

Article

Smart Meter Forecasting from One Minute to One Year Horizons

Luca Massidda *  and Marino Marrocu

CRS4, Center for Advanced Studies, Research and Development in Sardinia, loc. Piscina Manna ed. 1, 09010 Pula (CA), Italy; marino.marrocu@crs4.it

* Correspondence: luca.massidda@crs4.it; Tel.: +39-070-9250-285

Received: 20 November 2018; Accepted: 13 December 2018; Published: 18 December 2018



Abstract: The ability to predict consumption is an essential tool for the management of a power distribution network. The availability of an advanced metering infrastructure through smart meters makes it possible to produce consumption forecasts down to the level of the individual user and to introduce intelligence and control at every level of the grid. While aggregate load forecasting is a mature technology, single user forecasting is a more difficult problem to address due to the multiple factors affecting consumption, which are not always easily predictable. This work presents a hybrid machine learning methodology based on random forest (RF) and linear regression (LR) for the deterministic and probabilistic forecast of household consumption at different time horizons and resolutions. The approach is based on the separation of long term effects (RF) from short term ones (LR), producing deterministic and probabilistic forecasts. The proposed procedure is applied to a public dataset, achieving a deterministic forecast accuracy much higher than other methodologies, in all scenarios analyzed. This covers horizons of forecast from one minute to one year, and highlights the great added value provided by probabilistic forecasting.

Keywords: load forecasting; smart meter; time series forecasting; machine learning; energy prediction

1. Introduction

New challenges for the efficient management of the power distribution system are posed by the ongoing deregulation of power distribution in many countries, the growing distributed power generation from renewable sources, the introduction of distributed energy storage systems and the increasing diffusion of electric vehicles. Smart grids offer a viable solution through unprecedented flexibility in energy generation and distribution [1].

Over the last decade, a growing number of smart meters have been installed worldwide. These, together with the communication and data management network, constitute the advanced metering infrastructure (AMI) that will play a fundamental role in electrical distribution systems, by recording user load profiles, enabling two-way communication between the user and the distributor and allowing smarter systems for the management of energy resources [2].

How to use the volume of data from smart meters to promote and improve efficiency and sustainability of demand has become a major research topic worldwide [3]. Control decisions for the smart grid should be made continuously at both the aggregate and granular levels. To achieve this and ensure network reliability, the ability to predict future demand is of paramount importance.

Load forecasts have been widely used by the electrical sector. Power distribution companies rely on forecasts with different time horizons to support both system operability and planning. Retail electricity suppliers are making pricing, procurement and hedging decisions based largely on the expected load of their customers.

In recent years, moreover, there has been a steady trend towards the electrification of energy consumption, linked to the need for greater use of renewable energy sources, which results in a load profile increasingly characterized by the presence of peaks in consumption due to human behavior, that can lead to problems for electricity providers [4]. The presence of peaks in consumption and generation linked to human behaviour and renewable sources leads to the research of mitigation techniques, for example using a variable pricing system that pushes users to plan as much as possible the use of energy resources, in addition to the use of distributed storage systems, as tested in the project H2020 NETfficient, in which the techniques described in this work have found application [5].

The value that smart meters bring to load forecasting and more generally to the energy distribution system is manifold. First, smart meters enable distribution companies and electricity retailers to better understand and predict the load on a single house or building. Secondly, the high granularity of the load data provided by smart meters offers great potential for improving aggregate forecast accuracy [3]. In addition, the forecast of energy consumption over time allows property and building managers to plan energy consumption over time, shifting energy use to off-peak periods, improving energy purchase plans, and allowing them to assess their consumption habits, identifying possible margins.

Since the household loads are more volatile than the aggregate load, the higher the load level, the more uniform the profile is and the less uncertain the forecast is. The energy forecast at the smart meter level is not a trivial problem, as it depends on the complexity of the energetic behavior of the building, in turn related to climatic conditions and to the operation of lighting systems and HVAC (heating, venting and air conditioning), but especially for the difficulty in predicting the behavior of the occupants, influenced by multiple social factors [6,7].

To address the problem of load forecasting at the level of smart meters, the research community has attempted different approaches, from adapting techniques already widely used for aggregate load forecasting, to developing new techniques or using a combination of the same [3]. Methods such as the semi parametric additive model [8], exponential smoothing [9], and classical seasonal time series methods have been applied to load forecast at the building level [10], as well as methods based on artificial neural networks (ANN) and support vector machines (SVM) [11]. A 2012 [12] study compared several existing techniques, including linear regression (LR) as well as different types of ANNs and SVMs on two datasets: one for two commercial buildings and the other for three residential homes. The results showed that the techniques used could provide reliable forecasts in the first case but not in the second, because of the greater variability of the load. In [13] the load forecast is studied both at the building level and at the state and provincial level through a self-recurrent wavelet neural network.

The recent trend is to use deep learning techniques, with recursive or convolutional or hybrid neural networks. The conditional restricted Boltzmann machine (CRBM) and factored conditional restricted Boltzmann machine (FCRBM) have been evaluated in [6] to estimate the energy consumption of a household. The FCRBM achieves the highest accuracy in load forecasting compared to ANN, RNN (Recurrent Neural Network), SVM and CRBM. Different resolutions ranging from one minute to one week have been tested. The same dataset has been analyzed in two successive works in which the effectiveness of recursive networks of the type long short term memory (LSTM) has been investigated both in the standard form and in the form sequence to sequence (S2S) [14] and in which the accuracy of convolutional neural networks (CNNs) [15] has been evaluated, obtaining comparable or superior results to those obtained with the FCRBM algorithm.

Due to the high variability of smart meter measurements, it may be necessary to perform probabilistic forecasting in operational practice. The interested reader can find in [16] a review of the different methods proposed in this area for aggregate load forecasting. Probabilistic load forecasting has also been carried out on individual load profiles; in [17] a method combining gradient boosting (GB) and quantile regression has been proposed to quantify uncertainty and generate probabilistic forecasts, while in [18] the conditional kernel density (CKD) method have been tested. Recently, a point and probabilistic forecast of the load for 100 low voltage (LV) feeders has been conducted in [19]

comparing several methods such as Holt-Winters-Taylor seasonal exponential smoothing, kernel density estimation, seasonal linear regression and two autoregressive (AR) methods.

The aim of this work is to implement a forecasting procedure for household consumption using only smart-meter load data, with different time horizons and producing both deterministic and probabilistic forecasts. The most recent research sees an increased use of deep learning techniques for load forecasting, their real effectiveness for forecasting time series with seasonality is however at least partially questioned on the basis of the results actually obtained with these methods when compared with classical statistical methods. S. Makridakis in a recent article [20] analyses the performance of the methods proposed for M3 competition, highlighting how statistical methods allow obtaining, on average, more accurate forecasts and with a lower computational cost. It also provides suggestions on how to exploit the undoubted potential of machine learning (ML) techniques also in the context of time series forecasting. In particular, it highlights, among other things, the need for data pre-treatment, applying transformations and procedures of de-trending and de-seasonalization that allow obtaining a stationary signal, and the accurate assessment of the risk of over-fitting for ML procedures.

In our case, these indications translate into a separation between the long-term components of the signal, associated with the trend and seasonality, and the short-term components linked to stochastic variations in the average trend of consumption of a consumer, which we will treat as two distinct problems in a hybrid methodology. The method we propose in fact provides for an estimate of the average long-term behavior of users, to which a short-term auto-regressive forecast is superimposed, linked to the deviation of the previous estimate from the most recent consumption measures. A similar hybrid approach, in which the effects of long and short term are separated, has been recently used by S. Smyl in the winning procedure of the recent M4 competition [21].

The proposed methodology is implemented using ML techniques with limited computational requirements in order to allow an implementation also in low cost devices installed directly at the user's household. In particular, we will use the random forest (RF) technique for long-term forecasting and a simple linear regression for short-term forecasting. We will also realize a probabilistic forecast through a simple persistence of the distribution of forecast errors measured in the training dataset.

This work will highlight:

- The convenience of a hybrid approach, which separates long-term and short-term effects for load forecasting when using machine learning techniques;
- The effectiveness of the proposed procedure for predicting smart meter loads;
- The relative contribution of these two components to the accuracy of forecasting;
- The importance and added value of a probabilistic forecast for household load prediction.

The application of the proposed procedure to a public dataset already studied in the literature allows direct comparison with the results obtained with other methodologies. The results obtained show a much higher accuracy than all the methods applied so far for all time resolutions and forecast horizons.

The paper is structured as follows: The proposed methodology is described in the Section 2; the dataset is analysed and the metrics used are described in the Section 3; the results are presented and discussed in the Sections 4 and 5, finally the Section 6 presents our conclusions and expected developments for this activity.

2. Proposed Method

The electricity consumed by a family can be decomposed into a long-term and a short-term component. The first component includes the effect of any upward or downward trends in consumption and, above all, the effect of the periodicity of average demand. In the case of electricity consumption, a number of daily, weekly and annual frequencies are distinguished. This component describes the habits of the family. However, seasonal behaviour can only describe average behaviour and not short-term fluctuations linked to deviations from the standard routine. This second component

is clearly difficult to predict, it can be linked to exceptional events, to short-term changes linked for example to the ignition of a energy-intensive appliance and, more generally, to deviations from the routine. It is assumed that these short-term variations can be described by means of an auto-regressive function, i.e. linking consumption in the immediate future, with respect to the instant of emission of the forecast, to the load measured in the immediately preceding times.

The total load can therefore be described by means of an additive model:

$$y(t) = \hat{y}_{lt}(t) + \hat{y}_{st}(t) + \epsilon, \quad (1)$$

where $y(t)$ is the load at the time t , \hat{y}_{lt} and \hat{y}_{st} are the long term and short term components, and ϵ is the forecast error.

In this paper, the long-term component of the forecast is estimated using a random forest (RF) regression model. The RF methodology has been used for short-term load forecasting in the literature, for example in [22] the authors use this technique to predict the hourly electrical load data of the Polish electrical system, in [23] this technology is applied to load forecasting at a university campus in Cartagena (Spain), in [24] the accuracy of consumption forecasts of residential customers one day in advance is analyzed as a function of time, granularity, and size of residential customers.

The random forest method, introduced by Breiman et al. [25] is an ensemble learning methodology that can be used for both classification and regression. It is based on the construction of a forest of unrelated decision trees, which corrects the trend towards over-fitting on the training set of decision trees.

In the training phase the technique of bootstrap aggregating (bagging) is applied to tree learning, through which a random subset with replacement from the training set is selected B times, the subset of samples from the training set X is identified with $X_b \subset X$ and the corresponding label $Y_b \subset Y$. For each of these subset a tree f_b is fitted. In the decision trees training process an additional randomization, called feature bagging, is used, which consists in considering for each candidate split a random subset of features.

After training, the prediction for a x sample is obtained by averaging the predictions of all the generated regression trees:

$$\hat{y}_{lt} = \sum_{b=1}^B f_b(x) \quad (2)$$

Our procedure uses a forest of 100 trees, whose optimal depth is established through a cross-validation procedure. The features used are only of temporal type and uniquely identify each measurement: the year, the day of the year, the day of the week, and the time of the day, expressed as a real number including fractions of an hour.

Many of the recently proposed algorithms for predicting household consumption use a purely auto-regressive approach using non-linear methods based on neural networks. For the short term component of the forecast we will use an approach based on a simple and fast step-wise multiple linear regression (MLR). It is given by:

$$\hat{y}_{st}(t + j\Delta t) = \beta_{0,j} + \sum_{i=1}^{n_i} \beta_{i,j} r(t - (j - 1)\Delta t), \quad (3)$$

with

$$r(t) = y(t) - \hat{y}_{lt}(t). \quad (4)$$

The short-term forecast \hat{y}_{st} for each time $t + j\Delta t$, with $j = 1, \dots, L$, is a linear combination of the residual r of the long-term forecast relative to the actual value of power consumption for the most recent n_i steps. $\beta_{i,j}$ denote the regression parameters. The forecast is obtained simultaneously for n_j future steps.

Given the great variance of domestic consumption, it is important to provide an estimate of the forecast error for each forecast time. This can be obtained in a simple way by analyzing the distribution of the error that the model has made in the training set and applying the same distribution of the error also in the forecasting phase. The distribution of the error is described by 19 quantiles ($q = 0.05, 0.10, \dots, 0.95$), and is parameterized according to the time of forecast j and for each hour of the day h , thus creating a look-up table $E_q(j, h)$ to be applied even in the forecast phase.

The procedures described were implemented using only public domain tools. The code is written entirely in python, using the pandas libraries [26] and scikit-learn [27] for machine learning tools. All libraries have been used in the latest revision available at the time of writing. The procedure has minimal requirements for its execution, the time required for training and forecasting are listed in Table 3, a PC with 16 GB of RAM and a quad core Intel i5 processor at 3.2 GHz has been used for the analysis.

3. Evaluation Setup


The presented method was evaluated on a reference dataset of electricity consumption for a single residential customer, called “Individual household electric power consumption Data Set” [28]. This archive contains 2,075,259 measurements gathered in a house located in Sceaux (7 km of Paris, France) between December 2006 and November 2010 (47 months). The dataset contains the household global minute-averaged active power measurement, as well as measurements of reactive power, current, voltage and energy consumed in three sub-meters intended respectively for the kitchen, a laundry-room with the electric water-heater, and an air conditioner circuit. In this paper we will only deal with the global active power.

Table 1 shows a description of the time series, varying with the aggregation time being used. The original series has a sampling rate of one minute, in the following, we will analyze forecast scenarios in which we average the measurement of consumption over periods of a quarter of an hour, of an hour or of a week. The time series is not complete with 1.25% of missing data. The missing measurements are replaced with the value measured a week earlier at the same time.

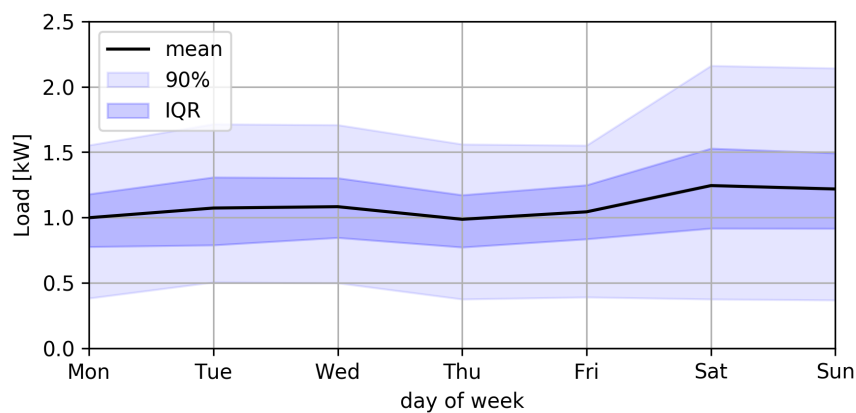
Table 1. Load data series description with different aggregation periodicity, the load is averaged in the selected period. The distribution is not normal, asymmetrical with a fat tail for high loads.

Period	1 min	15 min	1 h	1 Day	1 Week
Count (-)	2,049,280	136,639	34,168	1433	207
Mean (kW)	1.092	1.092	1.092	1.092	1.096
Std. (kW)	1.057	0.991	0.898	0.420	0.337
Min. (kW)	0.076	0.078	0.124	0.174	0.184
25% (kW)	0.308	0.321	0.342	0.817	0.878
50% (kW)	0.602	0.655	0.803	1.081	1.099
75% (kW)	1.528	1.563	1.579	1.324	1.329
Max. (kW)	11.122	8.566	6.561	3.315	2.505

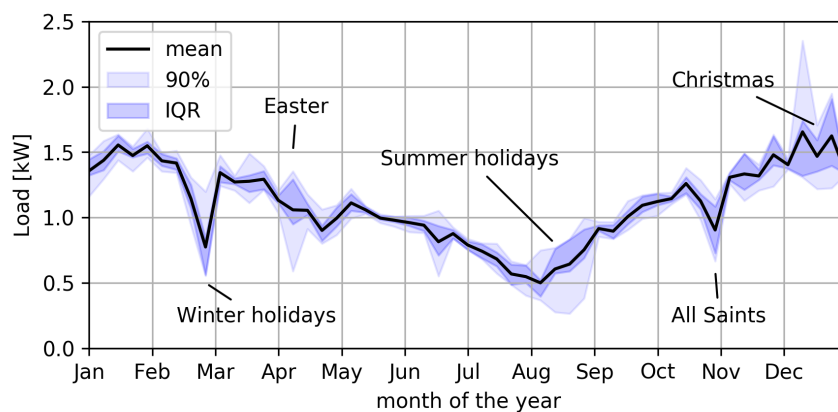
The daily, weekly and annual periodicity of the time series is illustrated in Figure 1. In Figure 1a the measurements are averaged over the hour, and the average value of the consumption of the whole series is shown as a function of the time of day, as well as the variation of the measurements around the average value, expressed by the inter quartile range (IQR) and by the interval between 5% and 95%. The weekly seasonality, with average daily consumption values, is shown in Figure 1b, while Figure 1c shows the annual seasonality with average consumption values on the week.



(a) Daily seasonality



(b) Weekly seasonality



(c) Yearly seasonality

Figure 1. Seasonality and load variability at different resolutions. The graphs show the average value of the load and the variability ranges of 90% (between 5% and 95%) and the inter quartile range (IQR) between 25% and 75%: (a) The hourly average load and its periodicity as a function of the time of day; (b) the daily average load and the periodicity as a function of the day of the week from Monday to Sunday; (c) the weekly average consumption and its periodicity as a function of the week of the year.

With almost four years of data available, it was decided to use the first three years of measurements for model training and to use the last year for verification, just like in [6,14,15]. The choice of the parameters of the RF model was made through a validation procedure, aimed at determining the optimal depth of the tree. Validation is carried out using the third year of measurements as a validation

set and training on the first two years of data. The optimal depth value is the one that determines the smallest fitting error in the validation set. The remaining parameters of the RF algorithm have been fixed, in particular, feature bagging is performed on half of the features and the number of trees in the forest is fixed at 100, having verified that a larger value does not bring any benefit in terms of accuracy for this dataset.

Some well-known accuracy metrics will be used in the following to assess the quality of the deterministic forecast. Specifically, the mean absolute error (MAE) and the root mean square error (RMSE):

$$\text{MAE} = \frac{1}{NL} \sum_{i=1}^N \sum_{l=1}^L |p_{i+l} - \hat{p}_{i,l}|, \quad (5)$$

$$\text{RMSE} = \sqrt{\frac{1}{NL} \sum_{i=1}^N \sum_{l=1}^L (p_{i+l} - \hat{p}_{i,l})^2}, \quad (6)$$

where N is the total number of power measures, L is the number of time steps Δt predicted in the future with respect to time t_i , p_{i+l} is the actual power at time $t_i + l\Delta t$, $\hat{p}_{i,l}$ is the power forecast l time steps in the future, predicted at time t_i .

Skill scores are widely used in assessing the performance of weather forecasting methods. They are defined as a measure of the relative improvement of a forecasting method over a reference. A commonly-used reference is the persistence forecast, that predicts that the power load on a given time will be the same measured at the same time one day before for lead time up to one day, one week before for one week lead time and one year before for one year lead times. Using the RMSE as a measure of accuracy, the skill score (SS) is defined as [29]:

$$\text{SS} = 1 - \frac{\text{RMSE}_f}{\text{RMSE}_p}, \quad (7)$$

where RMSE_f and RMSE_p are the root mean square error of the forecast method and of the persistence, respectively. The higher the skill score, the better.

To evaluate the accuracy of the probabilistic forecast we will use the continuous ranked probability score (CRPS) [30]:

$$\text{CRPS} = \frac{1}{N} \sum_{i=1}^N \sum_{h=1}^H \int_{-\infty}^{\infty} (F_{i+h}(x) - \hat{F}_{i,h}(x))^2 dx, \quad (8)$$

where $F_{i+h}(x)$ is the cumulative distribution function (CDF) of the probabilistic forecast for the i -th value, while $\hat{F}_i(x)$ is the CDF of the observations. Note that the CPRS coincides with the MAE for a deterministic forecast [31]. Small values of CRPS indicate good performances.

The analysis of the results will also show the variation of RMSE and CRPS according to the time of day and the forecast step in the future. Error measurements are obtained by grouping the data and applying the (6) and (8) on the sub-sets thus obtained; for example, the value of the RMSE at the time h of the day is obtained by applying the (6) only to load estimates made for times when the time of day is h .

4. Results

Some load forecast scenarios are analysed, based on the historical power series, differentiated on the basis of the sampling frequency used and the forecast horizon. The scenarios analysed are shown in Table 2. The selection of the scenarios has been carried out also according to the existing literature on the same dataset [6,14,15] to allow a direct comparison of the accuracy of the proposed method.

Table 2. The forecast scenarios analyzed, differing for the forecast horizon and for the temporal detail of the series.

Scenario	Sampling Rate	Forecast Interval	Lookback Steps	Forecast Steps
1	1 min	1 h	60	60
2	15 min	24 h	96	96
3	1 h	1 week	168	168
4	1 week	1 year	-	-

4.1. Forecast with One Minute Resolution

Scenario 1 provides for a time resolution of one minute with forecast times of up to one hour. This resolution was chosen because this time scale is used in the operation of the utility system and in real-time market activities, for example in automatic control of generation and resource redirection. Such a high resolution can also be used by home automation systems to prevent the risk of power cut-off due to overloading, as well as to operate small storage buffer systems.


The autoregressive short term component of the forecast is based on the measurements of the hour before the time of forecast. The number of lookback steps is therefore equal to 60, the same as the forecast steps, as shown in Table 2.

Table 3 shows the accuracy obtained for both deterministic and probabilistic prediction. Figure 2 shows the change in the accuracy of the method in terms of RMSE and CRPS as well as an example of the forecast result with the predicted confidence intervals.

Table 3. Results of the forecasting procedure for the various scenarios analysed. The optimal depth of the tree for the Random Forest algorithm used for long term forecasting, the values of the deterministic errors, mean absolute error (MAE) and root mean square error (RMSE), the skill score (SS) for the RMSE with respect to persistence, the accuracy of the probabilistic forecast measured with the continuous ranked probability score (CRPS), and the computational time required for the algorithm's training (including the search for the optimal depth of the trees) and for the emission of a forecast are listed.


Scenario	Optimal Depth (-)	MAE (kW)	RMSE (kW)	SS (%)	CRPS (kW)	Training Time (s)	Forecast Time (ms)
1	9	0.421	0.648	33.2	0.311	107	310
2	9	0.510	0.704	27.7	0.358	7.23	248
3	9	0.448	0.604	25.9	0.316	2.84	235
4	5	0.114	0.145	12.1	0.114	2.4	114

A persistence forecast is used to have an immediate comparison of the accuracy obtained. Given the limited extension of the time interval to be predicted in this scenario, the persistence is estimated by taking the measured power value one hour before the instant of prediction. The RMSE error is constantly lower than the value obtained with persistence and is logically smaller for the first instants of forecasting and then grows gradually as shown in Figure 2a. The accuracy of the forecast is shown in Figure 2b, where the value of the CRPS is reported as a function of the forecast lead time for persistence, for deterministic forecast (equivalent to the MAE) and for probabilistic forecast (where the estimate of the distribution of the forecast error is also used). The figure shows the effectiveness of the estimate of the error distribution, which results in a significant reduction in the CRPS measurement compared to a deterministic forecast and even more so with respect to persistence. Finally, Figure 2c shows an example of a forecast, carried out for the instants following a particularly high peak time in consumption. It can be seen that the forecast is quite accurate in predicting the trend of consumption even with a very high granularity, and how the estimated intervals for the probabilistic forecast increase in amplitude with the forecast time.



(a) RMSE variation in scenario 1

(b) CRPS variation in scenario 1



(c) Forecast for scenario 1


Figure 2. Results for scenario 1 (forecast interval 1 h, sampling rate 1 min): (a) Variation of the RMSE as a function of the forecast horizon, 60 steps correspond to 1 h; (b) variation of the probabilistic and deterministic CRPS as a function of the forecast horizon; (c) measured load, consumption forecast and probabilistic prediction intervals with forecast issued at 19:15 on 26 May 2010.

4.2. Forecast with Quarter Hour Resolution

The 15-min period is commonly used in literature and in operational practice for describing daily consumption profiles with sufficient detail. This resolution can be used in load shifting and peak clipping applications and for optimizing domestic storage systems. In scenario 2 a forecast with a horizon of up to one day and a resolution of 15 min, equal to 96 steps, is realized. For the short term part, the last 96 most recent measures are used at the time of emission of the forecast, see Table 2. The consumption measured at the same time the day before is used as persistence forecast.


The measurement of the RMSE as the forecast step changes is shown in Figure 3a, for scenario 2 which shows the greater accuracy of the model compared to persistence. Figure 3b shows the value of the CRPS according to the forecast step. It can be seen that in this scenario, compared to scenario 1,

the accuracy decreases rapidly after a few steps, equivalent to about 3 h and stabilizes on a constant value for the rest of the forecast interval. In this case too, the advantage guaranteed by the probabilistic forecast is appreciable. Figure 3c shows an example of a forecast made at the same time analysed in the previous scenario. The forecast follows very well the daily seasonality, the estimated forecast intervals grow very rapidly with the forecast time in accordance with Figure 3b.



(a) RMSE variation in scenario 2

(b) CRPS variation in scenario 2




(c) Forecast for scenario 2


Figure 3. Results for scenario 2 (forecast interval 1 day, sampling rate 15 min): (a) Variation of the RMSE as a function of the forecast horizon, 96 steps correspond to 24 h; (b) variation of the probabilistic and deterministic CRPS as a function of the forecast horizon; (c) measured load, consumption forecast and probabilistic prediction intervals with forecast issued at 19:15 on 26 May 2010.

Figure 4 shows the variation of the CRPS with the time of the forecast and with the forecast steps. Figure 4a shows the value of the CRPS grouped according to the hour of the day for a forecast horizon of 15 min, equal to one step in the future. Figure 4b,c shows the same graph for one-hour and two-hour forecast horizons. In this case too, it is noticeable that the accuracy of both the deterministic and the probabilistic forecast decreases as the forecast horizon increases, but also how the accuracy depends on the forecast time: the hours of greatest consumption, in the morning and in the evening are also the hours in which there is greater variability and a greater uncertainty in the forecast is obtained.


The CRPS of the deterministic prediction (equivalent to the MAE) is constantly lower than the value obtained with persistence, while the improvement obtainable with the probabilistic prediction, present for all time horizons, is more evident with the growth of the forecast horizon, in Figure 4b,c.



(a) Forecast horizon, 1 step equivalent to 15 min



(b) Forecast horizon, 4 steps equivalent to 1 h



(c) Forecast horizon, 8 steps equivalent to 2 h

Figure 4. CRPS results as a function of the forecast horizon and the time of the forecast for scenario 2 (forecast up to one day with sampling rate 15 min).

4.3. Forecast with One Hour Resolution

The time resolution of one hour is commonly used in the literature for longer forecast horizons, and can be of interest for domestic applications mainly for the monitoring of consumption and for the detection of any anomalies. In scenario 3, a forecast of up to one week, equivalent to 168 steps, is realized. For the short term forecast the 168 most recent measurements are used, while the

reference forecast of persistence is obtained with the measurement of consumption one week before the instant predicted.

The measurements of the RMSE and CRPS prediction error given in Table 3 and in Figure 5a,b highlight how the short term prediction improves accuracy significantly only for the hours immediately after the forecast was issued. There is also minimal impact in the following days for both the deterministic and the probabilistic prediction. Figure 5c shows an example of a one-week horizon forecast made at the same time selected for the previous scenarios. It can be seen that the proposed method is effective in forecasting both daily and weekly seasonality, both in terms of deterministic forecasting and in estimating its accuracy.

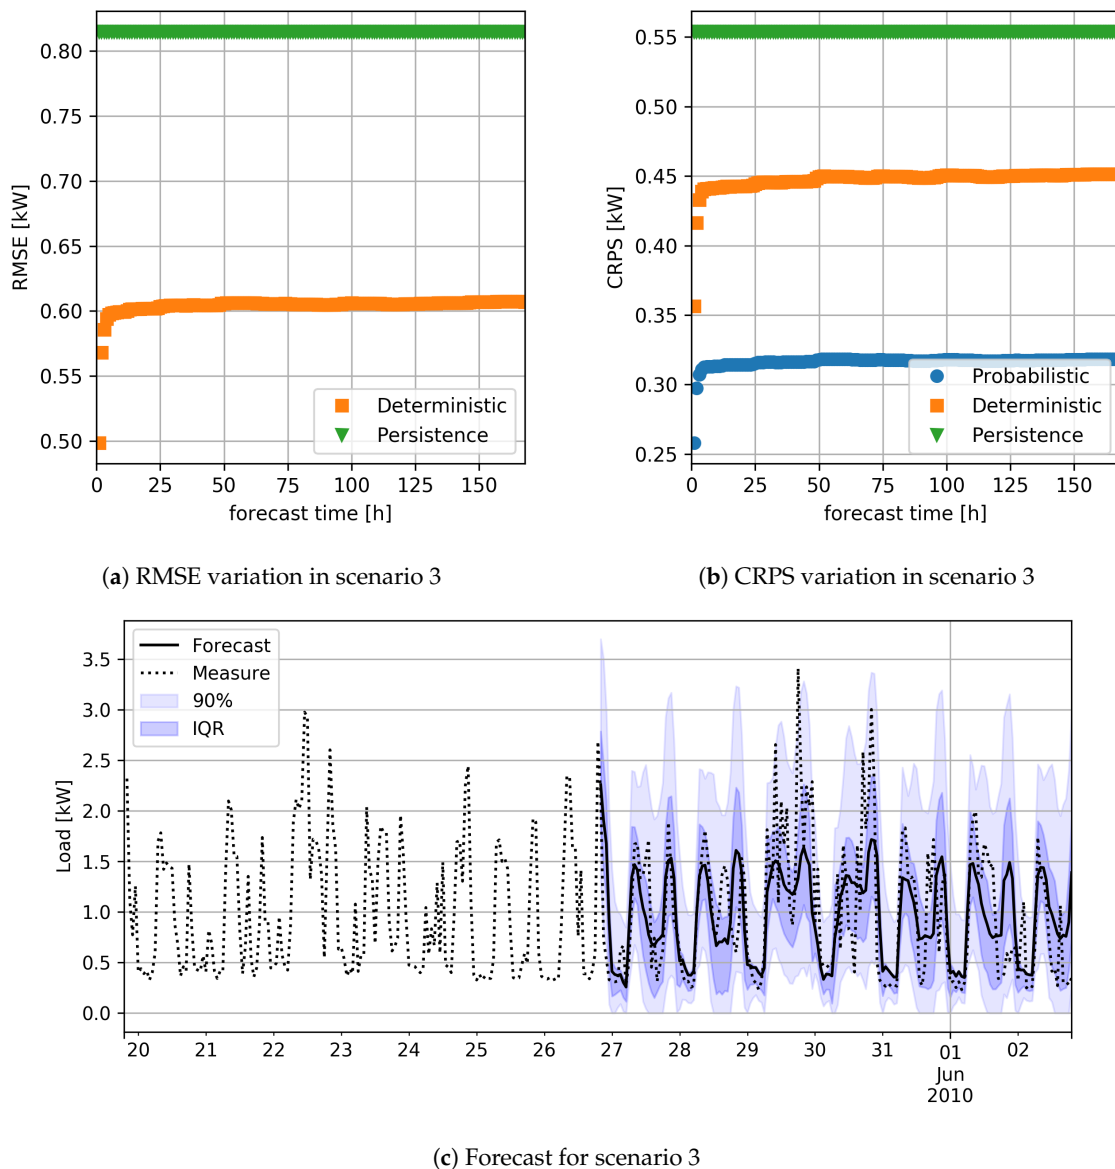


Figure 5. Results for scenario 3 (forecast interval 1 week, sampling rate 1 h): (a) Variation of the RMSE as a function of the forecast horizon, 168 steps correspond to 7 days; (b) variation of the probabilistic and deterministic CRPS as a function of the forecast horizon; (c) measured load, consumption forecast and probabilistic prediction intervals with forecast issued at 19:15 on 26 May 2010.

4.4. Forecast with One Week Resolution

Scenario 4 uses a one-week resolution for a forecast with a very long horizon of one year. In addition to assessing the accuracy of the proposed methodology and its applicability to a wide

range of time resolutions and forecast horizons, this scenario can be used in the household to verify the trend of its consumption and assess the effectiveness of any changes in habits that may lead to a more rational use of energy resources. In this case, the short term component of the forecast was not used, and the forecast with persistence used the measurement of consumption one year before the expected time.

The method accuracy measurement is shown in Table 3 while Figure 6 shows both the fitting obtained with the model and the forecast for the final year compared with the aggregate week measurements of the dataset.




Figure 6. Results for scenario 4 (average weekly consumption from 1 January 2010). The plot shows the measurements available in the dataset, the fit of the model in the training set consisting of the first 3 years of measurements and the forecast obtained for the last year.

5. Discussion

The public availability of the examined dataset allows an indirect comparison of the performances obtained with the proposed methodology with previous works that have analysed the same dataset. The selection of the scenarios was partially linked to the possibility of making such a comparison.

Table 4 shows the values obtained in literature using machine learning methodologies and deep learning techniques with recursive and convolutional neural networks already cited in the introduction. The considerably better result obtained with the proposed methodology is certainly linked to the choice of a hybrid approach in which the forecasting of phenomena in the long term, characterized by more seasonality (daily, weekly and annual) were clearly separated from the forecasting of the load in the short term, which was instead estimated with a simple auto-regressive methodology. The better performances can be only partially attributed to the chosen regression algorithm, it would probably be possible to obtain similar performances using neural networks suitably trained both for the long term component and for the short term component of our model, at the price of a more complicated feature pre-processing and of longer calculation times for the training. The use of a hybrid methodology, as shown in [32], allows instead to effectively isolate the average behavior from the alterations of the same and to train a regression model more quickly and effectively.

Table 4. Comparison of RMSE with the results of other forecasting procedures available in the literature. Artificial neural network (ANN), support vector machine (SVM), recurrent neural network (RNN), conditional restricted Boltzmann machine (CRBM) and factored conditional restricted Boltzmann machine (FCRBM) results are from [6]; sequence to sequence long short term memory (S2S LSTM) results are from [14] and convolutional neural network (CNN) results are from [15]. Random forest (RF) and linear regression (LR).

Scenario	Description	Method	RMSE (kW)
1	1 h forecast, 1 min resolution	ANN	0.732
		SVM	1.995
		RNN	0.939
		CRBM	0.903
		FCRBM	0.666
		S2S LSTM	0.667
		RF + LR	0.648
2	1 day forecast, 15 min resolution	ANN	0.907
		SVM	1.344
		RNN	1.009
		CRBM	1.030
		FCRBM	0.899
		RF + LR	0.704
3	1 week forecast, 1 h resolution	ANN	0.785
		SVM	0.791
		RNN	0.916
		CRBM	0.691
		FCRBM	0.663
		S2S LSTM	0.625
		RF + LR	0.604
4	1 year forecast, 1 week resolution	ANN	0.246
		SVM	0.188
		RNN	0.457
		CRBM	0.182
		FCRBM	0.170
		RF + LR	0.145

The proposed methodology can naturally be improved both for the long term and the short term part. For the long term, we are planning to introduce predictors based on the weather conditions and holidays scheduled in the calendar. For the short term it may be interesting to evaluate the use of a non-linear regressor, which could be based on RF as well. For both components additional features can be represented by the sub-meter readings provided as additional information in the dataset. We expect that a possible improvement can also come from the adoption of pattern recognition techniques for the identification of the appliances activation, which can allow a higher precision in identifying consumption habits by the user [33]. These techniques typically require a much richer information content in the dataset, including the identification of the power-on periods of individual appliances, which are difficult to obtain and are not included in the dataset analyzed in this paper. In any case, given the great variability of the data to be predicted, we think of developing the forecasts from a probabilistic point of view, with a better estimate of the accuracy of the forecast, thus including the effects of the weekly and annual seasonality in the forecast interval clearly visible in Figure 1b,c. For example, the effects of alteration and high load variability during Easter, All Saints' Day and Christmas, winter and summer holidays are clearly detectable (see annotations in Figure 1c). These are typical peak load periods for European countries [4,34].

6. Conclusions

This work presents a hybrid machine learning methodology based on random forest and linear regression for the deterministic and probabilistic prediction of household consumption at different time horizons and resolutions. The approach is based on the combined forecasting of long and short periods, using in the first case temporal features for the identification of trends and various seasonalities of the time series, and, in the second case, an auto-regressive approach using the most recent load measurements available at the time of emission of the forecast. Finally, through the analysis of the forecast error of the model, a probabilistic load forecast is realized.

The analysis highlights the relationship between the accuracy of the forecast and the lead time of the forecast and the time of the forecast, and compares the result obtained both with a reference based on persistence and with the results of the analysis of the same dataset with machine learning and deep learning methodologies published in the literature.

The method proves to be very effective in terms of absolute precision and calculation times required for model training and forecasting. The method also offers opportunities for development for both the long term and short term components, using additional predictors, experimenting with non-linear methods for short-term forecasting and refining the probabilistic forecasting methodology.

Author Contributions: The authors jointly conceived of and designed the methodologies, performed the analysis and wrote the paper.

Funding: This research was funded by Regione Autonoma della Sardegna, Delibera 66/14 del 13.12.2016, Progetto Complesso area ICT and by European Union's Horizon 2020 research and innovation program under Grant Agreement No. 646463, project NETfficient, which also funded the APC.

Conflicts of Interest: The authors declare no conflict of interest. The founding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, and in the decision to publish the results.

References

1. Farhangi, H. The path of the smart grid. *IEEE Power Energy Mag.* **2010**, *8*. [[CrossRef](#)]
2. Mohassel, R.R.; Fung, A.; Mohammadi, F.; Raahemifar, K. A survey on advanced metering infrastructure. *Int. J. Electr. Power Energy Syst.* **2014**, *63*, 473–484. [[CrossRef](#)]
3. Wang, Y.; Chen, Q.; Hong, T.; Kang, C. Review of smart meter data analytics: Applications, methodologies, and challenges. *IEEE Trans. Smart Grid* **2018**. [[CrossRef](#)]
4. Gajowniczek, K.; Nafkha, R.; Ząbkowski, T. Electricity peak demand classification with artificial neural networks. In Proceedings of the 2017 Federated Conference on Computer Science and Information Systems (FedCSIS), Prague, Czech Republic, 3–6 September 2017; pp. 307–315.
5. Massidda, L.; Marrocu, M. Decoupling Weather Influence from User Habits for an Optimal Electric Load Forecast System. *Energies* **2017**, *10*, 2171. [[CrossRef](#)]
6. Mocanu, E.; Nguyen, P.H.; Gibescu, M.; Kling, W.L. Deep learning for estimating building energy consumption. *Sustain. Energy Grids Netw.* **2016**, *6*, 91–99. [[CrossRef](#)]
7. Yu, C.N.; Mirowski, P.; Ho, T.K. A sparse coding approach to household electricity demand forecasting in smart grids. *IEEE Trans. Smart Grid* **2017**, *8*, 738–748. [[CrossRef](#)]
8. Fan, S.; Hyndman, R.J. Short-term load forecasting based on a semi-parametric additive model. *IEEE Trans. Power Syst.* **2012**, *27*, 134–141. [[CrossRef](#)]
9. Taylor, J.W. Exponentially weighted methods for forecasting intraday time series with multiple seasonal cycles. *Int. J. Forecast.* **2010**, *26*, 627–646. [[CrossRef](#)]
10. De Livera, A.M.; Hyndman, R.J.; Snyder, R.D. Forecasting time series with complex seasonal patterns using exponential smoothing. *J. Am. Stat. Assoc.* **2011**, *106*, 1513–1527. [[CrossRef](#)]
11. Wong, S.L.; Wan, K.K.; Lam, T.N. Artificial neural networks for energy analysis of office buildings with daylighting. *Appl. Energy* **2010**, *87*, 551–557. [[CrossRef](#)]
12. Edwards, R.E.; New, J.; Parker, L.E. Predicting future hourly residential electrical consumption: A machine learning case study. *Energy Build.* **2012**, *49*, 591–603. [[CrossRef](#)]

13. Chitsaz, H.; Shaker, H.; Zareipour, H.; Wood, D.; Amjady, N. Short-term electricity load forecasting of buildings in microgrids. *Energy Build.* **2015**, *99*, 50–60. [[CrossRef](#)]
14. Marino, D.L.; Amarasinghe, K.; Manic, M. Building energy load forecasting using deep neural networks. In Proceedings of the IECON 2016—42nd Annual Conference of the IEEE Industrial Electronics Society, Florence, Italy, 23–26 October 2016; pp. 7046–7051.
15. Amarasinghe, K.; Marino, D.L.; Manic, M. Deep neural networks for energy load forecasting. In Proceedings of the 2017 IEEE 26th International Symposium on Industrial Electronics (ISIE), Edinburgh, UK, 19–21 June 2017; pp. 1483–1488.
16. Hong, T.; Fan, S. Probabilistic electric load forecasting: A tutorial review. *Int. J. Forecast.* **2016**, *32*, 914–938. [[CrossRef](#)]
17. Taieb, S.B.; Huser, R.; Hyndman, R.J.; Genton, M.G. Forecasting uncertainty in electricity smart meter data by boosting additive quantile regression. *IEEE Trans. Smart Grid* **2016**, *7*, 2448–2455. [[CrossRef](#)]
18. Arora, S.; Taylor, J.W. Forecasting electricity smart meter data using conditional kernel density estimation. *Omega* **2016**, *59*, 47–59. [[CrossRef](#)]
19. Haben, S.; Giasemidis, G.; Ziel, F.; Arora, S. Short Term Load Forecasts of Low Voltage Demand and the Effects of Weather. *arXiv* **2018**, arXiv:1804.02955.
20. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. Statistical and Machine Learning forecasting methods: Concerns and ways forward. *PLoS ONE* **2018**, *13*, e0194889. [[CrossRef](#)] [[PubMed](#)]
21. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M4 Competition: Results, findings, conclusion and way forward. *Int. J. Forecast.* **2018**, *34*, 802–808. [[CrossRef](#)]
22. Dudek, G. Short-term load forecasting using random forests. In *Intelligent Systems' 2014*; Springer: Cham, Switzerland, 2015; pp. 821–828.
23. Ruiz-Abellón, M.; Gabaldón, A.; Guillamón, A. Load Forecasting for a Campus University Using Ensemble Methods Based on Regression Trees. *Energies* **2018**, *11*, 2038. [[CrossRef](#)]
24. Lusi, P.; Khalilpour, K.R.; Andrew, L.; Liebman, A. Short-term residential load forecasting: Impact of calendar effects and forecast granularity. *Appl. Energy* **2017**, *205*, 654–669. [[CrossRef](#)]
25. Breiman, L. Random forests. *Mach. Learn.* **2001**, *45*, 5–32. [[CrossRef](#)]
26. McKinney, W. Data structures for statistical computing in python. In Proceedings of the 9th Python in Science Conference, Austin, TX, USA, 28–30 June 2010; Volume 445, pp. 51–56.
27. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
28. Dheeru, D.; Karra Taniskidou, E. UCI Machine Learning Repository. 2017. Available online: <https://archive.ics.uci.edu/ml/datasets/individual+household+electric+power+consumption> (accessed on 17 December 2018)
29. Murphy, A.H. Skill scores based on the mean square error and their relationships to the correlation coefficient. *Mon. Weather Rev.* **1988**, *116*, 2417–2424. [[CrossRef](#)]
30. Wilks, D.S. *Statistical Methods in the Atmospheric Sciences*; Academic Press: Cambridge, MA, USA, 2011; Volume 100.
31. Hersbach, H. Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather Forecast.* **2000**, *15*, 559–570. [[CrossRef](#)]
32. Smyl, S.; Ranganathan, J.; Pasqua, A. M4 Forecasting Competition: Introducing a New Hybrid ES-RNN Model. Available online: <https://eng.uber.com/m4-forecasting-competition/> (accessed on 17 December 2018).
33. Singh, S.; Yassine, A. Big data mining of energy time series for behavioral analytics and energy consumption forecasting. *Energies* **2018**, *11*, 452. [[CrossRef](#)]
34. McSharry, P.E.; Bouwman, S.; Bloemhof, G. Probabilistic forecasts of the magnitude and timing of peak electricity demand. *IEEE Trans. Power Syst.* **2005**, *20*, 1166–1172. [[CrossRef](#)]

