for the decision tree using the second dataset. Furthermore, we have made experiments using sliding windows of 9 points long, SAX preprocessing of time series, fixing a state switch three points before the end of the series, or modifying the minimum and maximum lengths of candidate shapelets. However, the accuracy of the decision trees was lower, between 48% and 77%.

3.4. Type of Power Consumption

The curve of power consumption over a day depends on the purpose for which the energy is consumed. In this fourth classification problem, our objective is to classify whether power consumption serves one of these five purposes: (a) household, (b) industry, (c) services, (d) utilities, and (e) transportation. The demand curves are collected by meters of secondary substations owned by a distribution company or a high-demand client. Is the shapelets algorithm able to build a classifier capable of identifying the type of energy consumption from just the daily load curves?

The consumption data therefore refers to several sets of households or individual high-demand clients, for example: hospitals, retails, banks, post offices, government services, train and subway systems, waste water treatment plants, water pumping, gas utilities, and glass, furniture, rubber, plastics, porcelain, textiles, coating, and other types of industry. Table 4 shows us the number of meters available by class and the number of time series extracted from them. Note that the training data were obtained from meters other than the

ones used for testing. Therefore, the problem is challenging and difficult to classify, but this may be due to the low number of meters available. Additionally, the selected days were randomly chosen with the proviso that they were business days.

Table 4. The data available from 28 smart meters were separated into two subsets. From each of the 15 training meters, nine daily time series were randomly extracted to build up the training dataset. From each of the 13 testing meters, 19 daily time series were selected to make the testing dataset.

		Household	Industry	Services	Utilities	Transportation	Total
Meters	Train	5	3	4	2	1	15
	Test	4	3	3	2	1	13
Dataset	Train	45	27	36	18	9	135
	Test	37	57	57	38	18	207

The resulting decision tree—Figure 7—is fairly simple. Four shapelets are almost enough to classify the training dataset using five labels.

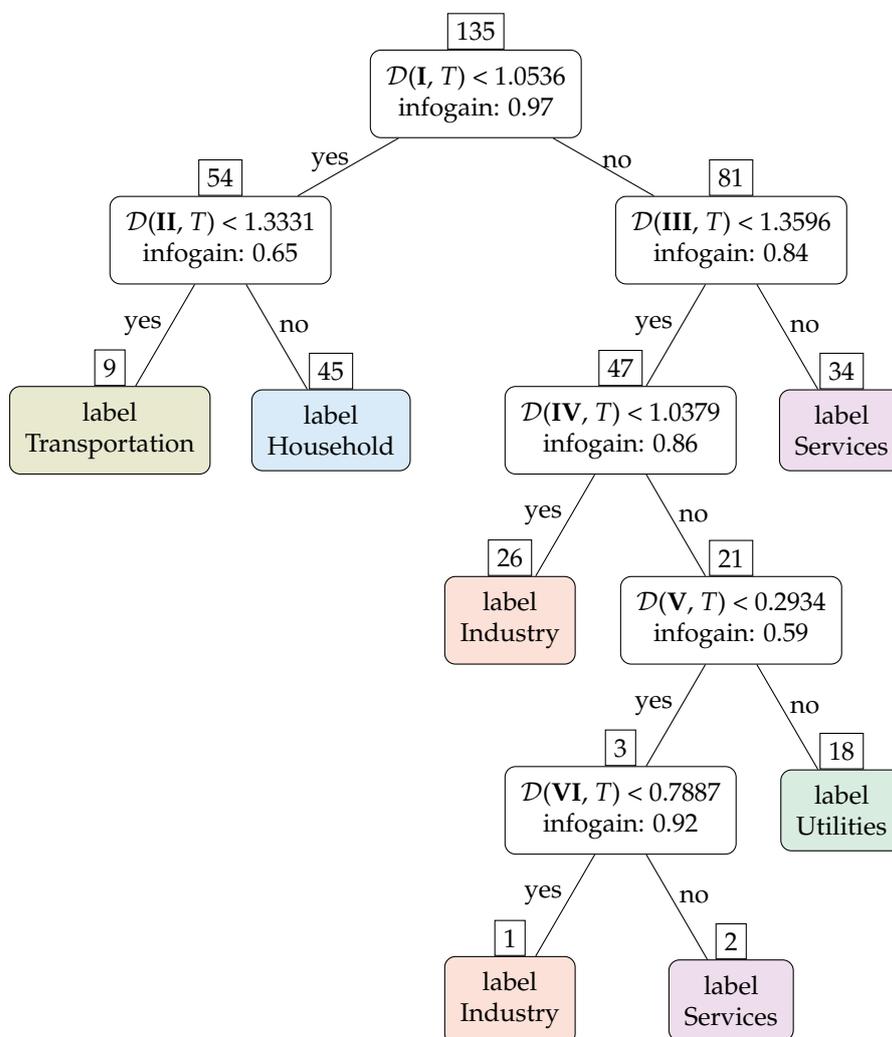


Figure 7. The decision tree classifies the type of power consumption.

The shapelet I splits the dataset into two groups: one for household and transportation consumption, and the other for services, utilities, and industry. As can be seen in Figure 8, this shapelet captures a peak in the evening as a pattern that favors “household” and “transportation” classes. Furthermore, the shapelet II splits these two labels—Figure 9a.

Since rush hour occurs twice a day, just before and after traditional work hours, it is natural that electric transportation consumes more energy during peak commuting hours.

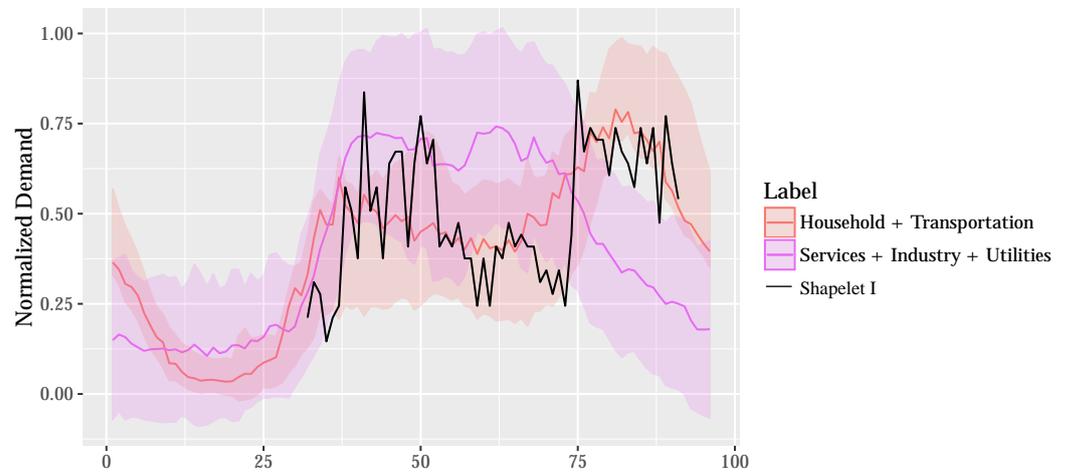


Figure 8. Normalized demand of two groups detached by shapelet I.

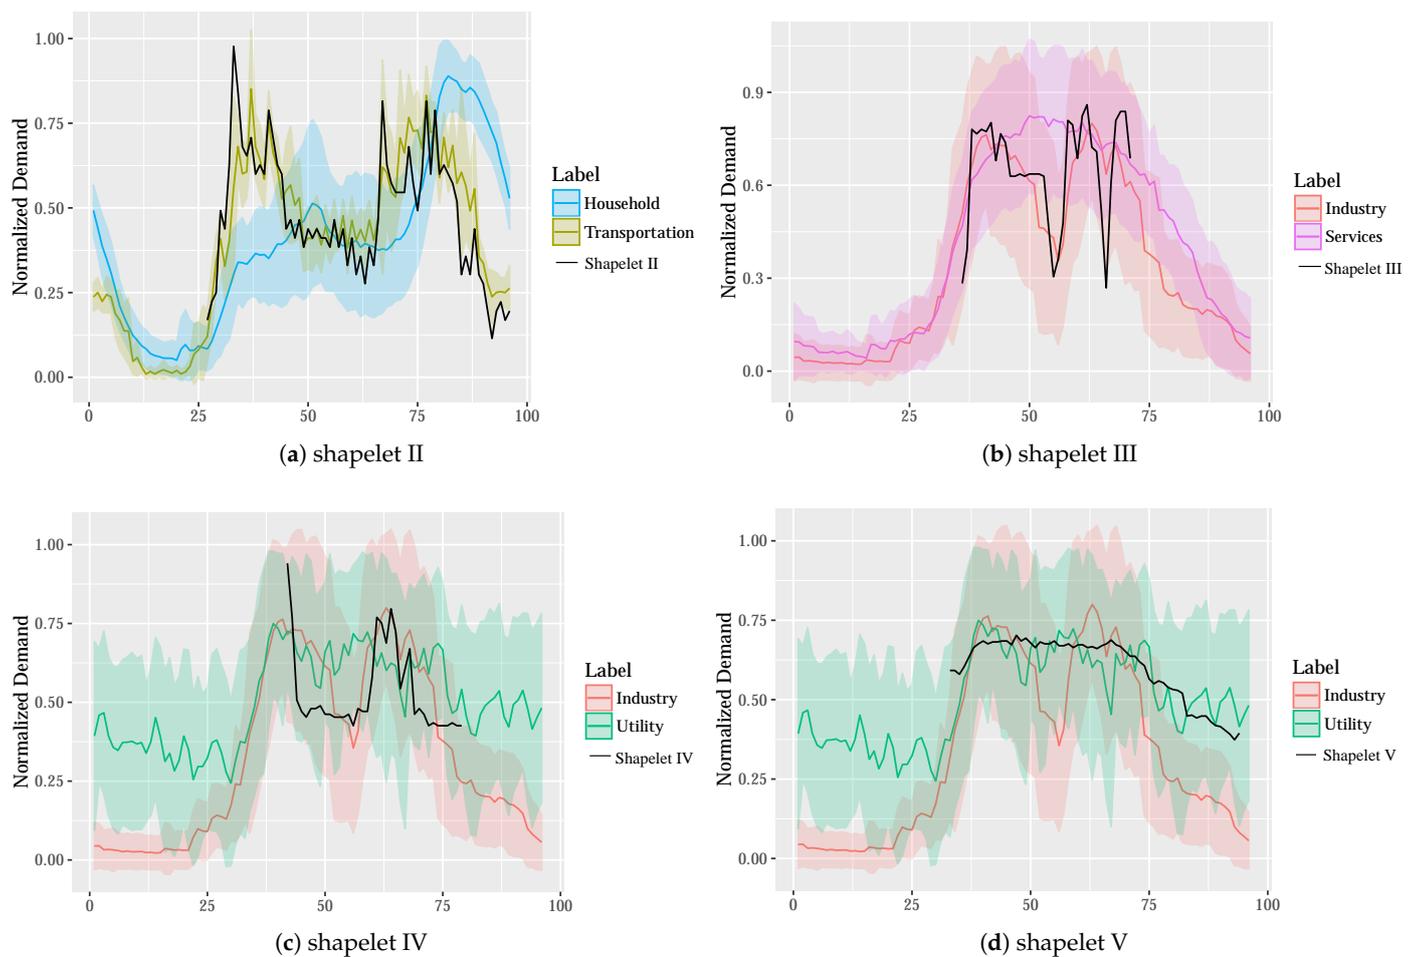


Figure 9. Normalized demand of classes detached by shapelets. The shapelet II detaches “transportation” from “household” through the demand patterns in the morning and evening. The shapelet III separates “industry” from “services” by a demand break around lunch time. The shapelet IV and V splits the time series into two groups: “industry” and “utility”.

On the other hand, shapelet III separates “industry” from the “services” demand series. In addition to the rise in the demand of both areas during the morning, and the decrease in the afternoon/evening, there is a break around lunch, which occurs only for industrial demand. Note that the descent phase of the demand for “services” extends into the evening.

Finally, the energy demand by industry and utilities presents different patterns. Utilities demand has a smoother fall in demand during the night compared to industry, and it does not get a break in demand towards lunch.

In conclusion, shapelets were able to identify the main patterns for each type of consumption, and they are quite interpretable from a domain knowledge point of view. However, the decision tree was evaluated against the training dataset and achieves only 60% accuracy. The low number of meters available may have biased the performance of shapelets in this scenario.

4. Discussion and Conclusions

In this study, the shapelet technique was applied to power load time series collected from three different hierarchical levels of the grid (national, primary power substations, and secondary power substations). Four use cases were defined considering the potential application of shapelets to answer a specific goal. This study is the first to analyze the impact of this technique in more than one use case in the energy domain, following the same methodology and at more than one grid level, while others focus on a single aspect and at a specific grid level. That was made possible by the access to real private data collected in the Portuguese smart grid and the real analytics challenges that were posed, establishing the hypothesis that power load curves alone might have inherent information capable of solving the challenges.

Shapelets have been shown to be useful for the study of pattern extraction and knowledge to understand human activities as energy consumers. This reinforces the importance of data-driven services in the energy sector that can monetize (extract additional value from) data collected within smart grids. The inability to recognize the value of the data in addition to the siloed application in which the data are collected was one of the barriers identified by the reviewers, as outlined in Section 1. Electrification (heating and mobility as examples) and integration of more variable renewable energy (VRE), towards the ambition of 100% renewable electricity, also depend on knowing the purpose of energy consumption and whether there is flexibility on the demand side and capacity of the grid to support those two goals. The fourth use case demonstrates the possibility to classify the purpose of energy consumption and which repeatable patterns the load follows beyond the general patterns highlighted in many studies at the system level. The use case in this paper is more meticulous in the study of consumption at the primary and secondary substation level.

In particular, shapelets have shown adequate accuracy in all four use cases:

- Weekends—A shapelet was enough to distinguish the daily load time series of weekends from the business days with an accuracy of 97% (Section 3.1). The interpretable shapelet is placed in the middle of the day, and grasps the demand break during lunch, a trend that exists on business days but not on weekends.
- Early Monday Morning—The resulting decision tree has a high accuracy of 89% (Section 3.2) capable of classifying the daily load time series as weekends, Mondays, and other working days. The two most important shapelets are: (i) a similar shapelet as the previous case with the same goal, the weekend group time series, and (ii) a shapelet that occurs in the early morning on Monday morning, with lower demand—a deeper valley in the curve—when compared with other business day time series. Compared with [16], we used data from the system level (instead of 33 buildings at the low-voltage level) with the same number of quarter-hourly training time series (namely 300 compared to 272) and a higher number of testing time series (namely, 2566 compared to 636). Although they have slightly different goals, both studies reach the same degree of accuracy (89%).

- **Load Dynamics in Substations**—The resulting decision tree has a high accuracy between 83% and 90% for the two testing datasets (Section 3.3), capable of identifying the substation load state, due to scheduled maintenance, failures or maneuvers, even that the load curve presents a high amplitude change by a recent maneuver.
- **Type of Power Consumption**—The resulting decision tree has an accuracy of 60% (Section 3.4), capable of classifying daily consumption according to the type of consumption without any other data beyond the own load curve. The amount of data available for testing may have biased the performance of shapelets in this use case. Once again, the shapelets are interpretable and reinforce the tacit and empirical knowledge practitioners might have about the consumption patterns present in different types of consumption. Compared with [17], we used data from 28 substations (instead of 10), with load curves detailed by 96 datapoints—quarter-hourly (instead of 24 datapoints—hourly) to classify them into one of the five classes (instead of two classes). Our accuracy results reflect the classification ability of the features extracted only by shapelets, while the other study combines the shapelet features with other features to build the classifier, preventing one from pondering shapelets alone).

Thus, this study extends the literature on shapelets applied to the energy domain. At the same time, the study has some limitations, including relying solely on the shapelet technique. However, such a decision was made because of the desire to study the potential of this technique and to reinforce that the power load time series has value beyond the primary functions for which those data are collected. The second limitation is that one relies solely on a database of Portugal power load time series. This is due to the lack of real public data collected within smart grids.

Moreover, the resulting shapelets, and the decision trees that utilize shapelets while node splitting, are interpretable. Interpretability is one of the crucial characteristics in the context of electricity companies in a high-regulated business in Europe, which seeks to regulate AI across several applications, particularly models that act on or support decisions over critical infrastructure. Interpretability eases the application of data-driven services once it must be approved by managers, understandable by grid operators, and defensible before the regulator. Additionally, the interpretability of shapelets increases the knowledge of domain practitioners and, cyclically, the fact that their domain knowledge is “recognized” by the model reinforces its acceptance and adoption. Shapelets fall into the category of techniques that deliver interpretable results.

We would like to reinforce the imperative to grid operators to boost their progress toward being more data-centric and provide load curves in higher detail. This promotes research and development on innovative technologies, and accelerates the adoption of new impactful energy models, such as the integration of more VRE, development of generation and demand side flexibility, integration of more decentralized energy resources (DER), electric vehicle charging and other electric vehicle models like vehicle-to-grid (V2G) and vehicle-to-everything (V2X), microgrids, energy communities, and peer-to-peer energy (P2P energy). There are questions as to how these various energy models work together toward the energy transition and how they would be integrated into an existing grid that wants to be resilient and sustainable, both economically and environmentally.

Author Contributions: Conceptualization, methodology, software, investigation, data curation, writing—original draft preparation, M.G.P.; writing—review and editing, M.G.P. and A.P.F.; supervision, A.P.F. and S.C.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Publicly available datasets and private third party datasets were analyzed in this study. The public available dataset (national electricity) can be found here: <https://datahub.ren.pt/> (REN Data Hub website). Restrictions apply to the availability of primary and secondary substations data. Data was obtained from E-Redes and are available from the author M.G.P. with the permission of E-Redes.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Terlouw, T.; AlSkaif, T.; Bauer, C.; van Sark, W. Multi-objective optimization of energy arbitrage in community energy storage systems using different battery technologies. *Appl. Energy* **2019**, *239*, 356–372. [[CrossRef](#)]
2. Schill, W.P. Electricity Storage and the Renewable Energy Transition. *Joule* **2020**, *4*, 2059–2064. [[CrossRef](#)]
3. Sanguesa, J.A.; Torres-Sanz, V.; Garrido, P.; Martinez, F.J.; Marquez-Barja, J.M. A review on electric vehicles: Technologies and challenges. *Smart Cities* **2021**, *4*, 372–404. [[CrossRef](#)]
4. Diahovchenko, I.; Kolcun, M.; Čonka, Z.; Savkiv, V.; Mykhailiushyn, R. Progress and challenges in smart grids: Distributed generation, smart metering, energy storage and smart loads. *Iran. J. Sci. Technol. Trans. Electr. Eng.* **2020**, *44*, 1319–1333. [[CrossRef](#)]
5. Psara, K.; Papadimitriou, C.; Efstratiadi, M.; Tsakanikas, S.; Papadopoulos, P.; Tobin, P. European Energy Regulatory, Socioeconomic, and Organizational Aspects: An Analysis of Barriers Related to Data-Driven Services across Electricity Sectors. *Energies* **2022**, *15*, 2197. [[CrossRef](#)]
6. Haben, S.; Arora, S.; Giasemidis, G.; Voss, M.; Greetham, D.V. Review of low voltage load forecasting: Methods, applications, and recommendations. *Appl. Energy* **2021**, *304*, 117798. [[CrossRef](#)]
7. Lim, B.; Zohren, S. Time-series forecasting with deep learning: A survey. *Philos. Trans. R. Soc. A* **2021**, *379*, 20200209. [[CrossRef](#)] [[PubMed](#)]
8. Aghabozorgi, S.; Shirkhorshidi, A.S.; Wah, T.Y. Time-series clustering—A decade review. *Inf. Syst.* **2015**, *53*, 16–38. [[CrossRef](#)]
9. Ismail Fawaz, H.; Forestier, G.; Weber, J.; Idoumghar, L.; Muller, P.A. Deep learning for time series classification: A review. *Data Min. Knowl. Discov.* **2019**, *33*, 917–963. [[CrossRef](#)]
10. Baydogan, M.G.; Runger, G. Learning a symbolic representation for multivariate time series classification. *Data Min. Knowl. Discov.* **2015**, *29*, 400–422. [[CrossRef](#)]
11. Yeh, C.C.M.; Zhu, Y.; Ulanova, L.; Begum, N.; Ding, Y.; Dau, H.A.; Zimmerman, Z.; Silva, D.F.; Mueen, A.; Keogh, E. Time series joins, motifs, discords and shapelets: A unifying view that exploits the matrix profile. *Data Min. Knowl. Discov.* **2018**, *32*, 83–123. [[CrossRef](#)]
12. Zhu, Y.; Imamura, M.; Nikovski, D.; Keogh, E. Matrix profile VII: Time series chains: A new primitive for time series data mining. In Proceedings of the 2017 IEEE International Conference on Data Mining (ICDM), New Orleans, LA, USA, 18–21 November 2017; pp. 695–704.
13. Imani, S.; Madrid, F.; Ding, W.; Crouter, S.E.; Keogh, E. Introducing time series snippets: A new primitive for summarizing long time series. *Data Min. Knowl. Discov.* **2020**, *34*, 1713–1743. [[CrossRef](#)]
14. Gharghabi, S.; Yeh, C.C.M.; Ding, Y.; Ding, W.; Hibbing, P.; LaMunion, S.; Kaplan, A.; Crouter, S.E.; Keogh, E. Domain agnostic online semantic segmentation for multi-dimensional time series. *Data Min. Knowl. Discov.* **2019**, *33*, 96–130. [[CrossRef](#)] [[PubMed](#)]
15. Hasan, M.M.; Chowdhury, D.; Khan, M.Z.R. Non-Intrusive Load Monitoring Using Current Shapelets. *Appl. Sci.* **2019**, *9*, 5363. [[CrossRef](#)]
16. Butunoi, B.P.; Frincu, M. Shapelet based classification of customer consumption patterns. In Proceedings of the 2017 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), Torino, Italy, 26–29 September 2017; pp. 1–6.
17. Zhang, F.; Bales, C.; Fleyeh, H. Feature Augmentation of Classifiers Using Learning Time Series Shapelets Transformation for Night Setback Classification of District Heating Substations. *Adv. Civ. Eng.* **2021**, *2021*, 8887328. [[CrossRef](#)]
18. Zhu, L.; Lu, C.; Sun, Y. Time series shapelet classification based online short-term voltage stability assessment. *IEEE Trans. Power Syst.* **2015**, *31*, 1430–1439. [[CrossRef](#)]
19. Zhu, L.; Lu, C.; Dong, Z.Y.; Hong, C. Imbalance learning machine-based power system short-term voltage stability assessment. *IEEE Trans. Ind. Inform.* **2017**, *13*, 2533–2543. [[CrossRef](#)]
20. Zhu, L.; Hill, D.J. Networked Time Series Shapelet Learning for Power System Transient Stability Assessment. *IEEE Trans. Power Syst.* **2021**, *37*, 416–428. [[CrossRef](#)]
21. Zakaria, J.; Mueen, A.; Keogh, E. Clustering time series using unsupervised-shapelets. In Proceedings of the 2012 IEEE 12th International Conference on Data Mining, Brussels, Belgium, 10–13 November 2012; pp. 785–794.
22. Ye, L.; Keogh, E. Time series shapelets: A new primitive for data mining. In Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, 28 June–1 July 2009; pp. 947–956.
23. Mueen, A.; Keogh, E.; Young, N. Logical-shapelets: An expressive primitive for time series classification. In Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, CA, USA, 21–24 August 2011; pp. 1154–1162.
24. Rakthanmanon, T.; Keogh, E. Fast shapelets: A scalable algorithm for discovering time series shapelets. In Proceedings of the Thirteenth SIAM Conference on Data Mining (SDM), Austin, TX, USA, 2–4 May 2013; pp. 668–676.

25. Grabocka, J.; Wistuba, M.; Schmidt-Thieme, L. Fast classification of univariate and multivariate time series through shapelet discovery. *Knowl. Inf. Syst.* **2015**, *49*, 429–454. [[CrossRef](#)]
26. Cheng, Z.; Yang, Y.; Wang, W.; Hu, W.; Zhuang, Y.; Song, G. Time2graph: Revisiting time series modeling with dynamic shapelets. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 3617–3624.
27. Hills, J.; Lines, J.; Baranauskas, E.; Mapp, J.; Bagnall, A. Classification of time series by shapelet transformation. *Data Min. Knowl. Discov.* **2014**, *28*, 851–881. [[CrossRef](#)]
28. Lines, J.; Bagnall, A. Alternative quality measures for time series shapelets. In *Intelligent Data Engineering and Automated Learning-IDEAL 2012*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 475–483.