



Forecast-informed power load profiling: A novel approach

Óscar García Hinde ^{a,*}, Vanessa Gómez Verdejo ^a, Manel Martínez-Ramón ^b



^a Department of Signal Theory and Communications, Universidad Carlos III, Spain

^b Department of Electrical and Computer Engineering, University of New Mexico, United States of America

ARTICLE INFO

Keywords:

Power load
Forecasting
Profiling
Clustering
Machine learning
Probabilistic model

ABSTRACT

Power load forecasting plays a critical role in the context of electric supply optimization. The concept of load characterization and profiling has been used in the past as a valuable approach to improve forecasting performance as well as problem interpretability.

This paper proposes a novel, fully fledged theoretical framework for a joint probabilistic clustering and regression model, which is different from existing models that treat both processes independently. The clustering process is enhanced by simultaneously using the input data and the prediction targets during training. The model is thus capable of obtaining better clusters than other methods, leading to more informative data profiles, while maintaining or improving predictive performance.

Experiments have been conducted using aggregated load data from two U.S.A. regional transmission organizations, collected over 8 years. These experiments confirm that the proposed model achieves the goals set for interpretability and forecasting performance.

1. Introduction

During the past decade, the use of renewable energy has continued to grow steadily, representing 18.9% of the energy consumed in the European Union (European Commission, 2019) and the 11% in the United States (U.S. Energy Information Administration, 2019) in 2018. This trend is expected to continue in the future due to the need to reduce carbon emissions worldwide. The global tendency towards green energy has the effect of turning energy production into a non deterministic process, as it depends mainly on the local availability of solar radiation and wind. Not only do both these sources add uncertainty to the energy production problem, but they also introduce hard constraints due to their limited availability and seasonal nature. This makes the balance between energy availability and demand a complex problem that can only be tackled by demand side-management and energy storage strategies. Accurate power load profiling and forecasting models have thus become vital aids in this context.

In particular, the use of load profiling can mitigate the uncertainties introduced by the diversity of stakeholders in generation, transmission, and distribution agents (Espinoza et al., 2005). In general, the characterization and profiling of electricity consumption or modelling of common behaviours at user level and its applications have become an important research topic in the literature related to renewable energies.

In Liao and Niebur (2003) the profiling methodology is justified by the need for an accurate customer billing assessment. In Heunis and Herman (2002), load profiling is presented with the purpose of designing efficient low voltage distribution networks in residential areas in

South Africa. A similar approach is found in Jardini et al. (2000), where load profiles were characterized in São Paulo, Brazil. Examples of the use of these profiles are “transformer rating selection and management, load diversity evaluation and to determine the expected load profile in any preset point of the distribution network”, among others. A different application can be found in Carpaneto et al. (2003), where authors group customers based on their different profiles in order to design customized tariffs based on their energy usage preferences. The topic was discussed also in Abreu et al. (2012), where machine learning is used to detect habits in energy consumption. A similar application, but restricted to modelling energy consumption in buildings is presented in Diao et al. (2017), Csoknyai et al. (2019). Load models based on load profiles that mimic observed load are presented in Mammoli et al. (2019).

Regarding power load forecasting, many efforts have been made to design reliable models. Usually, load prediction models are categorized into short-term (from minutes to less than one week) and long-term (from more than one week to several years) (Stoll and Garver, 1989; Willis, 2004). For the purpose of short-term load forecast, traditional methods include linear prediction models that use historical time-series of the observed load. The main structures used in these models are moving-average (MA), auto-regressive moving-average (ARMA) or auto-regressive integrated moving-average (ARIMA) (Box et al., 2015). The most widely used parameter optimization algorithm consists of the minimization of the mean square error (MMSE) (see, e.g. Gross

* Corresponding author.

E-mail addresses: oghinde@tsc.uc3m.es (Ó. García Hinde), vanessa@tsc.uc3m.es (V. Gómez Verdejo), manel@unm.edu (M. Martínez-Ramón).

and Galiana (1987), (Papalexopoulos and Hesterberg, 1990)). Further attempts at using AR or ARMA regression include the use of support vector machines (SVM) (Burges et al., 1999) due to their robustness (Rojo-Álvarez et al., 2004) and ability to construct nonlinear versions (Martínez-Ramón et al., 2006). Some approaches include an SVM-MA (Nie et al., 2012) or an SVM-ARIMA structure (Li et al., 2003). Both of these works include the use of kernels (Shawe-Taylor and Cristianini, 2004) in order to attain nonlinear properties. Other approaches, such as (Mori and Ohmi, 2005), include the use of Gaussian processes (GP) (Rasmussen and Williams, 2006). While the MMSE criterion can easily be implemented online if the structure is linear (Haykin, 1996), SVMs and GPs require block training, which limits the number of data that can be used.

The use of neural networks has become a widespread practice in power-load forecast. Early works using the standard multilayer perceptron neural network include (Yoo and Pimmel, 1999; Srinivasan et al., 1999). More recent works use recurrent networks adapted for time series analysis called long short-term memory (LSTM) networks (Hochreiter and Schmidhuber, 1997; Kong et al., 2017), convolutional neural networks (Dong et al., 2017) and others. In general, due to their complexity, these methods show excellent performance, but they need to be trained with very large data-sets to obtain satisfactory results. These techniques also allow online and batch training.

Most of these models can and have been enhanced with the inclusion of multi-source data (which is also allowed by the use of kernels (Camps-Valls et al., 2008)). In particular, besides historical load time-series, the most used data source in power-load forecast consists of weather parameters such as outdoor temperature, humidity, solar radiation intensity, dew point temperature, wind speed, rainfall and others.

Tying profiling and forecasting together, it has been shown that good profiling can also improve forecasting accuracy when it is used as a form of data selection for model training. The main efforts in this topic take two different approaches. The first one is to directly produce an interpretation of the behaviour of the load time-series in time, space or across users. The second one focuses on the usage of clustering simply for the improvement of load forecasting accuracy. The main idea behind the second approach consists of constructing prediction models specialized in each one of the clusters. This way, it is expected that the possible nonlinear relationships between the predictor (input data) and the regressor (forecast load) can be locally approximated by linear functions or, at least, by less complex nonlinear structures. Clustering techniques are the most popular for both load characterization as well as joint profiling and forecasting (Yilmaz et al., 2019; Rajabi et al., 2020), due to their easy implementation and direct interpretation.

Clustering is a classic idea in machine learning and it was first presented as an application for load forecast in Willis et al. (1983), in which the authors use K-means clustering and linear regression for load forecasting. In fact, most of the works related to load profiling use clustering based on the K-means algorithm (see e.g. Bishop (2017)). For example, Rhodes et al. (2014) uses K-means to cluster load data and produce a detailed analysis of the results by grouping the clusters into yearly seasons. In Espinoza et al. (2005), the authors assume that the data fits a linear ARMA model and show that, when the data is clustered, this assumption produces better prediction results for a large number of substations of the Belgian power grid. Nonlinear approaches have also been applied together with clustering. In Fahiman et al. (2017) the authors use K-means to cluster aggregate data supplied by smart metres. The test data is then classified as belonging to one of such clusters and then processed through traditional neural networks to produce a forecast. In this work it is assumed that the input-output relationship is still nonlinear, but the clustering allows less complex structures.

K-means can be seen as a simplification of the Gaussian mixture model (GMM), which is trained with the well known Expectation Maximization algorithm (Murphy, 2012). An example of GMM used to

identify typical daily electricity usage profiles of multiple buildings can be found in Li et al. (2018). Authors actually use two clustering levels, the first one (intra-building) being a GMM and the second one (inter building) a hierarchical clustering.

Other authors, such as (Pan et al., 2015), take advantage of projections of the data into higher dimension Hilbert spaces through the use of kernelized versions of clustering techniques (Rojo-Álvarez et al., 2018).

In all previous papers where clustering is used to enhance load prediction, the authors construct a cascaded scheme with two distinct stages: a clustering algorithm followed by a battery of regression functions that are trained independently. Alternative models attempt to train both stages at the same time. A variant of the GMM for its application to regression is the approach known as Gaussian Process Mixture Model (Tresp, 2001). A version of this method was applied to power-load forecasting in Li et al. (2019). This model produces a probabilistic interpretation of the output but, since the input space is not clustered, it cannot provide a meaningful profiling of the data.

In this paper we present a new approach to simultaneous clustering and regression with the aim of providing good profiling characteristics while maintaining solid forecasting performance. The main theoretical novelty of the approach lies in the definition of a probabilistic model that performs a joint training of the clustering and regression stages. This results in a clustering of the input data that is actively informed by the forecasting process. The clusters achieved by the model will therefore offer a better, more informative representation of the problem than if they were a function of the input data alone.

The remainder of this paper is organized as follows:

- Section 2 fully defines the theoretical framework for the proposed model. It also provides a simple synthetic example to illustrate its capabilities.
- Section 3 gives a detailed description of the experiments that have been carried out, in which we apply the model to power load forecasting. It includes a thorough analysis of the resulting profiles and predictive function distribution as well as a performance evaluation. Several reference models have been included in the experiments for comparison.
- In Section 4 we present our conclusions and final thoughts on the experiments as well as a brief discussion of future lines of work to improve and expand the models capabilities.

The full model, as used in this study, is implemented in python and is available at the following GitHub repository: https://github.com/OGHinde/Clusterwise_Linear_Model.

2. The clusterwise linear model

Following on the idea of using an initial clustering phase to select groups of samples to later train separate regressors, we propose a unified probabilistic model that integrates regression into the clustering process. For this reason, we have called it the Clusterwise Linear Model (CWLM). We first assume that the observations, or input data, are generated by a standard Gaussian mixture model (GMM) with K components. Next, we consider that each component of the mixture model is associated to a linear regression model that generates the output targets. The novelty of this approach lies in the fact that both stages are coupled: not only does the input space clustering influence the linear regression on the output space, but also the regression process affects the overall clustering of the data.

Let us consider a regression problem defined by an observation data matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T$ and a target vector $\mathbf{y} = [y_1, \dots, y_N]^T$, where $\mathbf{x}_i \in \mathbb{R}^D$ is the i th observation and $y_i \in \mathbb{R}$ is its corresponding target.

The proposed model starts by considering that the data distribution can be approximated by a mixture of K Gaussians,

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} | \mu_k, \Sigma_k), \quad (1)$$




Fig. 1. Graphical representation for the CWLM. Shaded nodes indicate variables that are observed during the training phase, whereas non-shaded nodes indicate latent variables. The smaller solid nodes indicate deterministic model parameters.

where μ_k and Σ_k are, respectively, the mean and covariance matrix of the k th Gaussian component, and π_k is the prior probability that sample x has been generated by the k th Gaussian component.

Next, we introduce the following set of K linear models,

$$y = w_k^T x + \epsilon_k, \quad (2)$$

where w_k are the linear regression weights of the k th component, including the bias term,¹ and ϵ_k is assumed to be Gaussian noise with zero mean and variance β_k^{-1} .

Thus, given that observation x has been generated by the k th Gaussian component, its corresponding target value y will be generated by the k th linear model. Therefore the probability distribution for y becomes

$$p(y|x, \theta) = \mathcal{N}(y|w_k^T x, \beta_k^{-1}). \quad (3)$$

The mixture distribution for the target variables can therefore be stated as

$$p(y|x, \theta) = \sum_{k=1}^K \pi_k \mathcal{N}(y|w_k^T x, \beta_k^{-1}), \quad (4)$$

where θ includes all model parameters: π_k , which are the prior cluster probabilities; μ and Σ , which contain all component mean vectors μ_k and covariance matrices Σ_k for the input clustering stage; and w and β contain all regression weight vectors w_k and estimation noise precisions β_k .

2.1. Probabilistic representation

From a probabilistic standpoint, this model can be represented by the graph depicted in Fig. 1.

The model assumes that a set of latent variables $Z = [z_1, \dots, z_N]^T$ exists, where each $z_i = [z_{i,k}]_{k=1}^K$ is modelled such that only the k th entry of these vectors equals 1 and the rest is zero, indicating that x_i has been generated by the k th Gaussian mixture component and, consequently, y_i has been generated by the k th linear regressor. The prior distribution of these variables is defined as

$$p(z_{i,k} = 1) = \pi_k. \quad (5)$$

where $0 \leq \pi_k \leq 1$ and $\sum_{k=1}^K \pi_k = 1$.

The graphical model in Fig. 1 leads us to the following complete-data likelihood function

$$p(Z, X, y | \theta) = p(Z | \theta) p(X | Z, \theta) p(y | Z, X, \theta), \quad (6)$$

¹ x is considered to be extended with a constant term of value 1 to account for the bias term.




Fig. 2. Synthetic data-set designed to illustrate the capabilities of CWLM.

where

$$p(Z | \theta) = \prod_{i=1}^N \prod_{k=1}^K \pi_k^{z_{i,k}} \quad (7)$$

$$p(X | Z, \theta) = \prod_{i=1}^N \prod_{k=1}^K \mathcal{N}(x_i | \mu_k, \Sigma_k)^{z_{i,k}} \quad (8)$$

$$p(y | Z, X, \theta) = \prod_{i=1}^N \prod_{k=1}^K \mathcal{N}(y_i | w_k^T x_i, \beta_k^{-1})^{z_{i,k}}. \quad (9)$$

Therefore, the complete-data likelihood becomes

$$p(Z, X, y | \theta) = \prod_{i=1}^N \prod_{k=1}^K [\pi_k \mathcal{N}(x_i | \mu_k, \Sigma_k) \mathcal{N}(y_i | w_k^T x_i, \beta_k^{-1})]^{z_{i,k}}. \quad (10)$$

From here we can now compute the complete-data log likelihood,

$$\begin{aligned} \ln p(Z, X, y | \theta) &= \\ &= \sum_{i=1}^N \sum_{k=1}^K z_{i,k} [\ln \pi_k + \ln \mathcal{N}(x_i | \mu_k, \Sigma_k) + \ln \mathcal{N}(y_i | w_k^T x_i, \beta_k^{-1})] \end{aligned} \quad (11)$$

2.2. Model inference through expectation maximization

The Expectation Maximization algorithm (EM) is a common approach to find the optimal values for the parameters of a model that depends on latent variables. The EM algorithm iterates between two distinct stages: it first performs the expectation step (E-step), in which the current values of the model parameters are used to evaluate the posterior probabilities of the latent variables; it then applies the maximization step (M-step), in which these posterior probabilities are used to maximize the complete log-likelihood and update the values of the model parameters. These two steps are performed until a convergence criteria is met.

During the E-step the posterior distribution of the latent variables, known as the responsibilities, is computed as

$$\begin{aligned} \gamma(z_{i,k}) &= \mathbb{E}_{Z|X,Y} \{z_{i,k} | \theta\} = p(z_{i,k} | x_i, y_i, \theta) = \\ &= \frac{\pi_k \mathcal{N}(x_i | \mu_k, \Sigma_k) \mathcal{N}(y_i | w_k^T x_i, \beta_k^{-1})}{\sum_{k'=1}^K \pi_{k'} \mathcal{N}(x_i | \mu_{k'}, \Sigma_{k'}) \mathcal{N}(y_i | w_{k'}^T x_i, \beta_{k'}^{-1})}. \end{aligned} \quad (12)$$

Note that in this model the responsibilities depend on both $p(x | \mu, \Sigma)$ and $p(y | w^T x, \beta^{-1})$. Therefore the clustering process is informed by both the *input space* (the input variables contained in the observation data matrix, X) and the *output space* (the output variables contained in the labels, y).

During the M-step we update the model parameters by maximizing the expected value of the complete log-likelihood under the posterior of the latent variables. That is,

$$\theta^{\text{new}} = \arg \max_{\theta} \mathbb{E}_{Z|X,Y} \{ \ln p(Z, X, y | \theta) \}, \quad (13)$$

where

$$\begin{aligned} E_Z \{ \ln p(Z, X, y | \theta) \} &= \\ &= \sum_{i=1}^N \sum_{k=1}^K \gamma(z_{i,k}) [\ln \pi_k + \ln \mathcal{N}(x_i | \mu_k, \Sigma_k) + \ln \mathcal{N}(y_i | w_k^T x_i, \beta_k^{-1})]. \end{aligned} \quad (14)$$

Thus, the derivatives of (14) with respect to each parameter give us all the necessary update rules.

The cluster weight update rule becomes

$$\pi_k = \frac{1}{N} \sum_{i=1}^N \gamma(z_{i,k}), \quad (15)$$

considering the constraints $0 \leq \pi_k \leq 1$ and $\sum_k \pi_k = 1$.

The cluster mean and cluster covariance matrix update rules are

$$\mu_k = \frac{1}{N_k} \sum_{i=1}^N \gamma(z_{i,k}) x_i \quad (16)$$

and

$$\Sigma_k = \frac{1}{N_k} \sum_{i=1}^N \gamma(z_{i,k}) (x_i x_i^T - \mu_k \mu_k^T) \quad (17)$$

respectively, where N_k is the number of members belonging to component k , $N_k = \sum_i \gamma(z_{i,k})$.

Whereas the update rules for the regression weights and noise precision become

$$w_k = (X^T \Gamma_k X)^{-1} X^T \Gamma_k y \quad (18)$$

and

$$\beta_k^{-1} = \frac{1}{N_k} \sum_{i=1}^N \gamma(z_{i,k}) (y_i - w_k^T x_i)^2 \quad (19)$$

respectively, where Γ_k is defined as $\Gamma_k = \text{diag}(\{\gamma(z_{1,k}), \gamma(z_{2,k}), \dots, \gamma(z_{N,k})\})$.

2.3. Predictive distribution

The predictive distribution enables us to obtain an estimation of the output, \hat{y}^* , given a new test observation x^* . In this case, the output of the k th regressor is given by $y_k^* = w_k^T x^* + \epsilon_k$. Therefore, the probability distribution of y^* given the test observation together with the training data and the inferred model parameters is given by

$$p(y^* | x^*, X, y, \theta) = \sum_{k=1}^K p(z_k^* | x^*, \theta) p(y^* | z_k^* = 1, x^*, X, y, \theta), \quad (20)$$

where

$$p(z_k^* | x^*, \theta) = \frac{p(z_k^*, x^* | \theta)}{\sum_{k'=1}^K p(z_k^*, x^* | \theta)} = \frac{\pi_k \mathcal{N}(x^* | \mu_k, \Sigma_k)}{\sum_{k'=1}^K \pi_{k'} \mathcal{N}(x^* | \mu_{k'}, \Sigma_{k'})} \quad (21)$$

and takes a similar role to that of the responsibilities, γ_k , in (12). The difference lies in the fact that here we do not have access to the real value of the target and therefore we cannot incorporate this information to the cluster assignment.

Now, since

$$p(y_k^* | z_k^*, x^*, X, y, \theta) = \mathcal{N}(y_k^* | w_k^T x^*, \beta_k^{-1}), \quad (22)$$

we have that

$$p(y^* | x^*, X, y, \theta) = \sum_{k=1}^K p(z_k^* | x^*, \theta) \mathcal{N}(y_k^* | w_k^T x^*, \beta_k^{-1}) \quad (23)$$

We can now obtain an estimation of y^* as the expected value of $p(y^* | x^*, X, y, \theta)$

$$\begin{aligned} \hat{y}_{MSE}^* &= \mathbb{E}\{y^* | x^*, X, y, \theta\} \\ &= \sum_{k=1}^K p(z_k^* | x^*, \theta) \mathbb{E}\{y_k^* | x^*, X, y, \theta\} \\ &= \sum_{k=1}^K p(z_k^* | x^*, \theta) w_k^T x^* \end{aligned} \quad (24)$$

This estimation is essentially the sum of the outputs of the K regressors weighted by $p(z_k^* | x^*, \theta)$, which translates to the minimum square error (MSE) estimator. The maximum a posteriori (MAP) estimator, \hat{y}_{MAP}^* can also be implemented by using the output of the regressor associated to the highest value of $p(z_k^* | x^*, \theta)$ as the estimated value for y^* .

2.4. Model extensions

We now present some minor modifications to the standard version of CWIM that can be easily applied to obtain a more powerful, robust and expressive model.

Regularization term. An L^2 regularization (Hoerl and Kennard, 1970) on the regression weights can be included in the model with barely any modifications to the algorithm. This regularization term is introduced to the model in (18), which now becomes

$$w_k = (X^T \Gamma_k X \beta_k + \eta I)^{-1} X^T \Gamma_k \beta_k y. \quad (25)$$

The free parameter η acts simply as a regularization constant, and must be cross-validated to determine its optimal value without over-fitting the training data. This result is akin to assuming a Gaussian prior on the regression weights, where $p(w|\eta) = \mathcal{N}(w|0, \eta^{-1}I)$.

Multi-output prediction. The model can be easily extended to perform basic multi-output prediction, provided we assume complete independence of the output variables among themselves. In this case, the target vector y becomes target matrix $Y = [y_1, \dots, y_N]^T$, where $y_i \in \mathbb{R}^T$ contains the T target variables for the i th observation. Note that in this case the number of responsibilities computed in the E-step grows linearly with T (see (12)), meaning that the number of matrix inversions performed by the regression weight update rule (18) also grows linearly with T . This leaves us with per-task responsibilities

$$\gamma^{(i)}(z_{i,k}^{(i)}) = \frac{p(x_i, y_i^{(i)}, z_{i,k} | \theta)}{p(x_i, y_i^{(i)} | \theta)} \quad (26)$$

and per-task regression weight update rule

$$w_k^{(i)} = (X^T \Gamma_k^{(i)} X)^{-1} X^T \Gamma_k^{(i)} y^{(i)}. \quad (27)$$

This severe increase in computational complexity can be avoided by averaging the responsibilities over all output targets for each observation,

$$\bar{\gamma}_{i,k} = \frac{1}{T} \sum_{i=1}^T \gamma^{(i)}(z_{i,k}^{(i)}) \quad (28)$$

resulting in a single regression weight update rule in which the new averaged responsibilities $\bar{\gamma}_{i,k}$ are arranged in matrix $\bar{\Gamma}_k = \text{diag}(\{\bar{\gamma}_{1,k}, \bar{\gamma}_{2,k}, \dots, \bar{\gamma}_{N,k}\})$.

Multiple input-space views. Another modification of the model allows us to use different characterizations or views of the data for the input mixture of Gaussians and the output linear regressors, provided there is a one to one correspondence between the samples in each view. By defining two distinct input data matrices, $X^1 \in \mathbb{R}^{(N,D1)}$ and $X^2 \in \mathbb{R}^{(N,D2)}$, we can easily reformulate the M-Step and E-Step equations by replacing all appearances of x_i with either x_i^1 or x_i^2 appropriately. As we shall see in Section 3, this multi-view approach can prove to be very useful when we wish to exploit different expressions of the input data in the Gaussian mixture and the linear regression portions of the model.

2.5. Model capabilities

To illustrate the capabilities of the CWLM we have generated a simple synthetic data-set with one-dimensional input and output spaces so we can visually analyse the problem. As can be seen in Fig. 2, the data is arranged in three distinct groups, each associated to a different cluster and regressor, whose weights, bias terms and observational noise have been set randomly.




Fig. 3. Model training and evaluation on the synthetic data-set. The subfigures on the top row show the results of fitting to the training set a Ridge Regression model (a), a K-means + Ridge Regression model (b) and the CWLM model (c), respectively. The subfigures on the bottom row show the predictions on a test set for the three models. Subfigure (f) also depicts the predictive distribution contour-plot for the CWLM model.




Fig. 4. Power Load prediction using the CWLM algorithm: forecast vs real values for a randomly selected week in the ISO test partition. Also shown is the probability density map for the forecast values.

We have fitted three models to this synthetic data-set: a ridge regression model, a K-Means clustering model with 3 components feeding 3 separate ridge regression algorithms, and the CWLM model described in Section 2.

Fig. 3 summarizes the results. It is obvious that, while Ridge Regression does its best to fit a linear model to the data, it fails to capture any of its structure. Meanwhile, the K-Means + Ridge Regression approach does a little better at acknowledging the complex nature of the data, but it still fails to provide an adequate description of the underlying structure due to its inability to use the information contained in the training targets. On the other hand, the CWLM model manages to correctly identify the three clusters and accurately estimate the values of the weights and bias terms. Given the mixed distribution nature of the model, and to the best of our knowledge, an analytical expression for the confidence intervals of the prediction cannot be explicitly obtained. However, the predictive model (see (23)) does provide a full probability distribution for the test target predictions. This can be used as an intuitive indicator of the confidence in the predictions for a given region of the output space. This can be seen in Sub Fig. 3(f), in which predictions that fall within areas delimited by strong contours will offer

higher confidence levels. We will illustrate the usefulness of this notion in Section 3.5 in the context of power load forecasting.

3. CWLM applied to power load forecasting

The goal of this section is to evaluate the ability of the CWLM to gain valuable insight into the structure of the data in the context of power load profiling, while achieving competitive forecasting performance scores. For this purpose we have worked with power load data-sets belonging to two North American Regional Transmission Organizations (RTO).

3.1. Data-set description

ISO New England (ISO) serves the states of Massachusetts, Connecticut, Maine, New Hampshire, Vermont and Rhode Island. It provides an online repository of historical power load data. We have used aggregated data from all its member utilities, spanning from January 2011 to December 2018. The data consists of hourly samples arranged as a time-series. The corresponding hourly measurements of ambient and




Fig. 5. Multi-view CWLM profile visualization for the ISO New England data set.

dew-point temperatures were included as input variables to enhance the overall performance of the models. This meteorological data was obtained from the public NOAA repository.

PJM Interconnection LLC (PJM) serves all or part of Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, New Jersey, North Carolina, Ohio, Pennsylvania, Tennessee, Virginia, West Virginia, and the District of Columbia. As with ISO New England, PJM provides a public repository that includes data from all their partner utilities, with each utility serving a distinct zone. For this experiment we have

selected four different zones: East Kentucky Power Cooperative (EKPC), Dayton Power and Light Company (DAY), Pennsylvania Electric Company (PN) and Commonwealth Edison Company (CE). In this case we used data from 2014 to 2017. Again, we are using hourly samples. Meteorological data was not available in this case.

For both data-sets we decided to focus on a daily structuring of the data to find out if specific daily behaviour patterns can be automatically identified and if these patterns can be exploited to improve




Fig. 6. K-Means + Ridge Regression profile visualization for the ISO New England data-set.

the accuracy of our forecast. To achieve this we rearranged the yearly time-series into successive 24 h time-series.

In the case of the ISO data, the ambient and dew-point temperature data are appended as input variables. Therefore each input sample for the ISO experiment is characterized by 72 variables: 24 successive power load values, 24 successive temperature values and 24 successive dew-point values.

In both experiments, the prediction targets for each daily sample are the 24 h of the following day. We are therefore using the multi-target version of the model as described in Section 2.4.

Note that, since the main goal of this study is to gain interpretability, we have prioritized data that was relevant to this task. Specifically, we have focused on years that were as close to the test set as possible and that had the most complete daily samples, to ensure that seasonal




Fig. 7. CWLM model visualization for the PJM data-set - EKPC.

Table 1
ISO data-set - Test performance for all models.

Model	Clusters	MAPE	R ²	RMSE	MAE
RR	–	4.53	0.81	661.1	135.7
SVR	–	5.14	0.79	654.2	148.5
KM-Reg	4	4.36	0.83	626.8	129.8
GMM-Reg	6	4.37	0.83	621.3	129.5
CWLM	6	4.34	0.84	620.5	128.8
Multi-View KM-Reg	6	4.36	0.83	626.8	129.7
Multi-View GMM-Reg	4	4.43	0.82	634.8	131.7
Multi-View CWLM	4	4.26	0.85	619.6	127.7

Table 2
PJM data-set - Test performance for all models.

Models	Clusters	MAPE	R ²	RMSE	MAE
Ridge Regression	–	5.97	0.978	1477.5	236.9
SVR	–	5.48	0.987	1392.5	214.9
KM-Reg	10	5.97	0.987	1467.0	233.2
GMM-Reg	7	5.70	0.989	1424.8	221.7
CWLM	13	5.49	0.990	1385.6	214.7

and daily patterns were as best represented as possible. This results in smaller data-sets than those used in other studies.

3.2. Baseline forecast models

In order to gauge the performance of the CWLM algorithm described in Section 2, we have established two distinct baseline model families. The first consists of standard models which do not feature any clustering components. The second consists of models that do introduce a clustering stage similar to that of the CWLM.

Beginning with the standard models, we first introduce the Ridge Regression (RR) algorithm (Golub et al., 1999). This model was chosen in order to have a linear regression baseline for reference. The second standard model is the Support Vector Machine adapted to regression (SVR) using the epsilon method (Drucker et al., 1997). A Gaussian

kernel was chosen to provide a non-linear baseline model. Since SVR cannot perform multitarget regression, 24 SVR models were trained simultaneously while sharing the same parameter values.

As for the clustering baseline models, the first is the K-Means + Ridge Regression (KM-Reg) algorithm, a simple approach to clustered regression in which K-Means is used to separate the training data into K different groups, which are then fed to K independent ridge regression models. Note that this approach allows us to use multiple input-space views in the same manner that is described in 2.4 for CWLM. This model is also used in the synthetic example from Section 2. The second of the clustering baselines is the Gaussian mixture model + Ridge Regression (GMM-Reg) algorithm. This is a more nuanced approach to clustered regression in which clustering is achieved by applying a Gaussian Mixture Model (GMM) to the input data and then training K ridge regression models with the full data-set, but using the likelihoods from the GMM model as sample weights. The final output is therefore the sum of the K regression outputs weighted by the GMM likelihoods for each sample and cluster. As is the case for the CWLM and KM-Reg approaches, this model admits multiple input-space views.

3.3. Experimental setup

The ISO data-set was split up into a training partition, containing all data from January 2011 to December 2016; a validation partition, containing all data from the year 2017, used to optimize all model hyperparameters; and a test partition, containing all data from the year 2018, used to compute the performance metrics.

As for the partitioning of the PJM data-set, samples from 2014 to 2016 were used for both training and validation. In this case, the validation set consists of randomly and uniformly selected samples, with a size equal to 20% of the total number of data-points from this time period. The rest of the samples from the 2014–2016 range were used as the training partition. All data from 2017 was used as the test partition.

The optimal values for all the relevant parameters were obtained after a thorough exploration using the validation partition:

- The regularization parameter, λ , for the Ridge Regression, KM-Reg and GMM-Reg models, was explored in the range $\lambda \in [10^{-4}, 10^2]$.
- The regularization parameter, η , for the CWLM algorithm, was explored in the range $\eta \in [10^{-4}, 10^2]$.
- The Gaussian kernel parameter, γ , for the SVR model was explored in the range $\gamma \in [10^{-4}, 10^3]$.
- The number of clusters for the CWLM algorithm as well as the KM-Reg and GMM-Reg models was explored in the range $K \in [2, 40]$.

All data was normalized row-wise so that every time-series lay between the values of 0 and 1. This row-wise normalization was applied in order to retain the shape of the time-series while ensuring an adequate scaling of the data.

For the ISO data-set, we applied both the standard and the multi-view versions of the CWLM, KM-Reg and GMM-Reg algorithms. The standard models received the full input data matrix, containing both the power load and meteorological data. For the multi-view models, we fed the full input-data matrix to the regression stage of the algorithms, whereas the input clustering stage of the algorithms only received the power load portion of the input-data.

Since the PJM data-set is not augmented with meteorological data, only the standard version of the clustering algorithms was used.

To evaluate the performance of the models described above, we have chosen to use the following metrics, where y_n is the true target value for the n th test sample, \hat{y}_n is the predicted target value for the n th test sample, $\bar{y} = \frac{1}{N_{test}} \sum_{n=1}^{N_{test}} y_n$ is the average of the true test target values and N_{test} is the size of the test set:

- Mean Absolute Percentage Error (*MAPE*), defined as:

$$MAPE = \frac{100\%}{N_{test}} \sum_{n=1}^{N_{test}} \left| \frac{y_n - \hat{y}_n}{y_n} \right| \quad (29)$$

A lower *MAPE* implies better performance.

- Root Mean Squared Error (*RMSE*) defined as:

$$RMSE = \left(\frac{1}{N_{test}} \sum_{n=1}^{N_{test}} (y_n - \hat{y}_n)^2 \right)^{1/2} \quad (30)$$

- Mean Absolute Error (*MAE*), defined as:

$$MAE = \frac{1}{N_{test}} \sum_{n=1}^{N_{test}} |y_n - \hat{y}_n| \quad (31)$$

- Coefficient of Determination (R^2), defined as:

$$R^2 = 1 - \frac{\sum_{n=1}^{N_{test}} (y_n - \hat{y}_n)^2}{\sum_{n=1}^{N_{test}} (y_n - \bar{y})^2} \quad (32)$$

Where the best possible score is 1.0 and negative scores are possible (because the model can be arbitrarily worse).

3.4. Performance analysis

Tables 1 and 2 summarize the performance results for the ISO and PJM data-sets for all methods under study, with their respective parameters validated according to the procedure described in Section 3.3.

While the differences in performance are not large, the clustered regression models do gain an edge over Ridge Regression. Overall, the CWLM algorithm comes on top in all metrics for both data-sets, with the multi-view version offering further performance gains in the case of the ISO data-set.

Of interest is the performance difference of the SVR model in both data-sets: in the case of ISO, SVR achieves the worst scores while for PJM it comes very closely tied to CWLM in *MAPE* and *MAE*, although CWLM achieves slightly better results in R^2 and *RMSE*. This suggests

that the forecast problem for the ISO data is far more linear. In such a scenario, non linear models like the SVR with a Gaussian kernel will tend to overfit. Meanwhile CWLM is able to exploit the linearity of the ISO data to generalize better, while offering the interpretability gains which will be described in Section 3.6.

3.5. Advantages of the predictive model

Fig. 4 shows the real and forecast power load values for a randomly selected week from the ISO test partition. At the same time, it maps the probability density from the predictive distribution defined in Eq. (23). Darker regions indicate high probability density, which translates into higher predictive confidence. Lighter regions indicate a more diffuse probability density, therefore suggesting a lower predictive confidence.

The figure shows consistently high predictive confidence during the night and early mornings. This is backed by how closely the forecast curve follows the real curve during these time periods. On the other hand, the probability density disperses during the busier times of the day and the afternoon. Again, this is reflected by a worsening of the quality of the forecast values.

As the predictive function allows us to visually establish an intuition of the confidence in our predictions, our model can be used to determine times of the day during which a higher volatility in the demand is to be expected, which in turn will influence the resource allocation strategies that need to be put in place.

3.6. Interpretability analysis: daily load profiling

The main goal of this study is to evaluate the interpretability gains for power load profiling provided by the CWLM algorithm. For this purpose, we focus on the nature of the data assigned to each cluster.

Beginning with the ISO data-set, Figs. 5 and 6 allow us to evaluate the profiling capability of the CWLM algorithm in the multi-view configuration and of the basic KM-Reg approach.² Each figure shows us three different visualizations for each of the four clusters: the first graph represents the cluster centroid together with all its member samples; the second one is a histogram representing the frequency with which each day of the week is represented in the cluster; finally we represent a second histogram showing the frequency with which each month of the year is represented in the cluster.

Fig. 5 clearly shows that there are strong patterns that have been automatically identified by the multi-view CWLM algorithm. The first and second clusters are dominated by days belonging to the colder months of the year, whereas the third and fourth clusters mostly contain warmer days. Furthermore, the first and third clusters are very highly populated by weekend days, whereas cluster two and four are mostly composed of weekdays.

On the other hand, Fig. 6 shows that the nature of the clusters selected by KM-Reg is far less clear cut. Some seasonal and daily patterns can be identified in these graphs, but we argue that the overall insight is not as "sharp" as that of the multi-view CWLM, which has very explicitly exploited seasonal and daily patterns, identifying interesting behavioural profiles in the data.

An analysis of the results obtained by the CWLM for the PJM data-set shows that most clusters are associated to specific utilities. For instance clusters 3 and 4 concentrate data from EKPC, whereas clusters 2 and 7 are mostly comprised of data from CE. We have therefore split our analysis between Figs. 7–9, each corresponding to a utility-specific set of clusters. Each figure features three subfigures, (a) to (c): the first shows the number of members for a given utility in each cluster, justifying the utility profile for each cluster; the second reflects the

² For simplicity's sake and to save space, we cannot include figures for all the models under evaluation. We have chosen to only represent the multi-view CWLM and the KM-Reg algorithm, since they both selected the same number of clusters.




Fig. 8. CWLM model visualization for the PJM data-set - CE.




Fig. 9. CWLM model visualization for the PJM data-set - PN DAY.

proportion of weekdays and weekends per cluster; the third subfigure shows how represented each month of the year is in each cluster. From these figures we can see that:

- Clusters associated to EKPC (Fig. 7) do not seem to distinguish between weekdays and weekends. They do however show a clear seasonal dependence: cluster 3 is composed mostly of spring and

autumn months, Cluster 4 is dominated by summer months and clusters 12 and 13 show a strong presence of winter months.

- In the case of clusters associated with CE (Fig. 8), we do find a strong distinction between weekdays and weekends as well as a seasonal component. For instance, cluster 2 is dominated by the warmer months and, as is to be expected due to the presence of summer vacations, weekends and weekdays appear to be mixed.




Fig. 10. GMM-Reg model visualization for the PJM data-set.




Fig. 11. KM-Reg model visualization for the PJM data-set.

Clusters 7 and 8 clearly show weekday behaviours for the rest of the year. Finally, cluster 9 models weekends throughout the year.

- Interestingly, Fig. 9 shows that utilities PN and DAY tend to share similar behaviours and therefore appear grouped in the same clusters. Going into more detail, clusters 1 and 11 again show hot summer months and spring-autumn months, respectively. Cluster 2 also models hotter months and cluster 5 focuses on summer and autumn weekends. Finally, clusters 6 and 10 are associated to cold months.

Visualizations for the PJM data-set results of the GMM-Reg and the KM-Reg models can be seen in Figs. 10 and 11, respectively. These figures represent the number of samples from each utility assigned to each cluster. Here we can see that GMM-Reg tends to also discriminate utilities quite well, particularly in the case of EKPC and PN. However, its lack of predictive performance (Table 2) suggests that these clusters are not as relevant as those obtained by CWLM. Finally, KM-Reg is incapable of even discriminating utilities in a meaningful way, as well as having the worst performance out of all the clustering models.

4. Conclusions

In this paper we propose a novel theoretical framework leading to a probabilistic model that features simultaneous data clustering and regression. The clustering of data samples is informed by the regression process, which enables the model to achieve a better characterization of the underlying nature of the data. Therefore, we obtain better, more informative clusters while maintaining or improving predictive performance. We propose that this model can improve problem interpretability in the context of day-ahead electric power-load forecasting, by generating useful and insightful daily load profiles.

At the same time, the model presents a probabilistic predictive function that is capable of providing an intuition for the forecasting confidence, which in turn can be used to improve the interpretation of

the predicted power-load values. We suggest that this could be of great use in the context of power-grid management and efficiency.

Experimental results in the context of power load forecasting applied to data from two major Regional Transmission Organizations confirm the usefulness of our model in terms of interpretability, as it is shown to generate insightful load profiles, while obtaining competitive forecasting performance when compared to other prediction models. In both data-bases, the automatically generated profiles reflect the relevance of regional and seasonal patterns, as well as the influence of weekdays and weekends: for the first data-set, very clear seasonal and daily patterns were obtained, with four strong clusters that segregated data into weekday and weekends during warm months and cold months; for the second data-set, strong regional separation was automatically achieved, with dedicated clusters for specific utilities at different times of the year. Meanwhile the predictive function points to very high confidence during the early hours of the morning, with said confidence dropping during the busiest hours of the day.

Two key improvements to CWLM can be introduced as future work. First, a fully Bayesian approach can be formulated in which the regression weights become a new latent variable with their own prior distribution. This can be solved using the Variational Inference approach known as Mean Field Approximation, which will result in a very similar formulation to the one presented in this work. The main advantage is that the need to validate the regularization term, η , would disappear, as it would become a part of the iterative model optimization algorithm. The second improvement we propose is the implementation of a complete multi-target model in which the possible correlations between the output variables are taken into account, while maintaining the same integration with the clustering process as the model presented in this paper.

CRediT authorship contribution statement

Oscar García Hinde: Conceptualization, Methodology, Software, Validation, Investigation, Writing - original draft, Writing - review &

editing. **Vanessa Gómez Verdejo:** Project administration, Conceptualization, Methodology, Resources, Funding acquisition, Supervision, Writing - review & editing. **Manel Martínez-Ramón:** Conceptualization, Methodology, Data curation, Funding acquisition, Supervision, Writing - review & editing.

Declaration of competing interest

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

Acknowledgements

This work is partially supported by the National Science Foundation EPSCoR Cooperative Agreement OIA-1757207 and the Spanish MINECO grants TEC2014-52289R and TEC2017-83838-R.

References

Abreu, Joana M., Pereira, Francisco Cámera, Ferri, Paulo, 2012. Using pattern recognition to identify habitual behavior in residential electricity consumption. *Energy Build.* 49, 479–487.

Bishop, Christopher M., 2017. Pattern Recognition and Machine Learning. Springer, pp. 4134–4141.

Box, George E.P., Jenkins, Gwilym M., Reinsel, Gregory C., Ljung, Greta M., 2015. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons.

Burges, Christopher J.C., Scholkopf, Bernhard, Smola, Alexander J., 1999. Advances in Kernel Methods: Support Vector Learning. MIT Press Cambridge, MA, USA.

Camps-Valls, Gustavo, Gómez-Chova, Luis, Muñoz-Marí, Jordi, Rojo-Álvarez, José Luis, Martínez-Ramón, Manel, 2008. Kernel-based framework for multitemporal and multisource remote sensing data classification and change detection. *IEEE Trans. Geosci. Remote Sens.* 46 (6), 1822–1835.

Carpaneto, Enrico, Chico, Gianfranco, Napoli, Roberto, Scutariu, Mircea, 2003. Customer classification by means of harmonic representation of distinguishing features. In: 2003 IEEE Bologna Power Tech. Conference Proceedings, vol. 3. IEEE.

Csoknya, Tamás, Legardeur, Jeremy, Akle, Audrey Abl, Horváth, Márk, 2019. Analysis of energy consumption profiles in residential buildings and impact assessment of a serious game on occupants' behavior. *Energy Build.* 196, 1–20.

Diao, Longquan, Sun, Yongjun, Chen, Zeyun, Chen, Jiayu, 2017. Modeling energy consumption in residential buildings: A bottom-up analysis based on occupant behavior pattern clustering and stochastic simulation. *Energy Build.* 147, 47–66.

Dong, Xishuang, Qian, Lijun, Huang, Lei, 2017. Short-term load forecasting in smart grid: A combined CNN and K-means clustering approach. In: 2017 IEEE International Conference on Big Data and Smart Computing. BigComp. IEEE, pp. 119–125.

Drucker, Harris, Burges, Christopher J.C., Kaufman, Linda, Smola, Alex J., Vapnik, Vladimir, 1997. Support vector regression machines. In: Advances in Neural Information Processing Systems, pp. 155–161.

Espinosa, Marcelo, Joye, Caroline, Belmans, Ronnie, De Moor, Bart, 2005. Short-term load forecasting, profile identification, and customer segmentation: a methodology based on periodic time series. *IEEE Trans. Power Syst.* 20 (3), 1622–1630.

Espinosa, M., Joye, C., Belmans, R., De Moor, B., 2005. Short-term load forecasting, profile identification, and customer segmentation: a methodology based on periodic time series. *IEEE Trans. Power Syst.* 20 (3), 1622–1630.

European Commission, 2019. Renewable energy statistics.

Fahimsa, Fateme, Erfani, Sarah M., Rajasegarar, Sutharshan, Palaniswami, Marimuthu, Leckie, Christopher, 2017. Improving load forecasting based on deep learning and K-shape clustering. In: 2017 International Joint Conference on Neural Networks. IJCNN. IEEE, pp. 4134–4141.

Golub, Gene H., Hansen, Per Christian, O'Leary, Dianne P., 1999. Tikhonov regularization and total least squares. *SIAM J. Matrix Anal. Appl.* 21 (1), 185–194.

Gross, George, Galiana, Francisco D., 1987. Short-term load forecasting. *Proc. IEEE* 75 (12), 1558–1573.

Haykin, Simon, 1996. Adaptive Filter Theory. Prentice-Hall, Inc.

Heunis, Schalk W., Herman, Ron, 2002. A probabilistic model for residential consumer loads. *IEEE Trans. Power Syst.* 17 (3), 621–625.

Hochreiter, Sepp, Schmidhuber, Jürgen, 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.

Hoerl, Arthur E., Kennard, Robert W., 1970. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12 (1), 55–67.

Jardini, José António, Tahan, Carlos M.V., Gouveia, Marcos R., Ahn, Se Un, Figueiredo, F.M., 2000. Daily load profiles for residential, commercial and industrial low voltage consumers. *IEEE Trans. Power Deliv.* 15 (1), 375–380.

Kong, Weicong, Dong, Zhao Yang, Jia, Youwei, Hill, David J., Xu, Yan, Zhang, Yuan, 2017. Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE Trans. Smart Grid* 10 (1), 841–851.

Li, Yuan-cheng, Fang, Ting-jian, Yu, Erkeng, 2003. Study of support vector machines for short-term load forecasting. *Proc. CSEE* 23 (6), 55–59.

Li, Kehua, Ma, Zhenjun, Robinson, Duane, Ma, Jun, 2018. Identification of typical building daily electricity usage profiles using Gaussian mixture model-based clustering and hierarchical clustering. *Appl. Energy* 231, 331–342.

Li, Ling-Ling, Sun, Jin, Wang, Ching-Hsin, Zhou, Ya-Tong, Lin, Kuo-Ping, 2019. Enhanced Gaussian process mixture model for short-term electric load forecasting. *Inform. Sci.* 477, 386–398.

Liao, Huawei, Niebur, Dagmar, 2003. Load profile estimation in electric transmission networks using independent component analysis. *IEEE Trans. Power Syst.* 18 (2), 707–715.

Mammoli, Andrea, Robinson, Matthew, Ayon, Victor, Martínez-Ramón, Manel, Chen, Chien-Fei, Abreu, Joana M., 2019. A behavior-centered framework for real-time control and load-shedding using aggregated residential energy resources in distribution microgrids. *Energy Build.* 198, 275–290.

Martínez-Ramón, Manel, Rojo-Álvarez, José Luis, Camps-Valls, Gustau, Muñoz-Marí, Jordi, Navia-Vázquez, Ángel, Soria-Olivas, Emilio, Figueiras-Vidal, Aníbal R., 2006. Support vector machines for nonlinear kernel ARMA system identification. *IEEE Trans. Neural Netw.* 17 (6), 1617–1622.

Mori, Hiroyuki, Ohmi, Masatatsu, 2005. Probabilistic short-term load forecasting with Gaussian processes. In: Proceedings of the 13th International Conference on Intelligent Systems Application To Power Systems. IEEE.

Murphy, Kevin P., 2012. Machine Learning: A Probabilistic Perspective. MIT press.

Nie, Hongzhan, Liu, Guohui, Liu, Xiaoman, Wang, Yong, 2012. Hybrid of ARIMA and SVMs for short-term load forecasting. *Energy Procedia* 16, 1455–1460.

Pan, E., Li, H., Song, L., Han, Z., 2015. Kernel-based non-parametric clustering for load profiling of big smart meter data. In: 2015 IEEE Wireless Communications and Networking Conference. WCNC, pp. 2251–2255.

Papalexopoulos, Alex D., Hesterberg, Timothy C., 1990. A regression-based approach to short-term system load forecasting. *IEEE Trans. Power Syst.* 5 (4), 1535–1547.

Rajabi, Amin, Eskandari, Mohsen, Ghadi, Mojtaba Jabbari, Li, Li, Zhang, Jiangfeng, Siano, Pierluigi, 2020. A comparative study of clustering techniques for electrical load pattern segmentation. *Renew. Sustain. Energy Rev.* 120, 109628.

Rasmussen, Carl Edward, Williams, Christopher K.I., 2006. Gaussian Processes for Machine Learning. MIT Press Cambridge, MA.

Rhodes, Joshua D., Cole, Wesley J., Upshaw, Charles R., Edgar, Thomas F., Webber, Michael E., 2014. Clustering analysis of residential electricity demand profiles. *Appl. Energy* 135, 461–471.

Rojo-Álvarez, José Luis, Martínez-Ramón, Manel, Marí, Jordi Muñoz, Camps-Valls, Gustavo, 2018. Digital Signal Processing with Kernel Methods. Wiley Online Library.

Rojo-Álvarez, José Luis, Martínez-Ramón, Manel, de Prado-Cumplido, Mario, Artés-Rodríguez, Antonio, Figueiras-Vidal, Aníbal R., 2004. Support vector method for robust ARMA system identification. *IEEE Trans. Signal Process.* 52 (1), 155–164.

Shawe-Taylor, John, Cristianini, Nello, 2004. Kernel methods for pattern analysis. Cambridge university press.

Srinivasan, Dipali, Tan, Swee Sien, Cheng, C.S., Chan, Eng Kit, 1999. Parallel neural network-fuzzy expert system strategy for short-term load forecasting: System implementation and performance evaluation. *IEEE Trans. Power Syst.* 14 (3), 1100–1106.

Stoll, Harry G., Garver, Leonard J., 1989. Least-Cost Electric Utility Planning. J. Wiley.

Tresp, Volker, 2001. Mixtures of Gaussian processes. In: Advances in Neural Information Processing Systems, pp. 654–660.

U.S. Energy Information Administration, 2019. Renewable sources.

Willis, H. Lee, 2004. Power Distribution Planning Reference Book. CRC Press.

Willis, H.L., Schauer, A.E., Northcote-green, J.E.D., Väistö, T.D., 1983. Forecasting distribution system loads using curve shape clustering. *IEEE Trans. Power Appar. Syst. PAS-102* (4), 893–901.

Yilmaz, Selin, Chambers, Jonathan, Patel, Martin Kumar, 2019. Comparison of clustering approaches for domestic electricity load profile characterisation-implications for demand side management. *Energy* 180, 665–677.

Yoo, Hyeonjoong, Pimmler, Russell L., 1999. Short term load forecasting using a self-supervised adaptive neural network. *IEEE Trans. Power Syst.* 14 (2), 779–784.