




Article

Short-Term Power Demand Forecasting for Diverse Consumer Types Using Customized Machine Learning Approaches

Asier Diaz-Iglesias ^{*}, Xabier Belaunzaran  and Ane M. Florez-Tapia 

Vicomtech Foundation, Basque Research and Technology Alliance (BRTA), 20009 Donostia-San Sebastián, Spain; xbelauzaran@vicomtech.org (X.B.); amflorez@vicomtech.org (A.M.F.-T.)

* Correspondence: adiaz@vicomtech.org

Abstract

Ensuring grid stability in the transition to renewable energy sources requires accurate power demand forecasting. This study addresses the need for precise forecasting by differentiating among industrial, commercial, and residential consumers through customer clusterisation, tailoring the forecasting models to capture the unique consumption patterns of each group. Feature selection incorporated temporal, socio-economic, and weather-related data obtained from the Copernicus Earth Observation (EO) program. A variety of AI and machine learning algorithms for short-term load forecasting (STLF) and very-short-term load forecasting (VSTLF) are explored and compared, determining the most effective approaches. With all that, the main contribution of this work are the new forecasting approaches proposed, which have demonstrated superior performance compared to simpler models, both for STLF and VSTLF, highlighting the importance of customized forecasting strategies for different consumer groups and demonstrating the impact of incorporating detailed weather data on forecasting accuracy. These advancements contribute to more reliable power demand predictions, with our novel forecasting approaches reducing the Mean Absolute Percentage Error (MAPE) by up to 1–3% for industrial and 1–10% for commercial consumers compared to baseline models, thereby supporting grid stability.

Keywords: load forecasting; machine learning; smart grid; clustering



Academic Editors: Álvaro Gutiérrez and Tek-Tjing Lie

Received: 16 July 2025

Revised: 9 September 2025

Accepted: 7 October 2025

Published: 10 October 2025

Citation: Diaz-Iglesias, A.; Belaunzaran, X.; Florez-Tapia, A.M. Short-Term Power Demand Forecasting for Diverse Consumer Types Using Customized Machine Learning Approaches. *Energies* **2025**, *18*, 5332. <https://doi.org/10.3390/en18205332>

Copyright: © 2025 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 European Union's ambition to become a climate-neutral economy by 2050 [1] is deeply intertwined with the need for effective grid management and accurate power demand forecasting. As the EU accelerates its transition towards renewable energy sources (RES), the variability and intermittency of these resources present significant challenges for maintaining a stable and reliable power grid. Accurate demand forecasting (often referred to as electric load forecasting) becomes crucial in this context, as it enables grid operators to anticipate fluctuations in power demand and supply, ensuring that the balance is maintained despite the inherent variability of RES. Effective grid management, supported by precise forecasting, is vital to preventing grid disruptions like imbalances or blackouts, making it essential to the EU's climate goals and a secure green energy transition.

Proper planning and effective applications of power demand forecasting require particular “forecasting intervals,” or “lead time” [2,3]. Based on the lead time, load forecasting can be categorized into four distinct types [4,5]:

- Very-short-term load forecasting (VSTLF): applicable for real-time control with a prediction period ranging from minutes to 1 h ahead.

- Short-term load forecasting (STLF): it covers forecasting within 1 h to 7 days.
- Medium-term load forecasting (MTLF): it covers forecasting ranging from a few months up to 2–3 years.
- Long-term load forecasting (LTLF): it covers forecasting typically ranging from 1 year to 10 years or more.

While LTLF presents the most difficult challenge in terms of forecasting [6], VSTLF and STLF are the most critical ones [7,8]. They have only become viable in recent years due to the widespread adoption of advanced metering infrastructure (AMI) and connected sensors capable of capturing electrical consumption at a high level of granularity [9,10].

Given the substantial increase in the use of AMIs, this paper focuses on VSTLF and STLF of individual consumers located in two different Spanish regions with varying climate conditions. The consumers are categorized into three types: industrial, commercial, and residential [11]. Each consumer type has unique temporal characteristics that must be taken into account when applying time-series forecasting algorithms; e.g., while commercial and industrial consumers typically exhibit pronounced seasonality patterns due to the workweek, the load of residential consumers is usually highly variable and irregular.

In addition to the temporal characteristics, socio-economic factors and the climate influence on each consumer need to be considered. Socio-economic factors significantly influence demand forecasting across different regions and consumer segments by shaping consumption patterns and energy usage [12,13]. Accurate forecasting models must account for these dynamics to reflect the variability in power usage effectively.

Weather also plays a crucial part in power demand forecasting, directly influencing energy consumption patterns. Weather variables can significantly impact the need for heating, cooling, and lighting, leading to variations in demand; therefore, accurate forecasting models must account for these variables to predict energy usage effectively [14]. Hence, spatially representative weather data are used to plan, design, size, construct, and manage energy systems for performance analysis and forecasting to enhance system efficiency [15]. Specifically, this article proposes leveraging the services of the Copernicus Earth Observation (EO) programme and the European Centre for Medium-Range Weather Forecasts (ECMWF), obtaining data from the ERA5 dataset [16], which provides comprehensive, high-resolution, and globally consistent datasets. Data from national weather stations was also used to enhance prediction accuracy.

Load forecasting models have been categorized into statistical, intelligent-computing-based, and hybrid models [17]. In recent years, Artificial Intelligence (AI) methods have emerged as powerful tools for improving the accuracy of VSTLF and STLF. These approaches include commonly used support vector machines (SVMs) or artificial neural networks (ANNs). More recently, emphasis has been placed on deep learning techniques such as convolutional neural networks (CNNs), long short-term memory (LSTM), or transformers [18–20]. While deep learning has shown outstanding performance, it requires large amounts of data and careful preprocessing of the data. This paper will therefore examine different AI-based approaches and elucidate a process of developing efficient and robust forecasting models for different types of customers with a focus on the 15 min and 24 h time horizons to achieve the goals set in terms of accuracy.

Our study, conducted within the framework of a Horizon Europe project, offers several significant contributions:

1. We propose a comprehensive and easily implementable framework for DSOs to deploy accurate STLF and VSTLF in a scalable manner.
2. We demonstrate that a simple, rule-based clustering technique is sufficient for effective consumer classification, achieving 85% accuracy without requiring complex algorithms or private consumer data.

3. We present a model-agnostic feature importance analysis, identifying and prioritizing the key drivers of power demand (e.g., calendar data, holidays, weather) unique to each consumer type.
4. We introduce novel, tailored forecasting approaches (model fusion for industrial/commercial, hybrid baseline for residential) that significantly enhance robustness and accuracy over standard single-model implementations.
5. We provide a practical comparison of various AI models, identifying LightGBM as a consistently high-performer, and validate all models in a realistic, production-like scenario with rolling forecasts.

The structure of this paper is as follows: First, a comprehensive analysis of the current state of the art is conducted. Next, the proposed methodology is outlined, followed by the presentation of the case study and a discussion of the results. Finally, the paper concludes with a discussion of future work and key conclusions.

2. State of the Art

This section presents an overview of the latest insights regarding VSTLF and STLF. Additionally, an overview of feature selection and customer classification is provided.

In certain applications, including this one under study, the process of load clustering prior to the forecasting step is to be considered. It is essential when tailoring forecasting models to the unique consumption patterns of different consumer groups, by accurately categorizing consumers, forecasters can account for the distinct factors that influence each group. The algorithms regarding load clustering can be divided into five groups: hierarchical, model-based, density-based, grid-based, and partition-based clustering [21]. Of these, the most popular algorithm is K-means, a type of partition-based algorithm, due to its outstanding computational efficiency on large-scale and high-dimensional load datasets [9]. A variant of this algorithm has recently been employed to categorise loads, as exemplified by the K-shape classification in [22]. It is also noteworthy that model-based algorithms have gained prominence, with the emergence of novel approaches such as a Gaussian Mixture Model (GMM) variation in [23] and a variational autoencoder proposed in [24]. In some cases, the performance of these models is enhanced by incorporating time-related features, which are extracted by applying domain knowledge. The load patterns are then identified per these features, as illustrated in [25,26].

A second line of research examines the feature selection process before carrying out the forecasting task, once the different consumers are clustered. The most commonly used technique for feature selection is Principal Component Analysis (PCA), a multivariate statistical analysis technique for data compression and feature extraction that can effectively remove linear correlations between data [27]. It is usually used to reduce dimensionality as in [28]. Another variety of techniques can be employed to calculate the influence of different variables on electric consumption; Ref. [29] employs a generalised additive model (GAM) to investigate the impact of various features on residential energy demand, identifying that month, day, and temperature are highly significant in predicting consumption. Another widely used method is the Pearson correlation coefficient, as demonstrated in [30,31], whereby the coefficient is employed to ascertain the impact of specific features on the electric load. An additional approach is illustrated in [32], whereby the coarse set of features is initially screened using the Maximal Information Coefficient (MIC), and subsequently, the fine set of key features affecting load forecasting is screened using implementations of Gradient Boosting algorithms, specifically the LightGBM and XGBoost methods, respectively. The corresponding fine set of features and historical loads are input into LightGBM and XGBoost with a robust prediction function for prediction, and the predicted value is employed to rectify the error and complete the load prediction. Other approaches consist

of model explainability techniques such as SHAP (SHapley Additive exPlanation). In [33], the technique is employed to reduce feature redundancy and overfitting, after which the selected features are leveraged to forecast day-ahead electricity prices.

Finally, regarding forecasting algorithms, there has recently been a notable shift towards the utilization focus of hybrid deep learning algorithms, including CNN-GRU [34], CNN-LSTM [35], and even GAN-enhanced models [36,37]. While these ANNs have become a prominent topic of discussion in the load forecasting literature over the last few years, they present significant challenges when applied to real-world STLF and VSTLF scenarios. This is largely due to the issue of model overfitting and the exponential increase in complexity associated with high-dimensional datasets [38]. Additionally, it is important to note that most deep learning models require a substantial number of experiments to identify an optimal configuration [39]. Therefore, while these hybrid deep learning approaches represent a significant area of research, they present considerable challenges for real-world STLF and VSTLF applications, including a high computational cost, a propensity for overfitting on smaller or noisier datasets, and a lack of transparency that can hinder adoption by industry practitioners who require interpretable and robust forecasts.

Consequently, there has been a gradual resurgence of interest in exploring more straightforward regression-based machine learning algorithms, which are less prone to these issues and often demonstrate superior efficacy. While evaluated, numerous alternative regressors have been demonstrated to exhibit superior efficacy in constructing STLF models compared to ANNs. For instance, the Support Vector Regressor (SVR) has been shown to outperform ANNs in [40]. In [41], the Multiple Linear Regression (MLR), ANNs, and SVR are subjected to a comparative analysis, wherein SVR emerges as the most prominent performer. In [42], a sliding window-based LightGBM model has been shown to outperform the popular LSTM deep learning algorithm.

An additional illustration of the efficacy of tree-based gradient boosting algorithms can be observed in [43], where the model is a combined method based on similar days matching and the XGBoost method for short-term load forecasting on holidays. This approach is proposed as a means of addressing the problem of load forecast errors on such days. One challenge associated with these gradient-boosting algorithms is the accurate specification of hyperparameters. This is addressed in [44] through the use of Bayesian Optimization, which is employed to optimise the hyperparameters for a LightGBM model prior to performing the forecasting.

3. Materials and Methods

In this section, we describe the developed load forecasting method in detail. The overarching framework of the load demand forecasting system is depicted in Figure 1.

This system begins by collecting a variety of input data sources, further discussed in Section 4, including historical energy consumption records provided by the stakeholders in the project, weather-related data obtained from the ERA5 dataset and regional meteorological stations, as well as socio-economic indicators and calendar-related data.

As the first step, the dataset undergoes a standard preprocessing pipeline. Outliers are identified and corrected using a Hampel filter. Missing values are imputed using the Multivariate Imputation by Chained Equations (MICE) algorithm to preserve the statistical properties of the dataset. Weather data power demand data is standardized to a range between -1 and 1 , while calendar features are categorically encoded for tree-based models and one-hot encoded for other models.

The forecasting process then involves classifying customers into distinct categories—residential, industrial, and commercial—based on their energy consumption features and the characteristics of their respective data. This classification is essential as each

customer category exhibits unique behavior and demand drivers that must be considered when forecasting load.

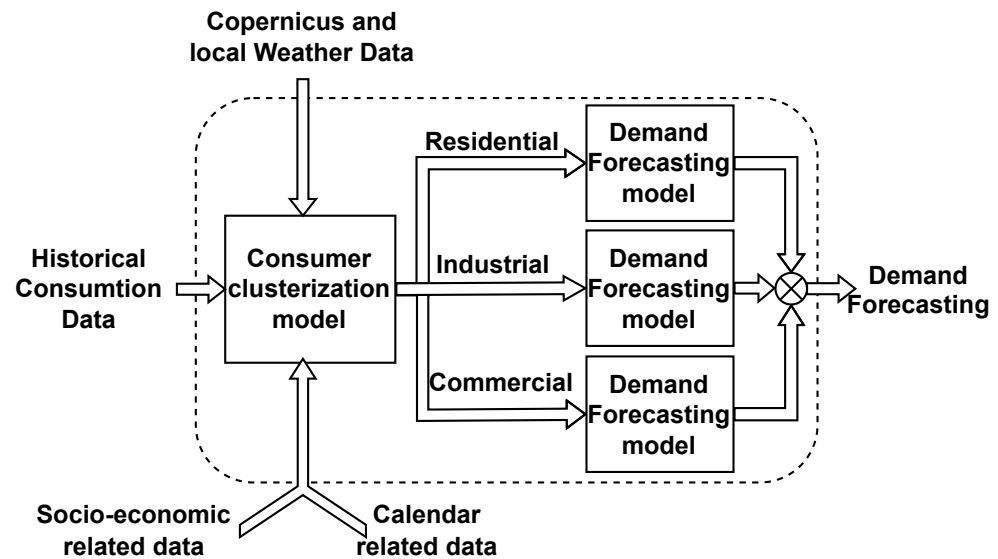


Figure 1. Framework of the load demand forecasting system. The dotted box separates the system (inside the box) from inputs and outputs (outside the box).

Once customers are classified, specific demand forecasting models are tailored and applied to each customer category. The model selection is critical, as it ensures that the most appropriate forecasting techniques are employed, considering each category's unique attributes and influencing factors. As a final step, the consumption forecasts for different customers at the same location are aggregated. By utilizing a targeted approach for each customer type, the system enhances the accuracy and reliability of the load forecasts, ultimately supporting better grid management and planning.

3.1. Customer Classification

One possible approach would be to implement a pattern recognition algorithm using time-domain features, such as the K-means clustering algorithm, as demonstrated in [25,26], to distinguish between different consumption patterns. Furthermore, if we had access to additional consumer characteristics such as household appliances data, consumer behaviour, or the size of the companies in the case of industrial and commercial customers, we could apply more sophisticated methods, as described in [45] or [46]. However, all the collected data is completely anonymous and only the customer type and locations are available as input data. Thus, having the ground truth, we developed a classifier based on descriptive characteristics and domain knowledge.

1. If the consumption during the holidays (C_H) is less than half the consumption during the working days (C_W), and the mean consumption on Saturdays (\hat{C}_6) is less than twice the mean consumption on Sundays (\hat{C}_7), the algorithm classifies the dataset as an *industrial* consumer:

$$\begin{aligned} C_H &< \frac{C_W}{2} \\ \hat{C}_6 &< 2\hat{C}_7 \end{aligned} \quad (1)$$

The first condition differentiates industrial and commercial consumers from residential ones, as they completely stop or reduce consumption to a minimum during holidays. The second condition separates the industrial consumers from the commercial ones, as the industrial consumers normally reduce their consumption on Saturdays too, in contrast to the commercial consumers, who tend to work on Saturdays.

- Once the industrial consumers are separated, a differentiation between the commercial and the residential consumers is sought. The dataset is classified as *commercial* if the standard deviation of mean consumption at different hours exceeds 0.1, and if mean consumption on holidays is lower than on workdays. The first condition identifies commercial datasets, which have consistent consumption peaks at specific hours, leading to a higher standard deviation compared to the more irregular residential consumption patterns. The second condition prevents misclassification of residential data that fits the first condition. Additionally, if mean consumption on Saturdays is more than 1.5 times that on Sundays, the dataset is also classified as *commercial*.

$$\begin{aligned} \text{std}(\hat{C}_i) &> 0.1 \\ \hat{C}_H &> \hat{C}_W \\ \hat{C}_6 &> 1.5\hat{C}_7 \end{aligned} \quad (2)$$

This condition is similar to the first applied rule, which compares the consumption during workdays and holidays, as well as the consumption during Saturdays and Sundays, which helps to not classify commercial data as residential.

- If the dataset does not fall into the two previous categories, i.e., does not meet the conditions established by Equations (1) and (2), it is classified as a *residential* consumer.

All of the thresholds present in Equations (1) and (2) were derived from exploratory data analysis and domain knowledge regarding typical consumption behaviour of each consumer type.

3.2. Model Training

This section aims to describe the process of designing the forecasting models for each consumer type, once they are classified; the reasoning behind model selection, parameter selection, and training process is also discussed. In Figure 2 the diagram of the proposed algorithm architecture is shown, where it can be seen that, based on historical consumption data and the most relevant features, further detailed in Section 3.2.1, and after being resampled adequately, each forecasting algorithm needs to perform two distinct forecasting tasks: one is a one-step forecast 15 min ahead in the future, and the other performs a prediction of the next 24 h at each time step of the testing data, i.e., hourly.

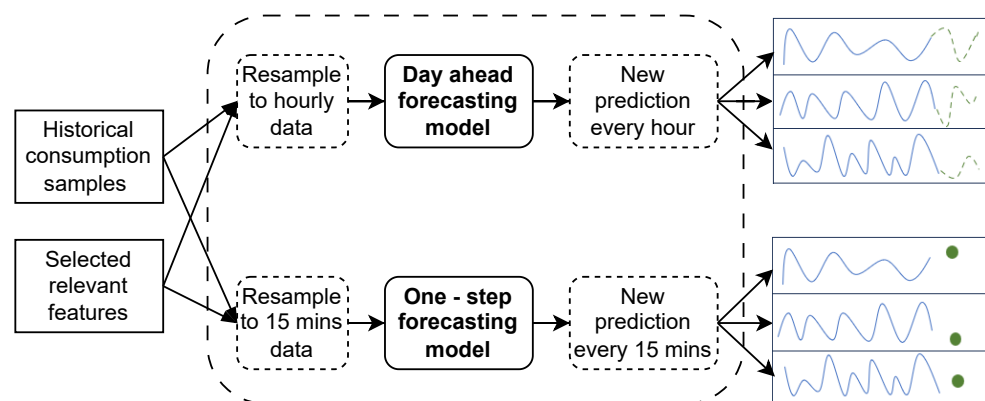


Figure 2. Design of the forecasting algorithms. The dotted box separates the system (inside the box) from inputs and outputs (outside the box). The blue lines represent historical data, and the green lines (and dots) represents, prediction.

A total of six forecasting models have been designed, with two models for each type of consumer (day ahead and 15 min ahead). It is also worth noting that before training the models, a baseline should be established for each model, to compare the results and

metrics obtained by the forecasting algorithms with the baseline's metrics. In the case of the one-step forecast, the baseline is determined by repeating the last lag, i.e., the power demand from 15 min earlier, and comparing it with the current power demand. As for the day-ahead forecast, the baseline is established by repeating the previous 24 samples (the power demand from the previous day) and comparing them to the current day's power demand.

3.2.1. Feature Selection

Regarding the input features of the forecasting models, a feature importance analysis is conducted for each consumer type to guarantee that only the most relevant features are incorporated as inputs. These features include the following:

- Calendar-related categorical features such as the month, day of the week, and time.
- Holiday binary feature.
- Weather data, in particular hourly temperature and humidity data.
- Socio-economical features, specifically the total population, population density, the territorial socio-economic index, and the gross disposable household income.

However, it is important to note that the available socio-economic data was provided as annual aggregates (a single value per year), resulting in an insufficient temporal resolution for the short-term and very-short-term forecasting tasks at hand. Consequently, these features were excluded from the model training and feature selection process. Despite this, a separate correlation analysis was conducted for residential consumers to explore potential relationships between these socio-economic indicators and consumption statistics (e.g., mean, median, min, max consumption), the results of which are discussed in Section 5.2.3.

It is imperative to note that all of these features are included as future covariates. This implies that they must be known in advance, with the requirement being to be known 15 min into the future and 24 h ahead. Furthermore, it is essential that the features are correctly handled to ensure that the models receive the data in the most appropriate way, for which a preprocessing step was included. The weather data was standardized between 1 and -1 , while the encoding of calendar-related features and the holiday feature are model-dependent, varying according to the forecasting model in use. For tree-based models (LightGBM, XGBoost, and RF), these variables were categorically encoded; for the rest (MLP, SVR, and LSTM), they were one-hot encoded.

The feature selection process, as depicted in Algorithm 1, is a Backward Elimination process that employs the training of LightGBM and XGBoost models with and without each feature. LightGBM and XGBoost were selected for this study due to their strong performance with limited data, as well as their ease of training, reliability, and speed [32,44]. These methods offer significant advantages, including robust regularization capabilities and fine-tuned control over model complexity, which are particularly valuable when working with smaller datasets or when overfitting is a significant concern.

However, we also incorporate SHAP to demonstrate the model-agnostic nature of our feature selection process. SHAP provides consistent, interpretable explanations of feature importance by assigning Shapley values that represent each feature's contribution to the model's predictions. Since SHAP is model-agnostic, it enables us to evaluate feature importance independently of any specific model structure. By ranking features based on their SHAP values, we ensure that our feature selection is robust across different models, supporting a transparent and objective approach that avoids reliance on any particular algorithm's behavior. This enhances the reliability and interpretability of the selected features.

Algorithm 1 Feature selection process.

```

1:  $S = \{s_1, s_2, \dots, s_n\}$  ▷ All features
2:  $S' = \{s'_1, s'_2, \dots, s'_m\}$  ▷ Important Features
3: Train single feature LightGBM and XGBoost models
4: Calculate error of the model,  $e_u$ 
5: for  $s_i$  in  $S$  do
6:   Train LightGBM and XGBoost with  $s_i$  as a feature
7:   Calculate the error of the model  $e_i$ 
8:   if  $e_i \gtrsim e_u$  then
9:      $s_i \in S'$ 
10:  else
11:     $s_i \notin S'$ 
12:  end if
13: end for

```

3.2.2. Model Selection

Once the most relevant features have been identified, a series of machine learning models is evaluated, extending beyond the previously employed LightGBM and XGBoost. The following machine learning models are tested: MLP (Multilayer Perceptron), LSTM (Long Short-Term Memory), RF (Random Forest), and SVR (Support Vector Regression). These algorithms were selected as they have consistently demonstrated their efficiency in STLF under different circumstances [40–42]. The majority of studies present either a comparison between different deep learning models (LSTM, MLP, etc.) [35] or a comparison between machine learning models (LightGBM, SVR, RF, etc.) [40].

These algorithms represent diverse approaches to time-series forecasting, including gradient boosting (LightGBM and XGBoost), feedforward neural networks (MLP), recurrent neural networks (LSTM), and a supervised max-margin model (SVR). The purpose of this comparison is to ascertain which technique performs better for which consumer type.

3.2.3. Bayesian Hyperparameter Optimization

In order to optimize the performance of the model in the validation set, it is often necessary to exhaustively explore a vast number of different combinations of hyperparameter values, which is a time-consuming process.

Bayesian optimization is a suitable approach for addressing the aforementioned issue, and can be divided into two main components: (1) a probabilistic surrogate model, which is employed to approximate the current black box objective function; and (2) an acquisition function, which is used to identify the optimal solution under the current data conditions [47]. By reasoning with past search results and attempting more promising values, Bayesian optimization reduces the time required for hyperparameter search and facilitates the identification of optimal values. The Bayesian hyperparameter optimization is represented in Equation (3).

$$x^* = \operatorname{argmin} f(x) \implies x \in \chi \quad (3)$$

where, $f(x)$ denotes the target value to be minimized for evaluation on the validation set, which is the root mean squared error (RMSE) in this case. χ is the domain of hyperparameters [44].

In our study a Tree-Structured Parzen Estimator (TPE) [48] is used as a surrogate model which is used to estimate the densities of good hyperparameters and bad hyperparameters.

3.3. Evaluation Criteria

This subsection describes the evaluation criteria that were derived from input provided by stakeholders involved in the Horizon Europe project, composed of Energy Services Companies (ESCOs), DSOs, and big industrial consumers.

3.3.1. Classification Algorithm Evaluation

The Key Performance Indicator (KPI) set for the classification algorithm is set in terms of accuracy. Drawing on results from the literature [45,46], we have set a threshold requiring the algorithm to achieve at least 75% accuracy in classifying the different customer types.

3.3.2. Forecasting Model Evaluation

The KPIs established for the forecasting models are also expressed in terms of accuracy; the 15 min ahead forecast needs to achieve an accuracy of at least 80%, and at least 85% in the case of the day-ahead forecast. Error metrics that provide information similar to accuracy in a prediction include MAE (Mean Absolute Error) and MAPE (Mean Absolute Percentage Error), shown in (4):

$$\begin{aligned} MAE &= \frac{1}{n} \sum_{i=1}^n |s_i - s_i^*| \\ MAPE &= \frac{100\%}{n} \sum_{i=1}^n \left| \frac{s_i - s_i^*}{s_i} \right| \end{aligned} \quad (4)$$

where s_i and s_i^* denote the actual and forecasted load values respectively.

Accuracy measures how closely forecasted values match actual values, typically expressed as a percentage. MAPE, which is scale-independent and interpretable as a percentage, is useful for comparing forecasts across datasets, while MAE provides error in the same units as the data and is robust to outliers. Optimizing MAPE can lead to demand underestimation, whereas optimizing MAE aims to balance errors, which can be problematic with extreme load variations such as in industrial consumers. Low load values can challenge MAE and MAPE evaluation, as occurs with residential consumers [49]. To address instances of significant error despite low average metrics, we use a quantitative score to measure the frequency of MAPE and MAE falling below a set threshold, highlighting periods of model underperformance potentially due to unaccounted seasonality or complex factors.

It is also important to note that in order to accurately assess the performance of the day-ahead forecaster, a production scenario is generated by predicting the next 24 h at each time step of the testing data, rather than using a single prediction every 24 h. This approach allows for the creation of overlapping predictions, which can then be individually evaluated to provide a more comprehensive assessment of the forecaster's accuracy.

4. Case Study

In this section, we will introduce the specific case study under investigation and provide a detailed description of the datasets utilized in the analysis, including their sources, and the types of data collected. This section aims to establish a clear understanding of the foundational elements of the study, setting the stage for subsequent discussions of the results.

4.1. Historical Load Data

The available dataset comprises 15 min interval data for two years of power consumption for real-use cases located in Spain, of four *industrial* and two *commercial* consumers, data provided by an SME specialized in global electrical solutions, control, electrical mainte-

nance, and data analysis. Moreover, a DSO also provided data for five *residential* consumers. Spain has different climate zones, as depicted in Figure 3, and as we are going to take weather-related variables into account, it is important to know the location of each customer. In our case, specifically, the *industrial* and *commercial* consumers are situated in the Basque Country, near San Sebastian, on the coast of the Bay of Biscay, dominated by a warm, humid, and wet oceanic climate. On the other hand, the *residential* consumers are located in Catalonia, near Barcelona, in the northeast of Spain, with a Mediterranean climate, Figure 3. Despite climate differences, these two regions have similar GDPs and socio-economic characteristics.

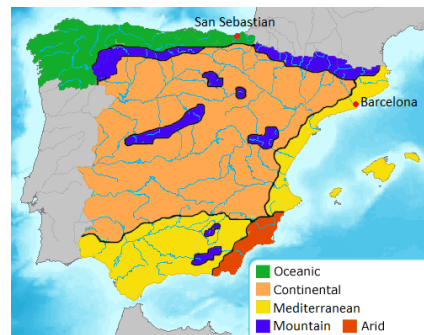


Figure 3. Different climate zones in Spain. The colors represent different climatic zones of Spain. Green: oceanic. Light orange: continental. Dark blue: Mountain. Yellow: mediterranean. Dark orange: arid. Light blue: water.

4.2. Weather Data

Weather-related data is utilized as input features for the forecasting models, with the impact of this data on forecasting accuracy analyzed for each consumer type. Different consumers have varying relationships with specific weather variables that affect their energy usage and environmental performance. According to [15] sixteen parameters are widely used for different energy applications, with dry-bulb temperature [$^{\circ}\text{C}$] being the most prevalent, followed by humidity [%], both of which are essential and highly influential across most energy applications throughout their life cycle. Wind speed [m/s] and global solar radiation [W/m^2] are also frequently used, particularly in studies focused on specific buildings [50,51].

However, in our study, since we lack information on specific buildings, we will focus solely on humidity and dry-bulb temperature. Although these variables were initially obtained from the ERA5 dataset [16], the resulting forecasting accuracy was not sufficiently high due to the spatial resolution limitations of the ERA5 and CAMS free services. To enhance the accuracy of the forecasts, we have incorporated additional data from the national meteorological agency, specifically Aemet [52].

4.3. Socio-Economic Data

Socio-economic factors play a significant role in demand forecasting across different regions. However, for industrial and commercial consumers, where the businesses are of similar size and located in the same region, these factors are excluded from the analysis as they provide limited additional information. Thus, this data is only used for residential consumers, which correspond to different towns near Barcelona, and for each of the towns, the following data was collected from the regional institute of statistics, Idescat [53]: (I) total population, (II) population density, (III) TSI (territorial socio-economic index), and (IV) GDHI (Gross Disposable Household Income). Income levels impact the affordability and use of energy-intensive appliances, with higher-income households more likely to invest in energy-efficient technologies, reducing overall consumption [12,13]. Addition-

ally, household size, employment rates, and daily routines influence energy consumption, with demand fluctuating based on periods of activity and rest.

5. Results

This section presents a comprehensive discussion of the results obtained from the customer classification, feature selection, and forecasting models. We will analyze the effectiveness of the classification process in accurately grouping customers into their respective categories, followed by an evaluation of the feature selection process employed to identify the most relevant variables influencing demand. Finally, we will conduct a detailed examination of the performance of the forecasting models, each of which has been specifically tailored to the respective customer groups. In cases where the results do not meet the established standards, alternative approaches will be introduced and compared, demonstrating their contribution to enhancing previous predictions. The insights derived from these analyses will help to inform the broader implications of the study.

5.1. Customer Classification Results

Results from the rule-based classification algorithm are shown in the confusion matrix of Figure 4. There is a significant imbalance in the number of available datasets for each category, which poses a challenge for conducting comprehensive analyses. To address this issue, we have divided the underrepresented datasets into more subsets (6 subsets in the case of residential datasets, and 30 in industrial and commercial datasets) to ensure a balanced representation across different categories. This approach allows for a more equitable distribution of data, enhancing the robustness of the analysis while reducing potential biases caused by the overrepresentation of certain dataset types. By doing so, we aim to provide a more thorough and accurate evaluation that accounts for the inherent data discrepancies. The algorithm performs correctly in overall 85% of the cases. The industrial datasets are classified correctly 77% of the time, accuracy in residential datasets is 87%, and accuracy in commercials is 89%. These values comply with the established KPIs, which state that the classification algorithm must cluster consumers (residential, commercial, industrial) with a precision of >75%.




Figure 4. Confusion matrix of dataset classifier.

5.2. Feature Selection Results

5.2.1. Industrial Consumers

The results presented in Table 1 for the day-ahead forecasting models indicate whether the model's performance improves (highlighted in green) or worsens (highlighted in red)

in terms of MAPE and MAE. These comparisons are made concerning the XGBoost and LightGBM methods following the introduction of various features.

These results indicate that weather does not significantly impact the power consumption of industrial consumers. Additionally, neither does the 15 min ahead forecaster. These results are coherent with the fact that the main power consumption source for industrial consumers is their heavy machinery, and although their overall consumption may be affected by temperature and humidity, this influence is negligible compared to the demand from the machinery.

On the other hand, the inclusion of *calendar information* (hour, day of the week, and month) is beneficial for the forecasting models as power consumption exhibits seasonal patterns. This allows the model to capture these seasonalities more accurately. This is again coherent, considering that the factories' production is directly related to the day of the week-, month-, and an hour-related schedule.

Finally, regarding the binary *holiday variable*, the forecasters show a significant improvement in accuracy, which is an expected result, as the majority of factories close up during holidays, or at least, work at a minimum capacity, so their consumption is directly altered by that factor.

We will compare these results with SHAP analyses to ensure that the feature importance and performance improvements are model-agnostic, aiming to confirm that the observed trends hold consistently across different models, providing greater confidence in the robustness and generalizability of our findings. With Figure 5, we can confirm that the most important features are the hour and holiday data, whereas the weather data and month/weekday features do not impact as much.

Table 1. Error change of industrial models to perform feature selection.

Feature	Model	Metric	Ind.1	Ind.2	Ind.3	Ind.4
Temperature and Humidity	XGBoost	MAPE	-1.1	0.5	0.3	0.0
		MAE	-0.8	1.5	0.7	-0.1
	LightGBM	MAPE	-0.7	-0.1	-1.3	-0.1
		MAE	-0.8	-0.1	-0.7	-0.2
Calendar Features	XGBoost	MAPE	-30.9	-8.1	-1.0	-1.2
		MAE	-6.0	-5.4	-2.7	-2.0
	LightGBM	MAPE	-28.7	-6.4	-5.5	-1.5
		MAE	-5.9	-6.3	-0.4	-1.6
Holiday Indicator	XGBoost	MAPE	-14.1	-11.0	-3.1	-1.1
		MAE	-3.8	-5.7	-1.4	-1.7
	LightGBM	MAPE	-20.4	-10.8	-4.4	-0.8
		MAE	-4.9	-5.5	-2.1	-1.2




Figure 5. Shap values for the industrial model features.

5.2.2. Commercial Consumers

In the same way, the results gathered in Table 2 indicate that although slightly, all the features improve the accuracy of the models; however, this improvement is not always significant. The case of *calendar information* and the *holiday variable* is similar to that of industrial consumers, as their consumption has seasonality related to the calendar schedule and the holidays.

Nevertheless, in this case, unlike in the industrial case, the weather-related variables do affect people’s commercially related behavioral patterns. On the one hand, reaching the temperature comfort levels inside a commercial building will directly affect power consumption. On the other hand, it also relates to encouraging people to visit these commercial sites, which in turn again directly impacts temperature comfort levels.

Again, we will compare these results with SHAP analyses. With Figure 6, we confirm that the most important features are the hour and holiday data. We can furthermore elaborate on the results in Table 2 related to weather variables, as temperature is the important one, while humidity is virtually irrelevant.

Table 2. Change in error of commercial models after feature selection.

Feature	Model	Metric	Com.1	Com.2
Temperature and Humidity	XGB	MAPE	−0.12	−2.21
		MAE	−0.01	−0.14
	LGBM	MAPE	−0.9	−1.43
		MAE	−0.54	−0.12
Calendar Features	XGB	MAPE	−1.59	−6.36
		MAE	−0.46	−0.22
	LGBM	MAPE	−1.35	−4.23
		MAE	−0.78	−0.10
Holiday Indicator	XGB	MAPE	−1.08	−0.92
		MAE	−1.01	−0.15
	LGBM	MAPE	−2.44	−2.36
		MAE	−1.35	−0.10




Figure 6. Shap values for the commercial model features.

5.2.3. Residential Consumers

Finally, the results for the residential consumers gathered in Table 3 show that the impact of the different features is not quite significant, although intuitively, we would expect a higher impact of *temperature* and *humidity*. We would expect residential consumers to be more likely to adjust their energy usage based on temperature changes, leading to

spikes in demand for heating or cooling, but in this case study that relationship does not seem so direct.

Regarding the inclusion of *calendar information* (hour, day of the week, and month) and the binary *holiday variable*, they too have a very small impact on the predictions.

However, if we compare the results with SHAP analyses in Figure 7, we see that the most important feature is again the hour. As expected, temperature ranks as the second most important feature, while other calendar-related data, including holiday information—which is the second most important feature for other consumers—appear to be less relevant in this case.

Table 3. Change in error of residential models after feature selection.

Feature	Model	Metric	Res.1	Res.2	Res.3	Res.4	Res.5
Temperature and Humidity	XGB	MAPE	−0.9	−0.55	−0.9	−1.0	−0.73
		MAE	0.0	0.0	0.0	−0.01	0.0
	LGBM	MAPE	−0.48	−0.47	−0.01	0.04	−0.1
		MAE	−0.01	0.0	0.0	0.0	0.0
Calendar Features	XGB	MAPE	−1.24	−1.13	−1.2	−2.55	−3.38
		MAE	−0.01	−0.01	−0.02	0.0	0.0
	LGBM	MAPE	−1.31	−1.97	0.1	−1.4	−1.87
		MAE	0.0	−0.02	−0.01	−0.01	−0.01
Holiday Indicator	XGB	MAPE	0.56	0.0	−0.59	−0.84	−0.44
		MAE	0.01	0.0	0.0	0.0	0.0
	LGBM	MAPE	−0.49	−0.82	0.0	−0.1	0.0
		MAE	0.0	0.0	0.0	0.0	0.0




Figure 7. Shap values for the residential model features.

As mentioned in Section 3.2.1, socio-economic features were excluded from the forecasting models due to their low temporal resolution. Nonetheless, an exploratory correlation analysis was performed to assess the relationship between annual socio-economic indicators—total population, population density, territorial socio-economic index (TSI), and Gross Disposable Household Income (GDHI)—and various aggregated consumption statistics (minimum, maximum, mean, and median power) for the residential datasets. The results of this analysis, presented in Table 4, indicate weak to moderate correlations. Population density showed the strongest negative correlations with consumption metrics (e.g., −0.60 with median consumption), suggesting that higher density areas might be associated with slightly lower per-consumer energy use, possibly due to factors like smaller living spaces or shared infrastructure. Correlations with TSI and GDHI were generally very weak or negligible. This preliminary analysis suggests that while some broad trends

may exist, the static, annual nature of this data limits its utility for LTLF, confirming our decision to omit it from the models.

Table 4. Pearson correlation coefficient between residential power demand and socio-economic variables.

	Total Population	Population Density	TSI	GDHI
P_min	−0.50	−0.60	−0.38	−0.39
P_max	−0.40	−0.52	0.02	0.11
P_mean	−0.51	−0.58	−0.20	−0.16
P_median	−0.53	−0.60	−0.22	−0.19

5.3. Model Selection Results

In this section, the models' results are evaluated and discussed, based on the KPIs established by the stakeholders of the Horizon project, which are the accuracies to be achieved. Although it can be challenging to precisely evaluate trained models regarding their accuracy, it can be done through diverse metrics that represent different error characteristics, such as the MAPE and MAE, which are valuable tools, especially when used in conjunction with quantitative assessments. The overall MAPE and MAE results gathered in Tables 5 and 6, and are therefore combined with a quantitative score, introduced in Section 3.3.

Table 5. Error metrics of 24 h ahead models. (Best metric for each consumer is highlighted in bold).

Model	Ind.1	Ind.2	Ind.3	Ind.4	Com.1	Com.2	Res.1	Res.2	Res.3	Res.4	Res.5
MAPE (%)											
Baseline	64.31	60.77	63.03	8.70	14.97	87.05	31.32	36.77	24.96	28.15	41.81
MLP	40.39	30.39	30.51	5.18	7.75	80.02	33.42	36.07	24.01	27.52	36.39
LSTM	22.28	13.45	17.95	4.58	5.77	26.80	33.77	35.15	26.57	30.14	33.79
RF	41.66	29.46	39.86	6.39	7.28	59.77	27.15	32.69	27.23	22.34	33.66
SVR	38.96	27.18	32.75	6.71	6.87	78.97	24.00	29.27	20.18	21.54	30.68
XGBoost	26.17	17.13	29.43	4.62	5.76	44.91	23.06	29.20	23.99	22.27	32.85
LightGBM	18.01	18.39	24.37	4.33	5.17	48.17	22.52	28.56	21.64	21.19	31.67
MAE (kW/h)											
Baseline	22.16	34.50	55.76	12.34	8.19	4.84	0.76	1.03	1.04	0.52	0.73
MLP	12.67	16.43	22.95	6.84	4.36	3.28	0.73	0.90	0.94	0.50	0.63
LSTM	10.24	10.49	17.36	6.19	3.36	1.52	0.69	0.86	1.12	0.48	0.58
RF	14.37	21.29	27.84	8.86	4.15	2.81	0.61	0.84	1.04	0.40	0.58
SVR	11.90	17.09	23.78	8.77	3.97	3.72	0.59	0.85	0.91	0.41	0.60
XGBoost	9.51	12.18	21.21	6.29	3.23	3.07	0.55	0.79	0.95	0.40	0.58
LightGBM	8.13	12.09	19.47	5.95	2.95	3.14	0.53	0.78	0.87	0.39	0.56

To calculate said quantitative score and evaluate the performance of the model, thresholds need to be established for MAPE and MAE. MAPE provides an overall assessment of the model's predictive accuracy. Since the target accuracy is over 80% for day-ahead forecasts and over 85% for 15 min ahead forecasts, we will set the MAPE threshold at 20% for day-ahead forecasts and at 15% for 15 min ahead forecasts. These thresholds will apply to all consumer types, as MAPE is scale-independent.

On the other hand, while MAE also measures the average magnitude of errors between predicted and actual values across an entire dataset, it is scale-dependent, meaning it will vary for each consumer. To establish the threshold for MAE, we will allow a 20% error for day-ahead forecasts and a 15% error for 15 min ahead forecasts. For each consumer,

the mean value is obtained, and the corresponding 20% and 15% values are calculated to establish the MAE thresholds discussed in Sections 5.3.1–5.3.3.

Table 6. Error metrics of 15 min ahead models. (Best metric for each consumer is highlighted in bold).

Model	Ind.1	Ind.2	Ind.3	Ind.4	Com.1	Com.2	Res.1	Res.2	Res.3	Res.4	Res.5
MAPE (%)											
Baseline	6.90	14.13	7.24	7.02	10.05	21.88	12.19	11.99	10.16	14.81	11.26
MLP	7.72	12.45	6.79	5.77	7.44	27.90	9.54	10.09	9.87	12.40	9.39
LSTM	9.79	13.45	7.66	6.28	9.02	26.79	9.84	10.65	10.74	12.27	9.54
RF	7.19	11.02	6.90	5.76	7.87	24.18	9.24	9.79	9.28	11.98	9.26
SVR	9.63	13.77	6.93	5.58	7.62	25.88	9.17	9.82	9.02	12.36	9.45
XGBoost	6.45	10.88	6.54	5.50	7.45	21.17	9.33	10.27	10.86	12.09	9.37
LightGBM	6.58	10.63	6.65	5.46	6.94	15.93	10.44	10.47	11.23	12.49	9.28
MAE (kW/h)											
Baseline	4.58	8.99	9.74	9.88	5.86	1.85	0.074	0.083	0.109	0.066	0.052
MLP	4.49	7.92	8.90	8.11	4.25	1.63	0.058	0.068	0.101	0.055	0.043
LSTM	4.98	9.82	7.99	8.72	5.18	1.66	0.059	0.070	0.108	0.053	0.043
RF	4.63	7.51	9.16	8.15	5.50	1.63	0.057	0.067	0.098	0.054	0.042
SVR	4.57	8.25	8.64	7.85	4.42	1.63	0.056	0.067	0.097	0.055	0.043
XGBoost	4.08	7.71	8.62	7.78	4.28	1.62	0.056	0.069	0.108	0.054	0.043
LightGBM	3.92	7.01	8.27	7.73	3.96	1.33	0.060	0.069	0.109	0.055	0.042

Finally, we will set the target quantitative score values to align with the accuracy benchmarks to be achieved. This approach ensures that the model is sufficiently robust and effective in consistently meeting the desired MAPE and MAE thresholds over extended periods. These values are valid for all types of consumers.

5.3.1. Industrial Consumers

For industrial consumers, the chosen model is LightGBM, as it is the best-performing one according to the results in Tables 5 and 6. However, if we plot the MAPE of the day-ahead forecasters over one month of testing data, as in the examples in Figure 8, it can be seen that it underperforms when the dates correspond to holidays: the marked area relates to the Easter holidays of 2021.




Figure 8. Evolution of MAPE in different industrial consumers. The pink color represents the section where the models underperform.

This issue primarily arises because, despite incorporating holiday data as an exogenous variable, the forecaster occasionally misclassifies a holiday as a working day, resulting in incorrect predictions of peak power demand. Additionally, industrial consumers exhibit distinctly bimodal behavior, with significant differences in consumption between weekdays and weekends or holidays. These discrepancies mean that if weekday consumption is predicted incorrectly, the resulting error is much larger compared to other consumer types, such as residential consumers, where consumption values do not vary as drastically.

Proposed Approach

To address this problem, another forecasting approach is proposed below, depicted in Figure 9, which combines, or fuses, two distinct models, both LightGBM-based, to perform the prediction: we train two distinct forecasting models, one only for holidays and the other one for working days. Then, at prediction time, both models perform the prediction but as the holiday information of the next 24 h is known, only the prediction for the correct model is taken.

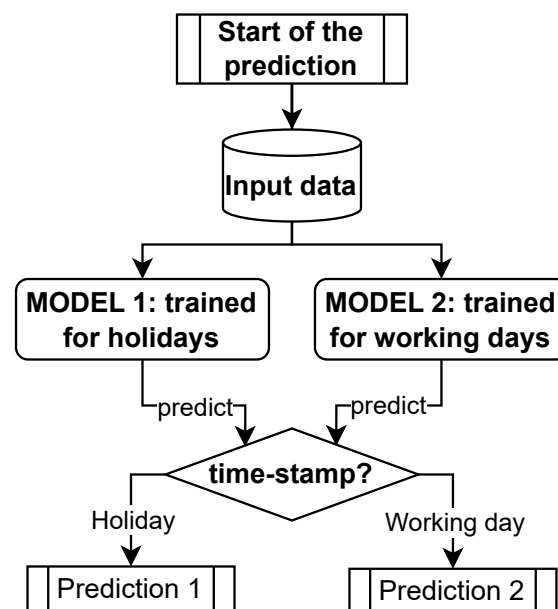


Figure 9. Proposed approach.

In Figure 10, it can be seen that, although some peaks can still be appreciated, a significant improvement is achieved in terms of MAPE, and how it holds over the tested month, even in holidays.

In Tables 7 and 8 results from applying a single model versus applying the proposed fusion approach are compared, both in terms of MAPE, MAE, and the quantitative score reached.

Table 7. Comparison of MAPE results for the industrial consumers with the proposed approach.

	Day Ahead				15 min Ahead			
	1 Model		Fusion		1 Model		Fusion	
	MAPE [%]	Score [%]	MAPE [%]	Score [%]	MAPE [%]	Score [%]	MAPE [%]	Score [%]
Ind.1	18.0	67.4	13.8	82.8	6.2	96.1	6.1	96.2
Ind.2	18.4	84.1	10.4	97.9	10.7	86.3	10.2	88.2
Ind.3	24.4	69.2	13.6	87.4	6.6	96.5	6.5	96.4
Ind.4	4.3	100.0	4.6	100.0	5.4	98.6	5.2	98.8

Table 8. Comparison of MAE results for the industrial consumers with the proposed approach.

	Day Ahead					15 min Ahead				
	Thr	1 Model		Fusion		Thr	1 Model		Fusion	
		MAE [kW]	Score [%]	MAE [kW]	Score [%]		MAE [kW]	Score [%]	MAE [kW]	Score [%]
Ind.1	10	8.1	90.3	7.3	92.6	7.5	3.9	84.4	3.7	84.6
Ind.2	25	12.1	72.1	9.5	86.1	18.8	7.1	92.5	6.8	92.9
Ind.3	22.5	19.5	54.9	16.1	82.8	16.9	8.5	86.5	8.4	86.5
Ind.4	15	5.9	98.7	6.2	98.7	11.3	7.6	78.1	7.4	80.8

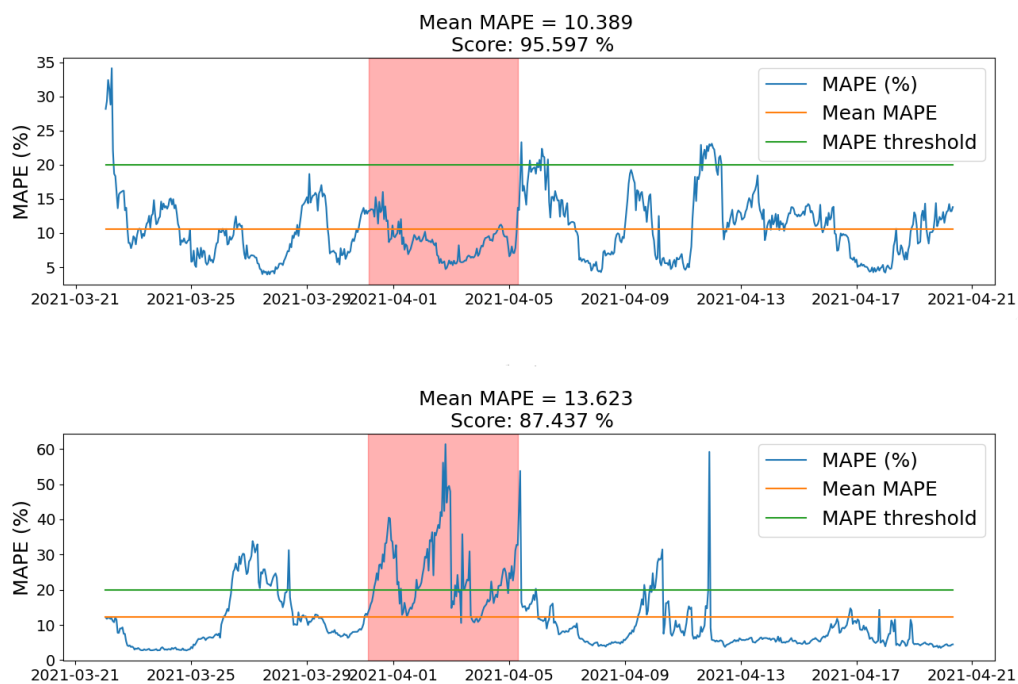


Figure 10. Evolution of MAPE in different industrial consumers after applying proposed approach. The pink color represents the section where the models underperform.

The results indicate that, although the MAPE and MAE values are below the established thresholds even when applying a single model, the quantitative score obtained does not meet the required standards, suggesting that the model lacks robustness. In contrast, the model fusion approach produces better outcomes overall, significantly lowering both MAPE and MAE while also increasing the quantitative score values. This improvement reflects the enhanced overall performance of the model over extended periods, demonstrating its superior effectiveness and stability.

This approach enables individual models to more effectively learn the intrinsic patterns of load and significantly outperforms other single-model approaches. However, a potential limitation of this approach is that it relies on accurate holiday information, which can compromise flexibility for the sake of precision.

5.3.2. Commercial Consumers

For commercial consumers too, the chosen model is LightGBM, as it is the best performing one according to the results in Tables 5 and 6. Nevertheless, when we again plot the MAPE of the day-ahead forecasters over one month of testing data. In Figure 11, we observe that the model underperforms for the second commercial customer on dates corresponding to holidays. The marked area highlights two Spanish national holidays on 6 and 8 December, as well as Christmas Day in 2023.




Figure 11. Evolution of MAPE for a commercial consumer. The pink color represents the section where the models underperform.

This, however, only happens for this second customer, and in the day-ahead forecaster as the results of the MAPE and MAE and the achieved scores are gathered in Tables 9 and 10 show. Again, the main error source is holidays, so to solve that, the same two-model fusion approach, depicted in Figure 9 has been applied to commercial consumers, in the same way as it was done for industrial consumers.

Table 9. Comparison of MAPE results for the commercial consumers with the proposed approach.

	Day Ahead				15 min Ahead			
	1 Model		Fusion		1 Model		Fusion	
	MAPE [%]	Score [%]	MAPE [%]	Score [%]	MAPE [%]	Score [%]	MAPE [%]	Score [%]
Com.1	5.17	99.7	5.92	100	6.9	90.6	7.09	91.1
Com.2	48.2	70.6	26.4	87.2	15.9	64.9	17.6	61.9

Table 10. Comparison of MAE results for the commercial consumers with the proposed approach.

	Day Ahead					15 min Ahead				
	Thr	1 Model		Fusion		Thr	1 Model		Fusion	
		MAE [kW]	Score [%]	MAE [kW]	Score [%]		MAE [kW]	Score [%]	MAE [kW]	Score [%]
Com.1	6	2.9	100.0	3.3	99.7	4.5	3.9	64.1	4.05	64.6
Com.2	8	3.1	79.77	2.2	95.6	6	1.3	95.0	1.44	95.8

The results indicate that, although the MAPE and MAE values are below the established thresholds in both consumers as well, the quantitative score obtained from a single model does not meet the required standards for the second consumer. In this scenario, the model fusion approach yields better outcomes, reducing both MAPE and MAE while also improving the quantitative score. For the first consumer, despite a slight increase in MAPE and MAE, the values remain well below the thresholds, and the quantitative score overall suggests strong model performance over extended periods.

5.3.3. Residential Consumers

Forecasting residential power demand presents unique challenges that differ significantly from those encountered in predicting commercial or industrial power demand, as can be perceived from the results gathered in Tables 5 and 6. One of the primary difficulties in residential demand forecasting is the high degree of variability and unpredictability in individual consumption patterns. Unlike commercial or industrial settings, where power usage is governed by more consistent schedules and operational needs, residential energy

consumption is influenced by a wide range of factors including individual behaviors, household routines, and the use of various home appliances—factors for which we lack detailed information. These factors can lead to sudden and irregular fluctuations in demand, making it harder to model and predict accurately.

Thus, the datasets are in general irregular, with peaks at given hours. Some of the peaks are very periodic and are produced at the same hour every day, but other peaks vary more in time and are more difficult to predict. This fact is described by the weekly and daily autocorrelation plots in Figure 12. In some municipalities, peaks are more regular and, therefore, the autocorrelation is considerable, but in other municipalities, autocorrelation is lower, resulting in irregularities in its demand peaks. That is why here again, it is not sufficient to apply a single model directly.




Figure 12. Autocorrelation plot of (a) a municipality with irregular peaks and (b) a municipality with regular peaks.

Proposed Approach

In this case, the issue is not related to identifying holidays or the influence of the bimodal behavior of customers, but rather lies in the high variability in consumption patterns. This is why this time we seek to reinforce the memory of the model, combining a time series forecasting model with a baseline prediction. The study also considers analogous baseline-refinement approaches, as exemplified in reference [54].

Adding this baseline prediction offers several advantages in enhancing predictive accuracy and robustness. A time series model captures temporal patterns, trends, and seasonalities inherent in historical data, providing nuanced and dynamic forecasts. However, baseline predictions, often simpler and based on historical averages or other straightforward methods, serve as a reliable reference point. By integrating these two approaches, the model benefits from the sophisticated pattern recognition of time series analysis while maintaining the stability and simplicity of baseline predictions. This combination helps mitigate the risk of overfitting, improves performance during anomalous events or data fluctuations such as the mentioned irregularities in peak demand, and ensures more consistent and reliable predictions.

Regarding the day-ahead forecast, for every point, the baseline method calculates the arithmetic mean of the consumption on the last day and the last week at the same hour,

as in Equation (5). Then, the series of the baseline prediction is used as an exogen variable in a gradient-boosting model.

$$x_i = 0.5 \cdot (x_{i-24} + x_{i-7.24}) \tag{5}$$

For the 15 min ahead forecast, the baseline has been calculated following Equation (6). The assigned weights are 0.6 to the last point, 0.2 to the mean consumption of the last 4 days at that same time, and 0.2 to the mean consumption of the last 4 weeks at the same time and weekday. The weights were empirically determined to maximize forecast accuracy on validation data, reflecting the higher relevance of recent observations while retaining seasonal patterns. This way, the baseline has a memory of consumption during the last days, not only the last value.

$$x_i = 0.6 \cdot \bar{x}_{i-1} + \frac{0.2}{4} \cdot (x_{i-24} + x_{i-2.24} + x_{i-3.24} + x_{i-4.24}) + \frac{0.2}{4} \cdot (x_{i-24.7} + x_{i-2.24.7} + x_{i-3.24.7} + x_{i-4.24.7}) \tag{6}$$

In general, the results gathered in Tables 5 and 6 suggest that the predictability of these consumers, with highly variable patterns, is further challenged by the relatively low magnitude of residential consumption. The model prioritizes accuracy in predicting the lowest consumption values because errors at low levels significantly increase the loss function in the case of MAPE, but then the peak demand predictions are missed. Conversely, prioritizing the accuracy in peak demand points highly increases the MAPE, as the error increases in the low consumption part.

Furthermore, as shown in Equation (4), in order to calculate the MAPE, it is necessary to divide by the actual value of the dataset. For *residential* consumers, these values are often less than 1 kW, which leads to significantly higher MAPE values due to the division by such small numbers. That is why, in the case of *residential* consumers, the thresholds of the MAPE are raised up to 30% for the day-ahead forecasts and to 25% for the 15 min ahead forecasts. The results comparing one model with the proposed hybrid approach are shown in Tables 11 and 12.

Table 11. Comparison of MAPE results for the residential consumers with the proposed approach.

	Day Ahead				15 min Ahead			
	1 Model		Hybrid		1 Model		Hybrid	
	MAPE [%]	Score [%]	MAPE [%]	Score [%]	MAPE [%]	Score [%]	MAPE [%]	Score [%]
Res.1	22.5	96.4	21.9	95.4	10.4	90.0	9.2	95.1
Res.2	28.6	67.9	28.6	68.4	10.5	92.8	9.9	93.5
Res.3	21.6	88.6	20.8	91.8	11.2	88.7	10.2	92.8
Res.4	21.2	97.4	20.6	98.0	12.5	88.6	11.7	91.2
Res.5	31.7	42.1	30.3	47.9	9.3	94.4	9.1	95.4

Table 12. Comparison of MAE results for the residential consumers with the proposed approach.

	Day Ahead					15 min Ahead				
	Thr	1 Model		Hybrid		Thr	1 Model		Hybrid	
		MAE [kW]	Score [%]	MAE [kW]	Score [%]		MAE [kW]	Score [%]	MAE [kW]	Score [%]
Res.1	0.61	0.53	88.2	0.53	83.4	0.46	0.06	99.8	0.05	99.8
Res.2	0.79	0.78	62.5	0.78	63.5	0.58	0.07	99.9	0.07	99.9
Res.3	1.20	0.87	93.4	0.87	93.1	0.90	0.11	99.9	0.10	100
Res.4	0.39	0.39	57.2	0.38	61.8	0.29	0.06	98.9	0.05	99.0
Res.5	0.43	0.56	15.6	0.55	19.9	0.32	0.04	99.4	0.04	99.5

The results indicate that, although the MAPE and MAE values are below the established thresholds, or very close to them, even when applying a single model, the quantitative score obtained does not meet the required standards for consumers 2 and 5 in the day-ahead forecast, suggesting that the model lacks robustness.

The hybrid approach yields improved results overall, reducing both MAPE and MAE while also increasing the quantitative score values, although not enough to meet the quantitative score thresholds for consumers 2 and 5. These modest improvements underscore the complexity of forecasting residential demand. However, the model's overall performance is enhanced over extended periods, demonstrating its superior effectiveness and stability.

6. Discussion of Results

The findings of this study are consistent with those in previous research, underscoring the significance of consumer-specific forecasting models. Overall, this study reinforces the value of tailored forecasting strategies and advanced machine learning techniques, as seen in [32,40], offering DSOs practical tools to enhance grid stability and efficiency. The key insight for practitioners is that high accuracy can be achieved without excessive complexity: simple clustering rules are effective, tree-based models like LightGBM often outperform more complex deep learning architectures, and incorporating domain knowledge (like separate models for holidays) can resolve significant error sources. The classification algorithm employed in this study attained accuracies that are analogous to those reported in studies such as [11,22]. These studies emphasised the advantages of customer segmentation, though we acknowledge the need for additional consumers to be included in future work to ensure more extensive validation. For industrial consumers, calendar-related features and holidays were key predictors, supported by findings in [41]. The fusion approach improved accuracy during holidays, addressing challenges noted in [43]. Commercial consumers showed similar dependencies on calendar features, with weather conditions, particularly temperature, playing a significant role, as seen in [14,50]. Residential forecasting proved most challenging due to high variability, consistent with [26,54]. Our hybrid approach, combining time series forecasting with baseline predictions, improved robustness but highlighted the complexity of residential demand. Overall, this study reinforces the value of tailored forecasting strategies and advanced machine learning techniques, as seen in [32,40], offering DSO tools to enhance grid stability and efficiency.

7. Conclusions

In conclusion, this study has achieved several milestones in the development and application of forecasting methodologies for different consumer types. We have successfully created a straightforward yet highly accurate customer classification system that effectively differentiates between industrial, commercial, and residential consumers. This classification system is foundational for tailoring forecasting models to the specific needs of each consumer type.

Our feature analysis has provided critical insights into the most significant variables influencing each consumer category. For industrial consumers, calendar-related features and holiday information emerged as the primary drivers of demand, with weather conditions proving to have negligible effects. In the case of commercial consumers, while calendar-related features and holidays remain crucial, weather conditions also play a notable role in influencing consumption patterns. Residential consumers, on the other hand, exhibit a more nuanced interplay of factors, where calendar-related features, holiday information, and weather variables all contribute, albeit to a lesser extent. Static socio-economic data,

while showing some broad correlations, proved insufficient for short-term forecasting due to its low temporal resolution.

The quantitative results presented in this study robustly support these conclusions. The efficacy of our simple classification rules is confirmed by an overall accuracy of 85% (Figure 4), providing a reliable foundation for model customization. Our feature importance analysis, detailed in Tables 1–3 and supported by SHAP values (Figures 5–7), empirically validates the distinct drivers for each consumer type. Most significantly, the superior performance of our tailored forecasting approaches is demonstrated by the notable improvements in key metrics. For instance, the model fusion strategy for industrial consumers (Tables 7 and 8) reduced the MAPE by up to 10% and significantly increased the quantitative robustness score, effectively mitigating holiday-related prediction errors. Similarly, for the challenging case of residential demand, the hybrid approach (Tables 11 and 12) yielded consistent, albeit modest, improvements in both MAPE and MAE, enhancing model stability. These data-driven results underscore that our methodology not only identifies critical patterns but also translates them into tangible gains in forecasting accuracy and reliability.

Furthermore, we explored a diverse array of AI and machine learning algorithms for load forecasting, rigorously comparing their performance to identify the most effective methods. When initial forecasting results did not meet the required standards, we introduced novel forecasting approaches that demonstrated superior performance compared to the previously tested models. In addition, it should be mentioned that we developed and tested the forecasters as if they were deployed in an online environment; in the case of the day-ahead forecaster, a production scenario was simulated by predicting the next 24 h at each time step of the testing data, rather than making a single prediction every 24 h. This approach more accurately reflects real-world conditions and enhances the reliability of the models. Moreover, the demonstrated improvements in forecasting accuracy, particularly the reduction of MAPE by 1–10% for commercial and industrial consumers, have direct practical implications for grid operators. Enhanced short-term forecasting allows DSOs to optimize energy procurement, reduce imbalance costs, and improve the scheduling of grid assets. For industrial consumers, more accurate predictions can facilitate participation in demand–response programs and enhance energy efficiency. While a precise quantification of these economic benefits requires access to proprietary grid operational and market data, which was beyond the resources of this study, the potential for significant cost savings and increased grid stability is a clear and valuable outcome of this work.

In conclusion, it is imperative to emphasize that the primary objective of the study was to formulate a comprehensive methodology for the classification and forecasting of load across diverse consumer types. The efficacy of grid planning and operation is contingent on the accuracy of load forecasts. Consequently, the study sought to ascertain the impact of distinct characteristics and forecasting models on STLF and VSTLF, respectively, for each consumer category. The insights derived from this research are designed to serve as a guide for DSOs, aiming to enhance their grid planning, synchronisation, and operational capabilities.

8. Future Work

In this section, we outline potential directions for advancing the research presented in this study. The primary goal is to deploy the developed models in an online environment, enabling real-time demand forecasting. This deployment involves several critical steps to ensure the system's effectiveness and reliability. First, the input data must undergo thorough preprocessing, including robust outlier detection using interquartile range methods and advanced data imputation techniques to handle missing values. Second,

the deployment architecture must incorporate fail-safes for data pipeline interruptions, including fallback forecasting models that can operate with reduced feature sets. Third, computational efficiency is paramount for real-time operation, necessitating model quantization and optimization for edge deployment where appropriate. Fourth, cybersecurity measures must be integrated throughout the pipeline to protect sensitive consumption data. Finally, we have implemented a monitoring framework that includes not only Change Point Detection and Concept Drift Detection as previously mentioned, but also performance degradation alerts and automated model versioning to ensure continuous service quality.

Once the input data is properly processed, the classification algorithms and forecasting models would be deployed in an online environment allowing the models to continuously receive and process new data, generating up-to-date predictions. To maintain and enhance the accuracy and relevance of the models over time, it is crucial to implement iterative refinement processes. These could include Change Point Detection, which identifies significant shifts in data patterns that could affect model performance; Concept Drift Detection, which monitors and adapts to changes in the underlying data distribution; and Model Retraining, which periodically updates the models based on the latest available data. By incorporating these refinement processes, the forecasting models will be better equipped to sustain high performance and adapt to evolving conditions, ensuring their long-term utility in managing the electrical grid effectively.

While our case study was constrained by the available consumer data from the two Spanish regions, we acknowledge that dataset imbalance represents a limitation of our study. Future work should incorporate larger and more balanced datasets to enhance the generalizability of the clustering and forecasting approaches across diverse consumer populations.

Additionally, future work could explore the fusion of high-performing models like RF and SVR for residential forecasting, particularly to assess whether such combinations can further improve accuracy beyond the current hybrid baseline approach.

Author Contributions: Conceptualization, A.D.-I.; methodology, A.D.-I. and X.B.; software, A.D.-I. and X.B.; validation, A.D.-I., X.B. and A.M.F.-T.; formal analysis, A.D.-I.; investigation, A.D.-I. and X.B.; resources, A.M.F.-T.; data curation, A.D.-I.; writing—original draft preparation, A.D.-I.; writing—review and editing, A.D.-I. and A.M.F.-T.; visualization, A.D.-I. and X.B.; supervision, A.M.F.-T.; project administration, A.M.F.-T.; funding acquisition, A.M.F.-T. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the European Union’s Horizon Europe research and innovation programme under HORIZON-EUSPA-2021-SPACE call, grant agreement No. 101082355, with the acronym RESPONDENT.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to institutional restrictions.

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

References

1. European Commission. The European Green Deal. Available online: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en (accessed on 2 January 2025).
2. Ahmad, T.; Zhang, H.; Yan, B. A review on renewable energy and electricity requirement forecasting models for smart grid and buildings. *Sustain. Cities Soc.* **2020**, *55*, 102052. [CrossRef]
3. Nti, I.K.; Teimeh, M.; Nyarko-Boateng, O.; Adegoye, A.F. Electricity load forecasting: A systematic review. *J. Electr. Syst. Inf. Technol.* **2020**, *7*, 13. [CrossRef]
4. Sheng, Z.; Wang, H.; Chen, G.; Zhou, B.; Sun, J. Convolutional residual network to short-term load forecasting. *Appl. Intell.* **2021**, *51*, 2485–2499. [CrossRef]
5. Hammad, M.A.; Jereb, B.; Rosi, B.; Dragan, D. Methods and models for electric load forecasting: A comprehensive review. *Logist. Sustain. Transp.* **2020**, *11*, 51–76. [CrossRef]

6. Sharma, N.K.; Patnaik, B. Long-Term Load Forecasting: A Systematic Review with Focus on Microgrid Application. In Proceedings of the 2024 Third International Conference on Power, Control and Computing Technologies (ICPC2T), Raipur, India, 18–20 January 2024; pp. 286–291. [\[CrossRef\]](#)
7. Holderbaum, W.; Alasali, F.; Sinha, A. *Energy Forecasting and Control Methods for Energy Storage Systems in Distribution Networks, Predictive Modelling and Control Techniques*; Lecture Notes in Energy; Springer: Cham, Switzerland, 2023. [\[CrossRef\]](#)
8. Khwaja, A.; Anpalagan, A.; Naeem, M.; Venkatesh, B. Joint bagged-boosted artificial neural networks: Using ensemble machine learning to improve short-term electricity load forecasting. *Electr. Power Syst. Res.* **2020**, *179*, 106080. [\[CrossRef\]](#)
9. Quilumba, F.L.; Lee, W.J.; Huang, H.; Wang, D.Y.; Szabados, R.L. Using Smart Meter Data to Improve the Accuracy of Intraday Load Forecasting Considering Customer Behavior Similarities. *IEEE Trans. Smart Grid* **2015**, *6*, 911–918. [\[CrossRef\]](#)
10. Ribeiro, A.M.N.C.; do Carmo, P.R.X.; Endo, P.T.; Rosati, P.; Lynn, T. Short- and Very Short-Term Firm-Level Load Forecasting for Warehouses: A Comparison of Machine Learning and Deep Learning Models. *Energies* **2022**, *15*, 750. [\[CrossRef\]](#)
11. du Toit, J.; Davimes, R.; Mohamed, A.; Patel, K.; Nye, J.M. Customer Segmentation Using Unsupervised Learning on Daily Energy Load Profiles. *J. Adv. Inf. Technol.* **2016**, *7*, 69–75. [\[CrossRef\]](#)
12. del Real, A.J.; Dorado, F.; Durán, J. Energy Demand Forecasting Using Deep Learning: Applications for the French Grid. *Energies* **2020**, *13*, 2242. [\[CrossRef\]](#)
13. Hu, H.; Wang, L.; Peng, L.; Zeng, Y.R. Effective energy consumption forecasting using enhanced bagged echo state network. *Energy* **2020**, *193*, 116778. [\[CrossRef\]](#)
14. Perera, A.T.D.; Nik, V.M.; Chen, D.; Scartezzini, J.L.; Hong, T. Quantifying the impacts of climate change and extreme climate events on energy systems. *Nat. Energy* **2020**, *5*, 150–159. [\[CrossRef\]](#)
15. Amin, A.; Mourshed, M. Weather and climate data for energy applications. *Renew. Sustain. Energy Rev.* **2024**, *192*, 114247. [\[CrossRef\]](#)
16. Copernicus Climate Change Service. *ERA5-Land Hourly Data from 1950 to Present*; Copernicus Climate Change Service: Reading, UK, 2019. [\[CrossRef\]](#)
17. Ali, S.; Bogarra, S.; Riaz, M.N.; Phyto, P.P.; Flynn, D.; Taha, A. From Time-Series to Hybrid Models: Advancements in Short-Term Load Forecasting Embracing Smart Grid Paradigm. *Appl. Sci.* **2024**, *14*, 4442. [\[CrossRef\]](#)
18. Huang, M.; Wei, Z.; Lin, Y. Forecasting-aided state estimation based on deep learning for hybrid AC/DC distribution systems. *Appl. Energy* **2022**, *306*, 118119. [\[CrossRef\]](#)
19. Said, D.; Elloumi, M.; Khoukhi, L. Cyber-Attack on P2P Energy Transaction Between Connected Electric Vehicles: A False Data Injection Detection Based Machine Learning Model. *IEEE Access* **2022**, *10*, 63640–63647. [\[CrossRef\]](#)
20. Huang, Y.; Mo, F.; Zhang, Z.; Li, C.; Li, K. Load Autoformer: A Transformer architecture for short-term load forecasting. In Proceedings of the 2023 IEEE Sustainable Power and Energy Conference (iSPEC), Chongqing, China, 28–30 November 2023; pp. 1–5. [\[CrossRef\]](#)
21. Si, C.; Xu, S.; Wan, C.; Chen, D.; Cui, W.; Zhao, J. Electric Load Clustering in Smart Grid: Methodologies, Applications, and Future Trends. *J. Mod. Power Syst. Clean Energy* **2021**, *9*, 237–252. [\[CrossRef\]](#)
22. Xiao, W.; Ding, T.; Chen, H. Load Characteristics Classification of Multiple Users Based on Time Series Clustering. In Proceedings of the 2021 11th International Conference on Power and Energy Systems (ICPES), Shanghai, China, 18–20 December 2021; pp. 569–574. [\[CrossRef\]](#)
23. Zhe, D.; Ya, W.; Chuanxiao, M.; Weihai, A. Load Clustering Method Based on Improved Density Peak and Gaussian Mixture Model. *J. Phys. Conf. Ser.* **2022**, *2171*, 012035. [\[CrossRef\]](#)
24. Cai, M.; Zheng, Y.; Peng, Z.; Huang, C.; Jiang, H. Research on load clustering algorithm based on variational autoencoder and hierarchical clustering. *PLoS ONE* **2024**, *19*, e0303977. [\[CrossRef\]](#)
25. Okereke, G.E.; Bali, M.C.; Okwueze, C.N.; Ukekwe, E.C.; Echezona, S.C.; Ugwu, C.I. K-means clustering of electricity consumers using time-domain features from smart meter data. *J. Electr. Syst. Inf. Technol.* **2023**, *10*, 2. [\[CrossRef\]](#)
26. Yu, K.; Cao, J.; Chen, X.; Yang, Z.; Gan, L. Residential load forecasting based on electricity consumption pattern clustering. *Front. Energy Res.* **2023**, *10*, 1–13. [\[CrossRef\]](#)
27. Ringnér, M. What is principal component analysis? *Nat. Biotechnol.* **2008**, *26*, 303–304. [\[CrossRef\]](#)
28. Huang, L.; Zhou, X.; Shi, L.; Gong, L. Time Series Feature Selection Method Based on Mutual Information. *Appl. Sci.* **2024**, *14*, 1960. [\[CrossRef\]](#)
29. Guo, Z.; O’Hanley, J.R.; Gibson, S. Predicting residential electricity consumption patterns based on smart meter and household data: A case study from the Republic of Ireland. *Util. Policy* **2022**, *79*, 101446. [\[CrossRef\]](#)
30. Wang, Y.; Chen, J.; Chen, X.; Zeng, X.; Kong, Y.; Sun, S.; Guo, Y.; Liu, Y. Short-Term Load Forecasting for Industrial Customers Based on TCN-LightGBM. *IEEE Trans. Power Syst.* **2020**, *36*, 1984–1997. [\[CrossRef\]](#)
31. Zhao, Y.; Ye, L.; Pinson, P.; Tang, Y.; Lu, P. Correlation-Constrained and Sparsity-Controlled Vector Autoregressive Model for Spatio-Temporal Wind Power Forecasting. *IEEE Trans. Power Syst.* **2018**, *33*, 5029–5040. [\[CrossRef\]](#)

32. Yao, X.; Fu, X.; Zong, C. Short-Term Load Forecasting Method Based on Feature Preference Strategy and LightGBM-XGboost. *IEEE Access* **2022**, *10*, 75257–75268. [CrossRef]
33. Liu, H.; Shen, X.; Tang, X.; Liu, J. Day-Ahead Electricity Price Probabilistic Forecasting Based on SHAP Feature Selection and LSTNet Quantile Regression. *Energies* **2023**, *16*, 5152. [CrossRef]
34. Triban, S.; Lawgali, A. Residential Short-Term Load Forecasting Based on CNN-LSTM with Consumer Behavior pattern. In Proceedings of the 2023 IEEE 11th International Conference on Systems and Control, ICSC 2023, Sousse, Tunisia, 18–20 December 2023; pp. 232–237. [CrossRef]
35. Kim, D.; Lee, D.; Nam, H.; Joo, S.K. Short-Term Load Forecasting for Commercial Building Using Convolutional Neural Network (CNN) and Long Short-Term Memory (LSTM) Network with Similar Day Selection Model. *J. Electr. Eng. Technol.* **2023**, *18*, 4001–4009. [CrossRef]
36. Wu, D.; Hur, K.; Xiao, Z. A GAN-Enhanced Ensemble Model for Energy Consumption Forecasting in Large Commercial Buildings. *IEEE Access* **2021**, *9*, 158820–158830. [CrossRef]
37. Liu, Y.; Liang, Z.; Li, X. Enhancing Short-Term Power Load Forecasting for Industrial and Commercial Buildings: A Hybrid Approach Using TimeGAN, CNN, and LSTM. *IEEE Open J. Ind. Electron. Soc.* **2023**, *4*, 451–462. [CrossRef]
38. Ceperic, E.; Ceperic, V.; Baric, A. A Strategy for Short-Term Load Forecasting by Support Vector Regression Machines. *Power Syst. IEEE Trans.* **2013**, *28*, 4356–4364. [CrossRef]
39. Shwartz-Ziv, R.; Armon, A. Tabular data: Deep learning is not all you need. *Inf. Fusion* **2022**, *81*, 84–90. [CrossRef]
40. Rai, S.; De, M. Analysis of classical and machine learning based short-term and mid-term load forecasting for smart grid. *Int. J. Sustain. Energy* **2021**, *40*, 821–839. [CrossRef]
41. Massana, J.; Pous, C.; Burgas, L.; Melendez, J.; Colomer, J. Short-term load forecasting for non-residential buildings contrasting artificial occupancy attributes. *Energy Build.* **2016**, *130*, 519–531. [CrossRef]
42. Park, S.; Jung, S.; Jung, S.; Rho, S.; Hwang, E. Sliding window-based LightGBM model for electric load forecasting using anomaly repair. *J. Supercomput.* **2021**, *77*, 12857–12878. [CrossRef]
43. Huang, A.; Zhou, J.; Cheng, T.; He, X.; Lv, J.; Ding, M. Short-term load forecasting for holidays based on similar days selecting and XGBoost model. In Proceedings of the 2023 IEEE 6th International Conference on Industrial Cyber-Physical Systems (ICPS), Wuhan, China, 8–11 May 2023; pp. 1–6. [CrossRef]
44. Lei, S.; Liang, X.; Wang, X.; Ding, J.; Ge, X.; Wang, F.; Feng, J. A Short-term Net Load Forecasting Method Based on Two-stage Feature Selection and LightGBM with Hyperparameter Auto-Tuning. In Proceedings of the 2023 IEEE/IAS 59th Industrial and Commercial Power Systems Technical Conference (I&CPS), Las Vegas, NV, USA, 21–25 May 2023; pp. 1–6. [CrossRef]
45. Wang, Y.; Bennani, I.L.; Liu, X.; Sun, M.; Zhou, Y. Electricity Consumer Characteristics Identification: A Federated Learning Approach. *IEEE Trans. Smart Grid* **2021**, *12*, 3637–3647. [CrossRef]
46. Dudek, G.; Gawlak, A.; Kornatka, M.; Szkutnik, J. Analysis of Smart Meter Data for Electricity Consumers. In Proceedings of the 2018 15th International Conference on the European Energy Market (EEM), Lodz, Poland, 27–29 June 2018; pp. 1–5. [CrossRef]
47. Zhou, Y.; Ding, Z.; Wen, Q.; Wang, Y. Robust Load Forecasting Towards Adversarial Attacks via Bayesian Learning. *IEEE Trans. Power Syst.* **2023**, *38*, 1445–1459. [CrossRef]
48. Ozaki, Y.; Tanigaki, Y.; Watanabe, S.; Nomura, M.; Onishi, M. Multiobjective Tree-Structured Parzen Estimator. *J. Artif. Intell. Res.* **2022**, *73*, 1209–1250. [CrossRef]
49. Chen, C.; Twycross, J.; Garibaldi, J.M. A new accuracy measure based on bounded relative error for time series forecasting. *PLoS ONE* **2017**, *12*, e0174202. [CrossRef]
50. Luo, X.; Oyedele, L.O. Forecasting building energy consumption: Adaptive long-short term memory neural networks driven by genetic algorithm. *Adv. Eng. Inform.* **2021**, *50*, 101357. [CrossRef]
51. Salata, F.; Ciancio, V.; Dell’Olmo, J.; Golasi, I.; Palusci, O.; Coppi, M. Effects of local conditions on the multi-variable and multi-objective energy optimization of residential buildings using genetic algorithms. *Appl. Energy* **2020**, *260*, 114289. [CrossRef]
52. Government of Spain. AEMET. 2025. Available online: <https://www.aemet.es/es/portada> (accessed on 2 January 2025).
53. Government of Catalonia. Statistical Institute of Catalonia. 2025. Available online: <https://www.idescat.cat/?lang=en> (accessed on 2 January 2025).
54. Xiao, J.W.; Liu, P.; Fang, H.; Liu, X.K.; Wang, Y.W. Short-Term Residential Load Forecasting with Baseline-Refinement Profiles and Bi-Attention Mechanism. *IEEE Trans. Smart Grid* **2024**, *15*, 1052–1062. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.