



Load and wind power forecasting methods for the optimal management of isolated power systems with high wind penetration

Geoff A. Dutton, Georges Kariniotakis, Jackie A. Halliday, Eric Nogaret

► To cite this version:

Geoff A. Dutton, Georges Kariniotakis, Jackie A. Halliday, Eric Nogaret. Load and wind power forecasting methods for the optimal management of isolated power systems with high wind penetration. Wind Engineering, 1999, 23 (2), p.69-87. hal-00538463

HAL Id: hal-00538463

<https://minesparis-psl.hal.science/hal-00538463>

Submitted on 20 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Load and wind power forecasting methods for the optimal management of isolated power systems with high wind penetration

A.G. Dutton¹, G. Kariniotakis², J.A. Halliday¹, E. Nogaret²

¹Energy Research Unit, CLRC Rutherford Appleton Laboratory, Chilton, Didcot, Oxon. OX11 0QX, UK

²Ecole des Mines de Paris, Centre d'Energetique, B. P. No. 207, F-06905 Sophia Antipolis, France

ABSTRACT

Load demand forecasting is essential for the management of autonomous power systems. Forecasts are required for the optimal scheduling of power units for the next 2-48 hours depending on the power system size. Where such systems are subject to increasing renewable energy penetration, this renewable generation capacity is often treated as a "negative load" due to its inherent unpredictability. A range of methods are presented and evaluated for the efficient forecasting of load and wind power generation in isolated power systems. The best of these methods have been implemented within an on-line advanced control advice system within the despatch centre of the island of Crete.

1. INTRODUCTION

In order to maximise the penetration of renewable energy plants, such as wind farms, within an isolated power system, an advanced control advice system called CARE (developed within EC JOULE III project no. JOR3-CT96-0119) has been developed. The purpose of this system is to advise the operators of the system when it will be necessary to alter the balance of conventional (fossil-fuelled) power plant (e.g. by starting or stopping units) and to allow the maximum savings to be made, whilst at the same time maintaining system security. A central part of such an advice system is the efficient prediction of the electrical demand and renewable energy generation for horizons between 1 and 48 hours ahead. A variety of forecasting techniques have been evaluated during the development of forecasting modules for the CARE software. These techniques and the results of the evaluation process are presented in this paper. The prototype system, incorporating the best of the algorithms described, has been implemented and is undergoing further evaluation on the island of Crete.

Conventional load forecasting models have been designed for interconnected systems, where load changes are relatively smooth, while forecasting for isolated systems has rarely been attempted. The target performance for forecasting models within interconnected systems is of the order of approximately 3% Mean Annual Percentage Error (MAPE). In isolated systems the load is characterised by more frequent non-predictable events such as faults and system overloads which lead, in turn, to power failures. The performance of different algorithmic forecasting techniques within this more complex situation is discussed in section 2.

Wind power is often seen by the operators as a "negative load", so that for a system with substantial penetration (e.g. 20%) we have the situation in which a large proportion of the load (e.g. 20 %) is predicted with a much lower accuracy than is usually considered acceptable. This presents a considerable challenge to the control advice system and there is a high premium on developing accurate forecasting techniques for renewable

energy resources. In this context, various advanced models for wind speed forecasting are presented in section 3.

2. LOAD FORECASTING

2.1. Introduction

Load forecasting is required on both short (several minutes up to several hours ahead) and long times scales (of the order of several hours ahead up to several days ahead), in order to provide input to the economic dispatch and unit commitment algorithms respectively of the CARE program.

Four different forecasting methods, with increasing level of sophistication, have been considered:

- (i) naive [DAY-1] – the load at any given hour in the next day will be the same as the load at the same time the previous day,
- (ii) naive [DAY-7] – the load at any given hour in the next day will be the same as the load at the same time the previous week,
- (iii) linear autoregressive (AR) model (with associated data detrending),
- (iv) fuzzy neural network.

In the case of Crete, the most suitable forecasting horizons were defined as 8 hours ahead at 20 minute intervals, and 24 or 48 hours ahead at one hour intervals. The results presented here concentrate on the 1-24 hours and 1-48 hours horizons.

Hourly time series data were made available from the island for four years (1994-1997). The data up to the end of 1996 were variously used for the estimation of model parameters and subsequent evaluation, but the data set for 1997 was reserved exclusively for evaluation. The original data was pre-processed to eliminate the effects of faults in data collection and power failures.

Figure 1 shows the mean weekly load on Crete for the four years to December 1997; note how the summer peak load is increasing quite steeply from one year to the next, due to the development of the tourist industry and the ever-increasing proportion of air-conditioned buildings. This graph emphasises very well the challenge facing the power station operators on the island. Figure 2 shows a typical weekly load cycle with higher loads during the weekdays and lower loads at the week-end.

The load forecasts presented within this paper are compared on the basis of the MAPE defined for each look ahead, k , as:

$$MAPE(k) = \frac{1}{N - K} \sum_{t=0}^{N-K} \left| \frac{Predicted_load(t+k) - Measured_load(t+k)}{Measured_load(t+k)} \right| \bullet 100\% \quad (1)$$

where N is the total number of data points and K is the maximum look ahead.

A model may have either a *fixed* or *moving* window. If the window is fixed then, for 24 hour look ahead schemes, the MAPE corresponds to a certain hour of the day.

2.2. Naïve [DAY-1] and Naïve [DAY-7] predictors

The performance of the simple, so-called “naïve” predictors, which utilise the load value from the same time the previous day [DAY-1] or the same time the previous week [DAY-7] can be used as a reference to set the maximum acceptable error from the more advanced methods (with the ultimate aim of approaching the acceptable load forecasting error of approximately 3% within interconnected systems).

For the Crete load data for 1997, figure 3 shows that the [DAY-1] predictor has an overall error of 5.90%, while the [DAY-7] predictor has an overall error of 6.31% (based on forecasts made every 24 hours at midnight for a 52 week period). However, the two methods vary considerably in relative effectiveness from day to day during the week. The [DAY-7] predictor is the best model for Sunday and Monday, while the [DAY-1] predictor performs best on all other days. It is further worth noting that on Wednesday through Friday, the [DAY-1] predictor achieves a MAPE of around 3.7% which is close to the target value for more sophisticated predictors in interconnected systems. A combined [DAY-1]/[DAY-7] prediction scheme would have achieved a MAPE of 5.02% in 1997, so this is the target for the more sophisticated predictors presented below.

2.3. Linear autoregressive (AR) load forecasting model

The linear autoregressive (AR) load forecasting model requires only the historic load time series (which is available from the SCADA system in any real application). This data must first be pre-processed to remove any data measurement errors and to smooth over any discontinuities due to load disconnections (the effects of which are definitely not wanted in the forecast!). Results are presented below from evaluation using historic time series data, which has been appropriately smoothed and corrected for load disconnections.

The algorithm used by the AR predictor model firstly detrends the load data to remove the weekly mean and the diurnal variation, and then fits an autoregressive model (containing both seasonal and non-seasonal terms) to the time series of residual values. The initial detrending of the data significantly improves the prediction accuracy.

The basic principles of the **detrending** procedure are explained in figure 4. The process is as follows:

- (i) the data is input in an integral number of N weeks (figure 4a) and the N weekly mean loads evaluated,
- (ii) the N weekly mean loads are then (optionally) subtracted from the initial data (figure 4b),
- (iii) the mean weekly load profile is then obtained (figure 4c),
- (iv) the mean weekly load profile is subtracted from the input data to produce a time series of residual load values (figure 4d) which are then fitted to a linear autoregressive model to produce a prediction for the next 24 hours (figure 4d),
- (v) this residual load prediction is often small by comparison with the daily mean load (figure 4e),
- (vi) the residual load, the daily mean load, and the weekly mean load are then reassembled to produce the load prediction, which, in time, can be compared with the actual load (figure 4f).

Clearly, there are a number of different ways that could have been chosen to detrend the data. The initial subtraction of the weekly mean load may be omitted (since there are conflicting results as to the effectiveness of this step – see below), or a single daily mean load profile evaluated instead of the weekly mean load profile (although current results suggest that this is not beneficial). However, the results clearly show that detrending is beneficial to the output from the AR model.

The stochastic component of the load for the next timestep is evaluated from an **autoregressive** (AR) model of the detrended time series, thus:

$$y'(t) = a_1 y(t-1) + a_2 y(t-2) + \dots + a_p y(t-p) + A_1 y(t-\delta) + A_2 y(t-2\delta) + \dots + A_q y(t-q\delta) + X(t) \quad (2)$$

where the a_i , $i = 1, \dots, p$ are p non-seasonal parameters to be estimated, the A_i , $i = 1, \dots, q$ are q “seasonal” parameters (e.g. load values from multiples of 24 hours previous to the current timestep) with some time period δ and $X(t)$ is a white noise term.

The best results using an AR model have been achieved with just one seasonal parameter (with a period of 24 hours) and no non-seasonal parameters. Results are presented in table I for three detrending options:

- (i) “full” detrending in which the N weekly means are subtracted *before* evaluating the weekly mean load profile,
- (ii) “daily” detrending in which the weekly mean load profile is evaluated from the raw data,
- (iii) no detrending.

The use of detrending is clearly advantageous (MAPE of 4.18% for the 1995 data, compared with a MAPE of 6.0% with no detrending), although the usefulness of removing the weekly means before removing the diurnal variation is not clear, since it gives marginally better results in 1997 and worse results in 1995. The results for 1997 with detrending by removal of the diurnal variation only is shown in figure 5. Note that, when the individual days are separated out, the method does slightly worse than persistence [DAY-1] on Wednesday to Friday, but considerably better at the week-end. The inclusion of a single non-seasonal parameter (figure 6) offers considerably better performance for very short look ahead times (up to 3 hours ahead), but worse performance overall.

It is important to understand therefore that the choice of the most appropriate model depends on the look ahead horizon required.

However, while it is clear that different models may be more appropriate for different look ahead times and even different days of the week, the relative performance of certain models cannot always be predicted even for a time series as consistent as the Cretan load data.

Any algorithmic load prediction is sensitive to the quality of the input data and extensive checking is required before the data is input to the module. It is interesting to note the large spike in the residual load data shown in figure 4d at around –300 hours. This actually occurs on a Sunday on which the morning peak load is unusually higher than the afternoon peak load (figure 4a). It is important to verify whether this is a genuine perturbation since its effects will persist in the forecast.

2.4. Fuzzy neural network based load forecasting model

A fuzzy autoregressive model can be described by a set of rules having the general form :

$$\begin{aligned} R^1 : & \text{ IF } (x_1 \text{ is } A_1^1, \dots, \text{ and } x_m \text{ is } A_m^1) \text{ THEN } \hat{y}^1 = g^1(x_1, \dots, x_m) \\ R^r : & \text{ IF } (x_1 \text{ is } A_1^r, \dots, \text{ and } x_m \text{ is } A_m^r) \text{ THEN } \hat{y}^r = g^r(x_1, \dots, x_m) \end{aligned} \quad (3)$$

where :

A_j^i is a fuzzy set characterised by the membership function $\mu_i^j(x_j)$ with $i=1, \dots, r$ and $j=1, \dots, m$ (where fuzzy sets are represented by Gaussian functions),

$\hat{y}^i$ is the output of the i -th rule (the function g^i implies the value of $\hat{y}^i$ when $x_1, \dots, x_m$ satisfy the premise; a linear function is taken $g^i(x) = p_0^i + p_1^i x_1 + \dots + p_m^i x_m$),

$x_1, \dots, x_m$ are past values of the process.

If x_i represents "load", then a number of fuzzy sets, representing vague terms like "low", "medium", or "high" may be associated to this variable (Kariniotakis et al 1996). Being in the form of IF-THEN rules, a fuzzy model can be used as an explanatory model for the process itself whose underlying dynamics are unknown. The dynamic process is described locally since at each rule a different model of the process is associated and is activated only for certain regions of values of the input variables. Non-stationarity and seasonalities, are favourable factors to apply local modelling.

Figure 7 presents sample results on load forecasting using fuzzy modelling. The average performance is 3.2 % for both ordinary and special days. These results are obtained by n -MISO models developed for each hour of the day.

2.5. Sources of error in load forecasting

The load forecasting module within CARE outputs the predicted load values for each prediction step with two further values representing the bounds of the 95% confidence interval on the AR forecast. An example of the performance to be expected for two (high load) summer days is shown in figure 8. The reliability of these forecasts in making secure unit commitments within CARE is currently under assessment.

2.6. Load forecasting : conclusions and discussion

A linear AR model can give a clear improvement over PERS[DAY-1] forecasts, when averaged across an entire week, but some days of the week (notably Wednesday to Friday) are better predicted by PERS[DAY-1].

The naïve predictor PERS[DAY-7] performs better than the naïve predictor PERS[DAY-1] at week-ends (particularly Sunday and Monday), but both are out-performed by the linear AR model.

The application of fuzzy modelling techniques and the development of different models for different hours of the day has achieved a further performance improvement, very close to that required for large, interconnected systems.

The growing sophistication of these methods implies increasing problems of portability between locations due to the requirements to tune the algorithm to specific local

conditions. Notwithstanding the requirement for robust on-line data checking and smoothing of the input time series, the robustness of the algorithms to data input errors and changing conditions on the island must be closely monitored.

3. WIND POWER FORECASTING

3.1. Introduction

As the penetration of wind power into a system increases, so there is an increasing benefit from improved methods of wind power forecasting. Typically, wind power is viewed by the system operators as a “negative load”, which is predictable with a much lower accuracy than is considered acceptable for more conventional loads. This results in the operators having to schedule a higher level of spinning reserve, with consequently higher expenditure.

Wind power forecasting is approached in the CARE control advice system by first forecasting the future wind speed and then converting this estimate to a power estimate through a filter representing the characteristics of the wind farm.

Four basic methods for wind speed forecasting are proposed for incorporation in the CARE software:

- (i) persistence – the wind speed at the current time is assumed to remain constant throughout the forecast period,
- (ii) linear autoregressive ARMA model,
- (iii) adaptive fuzzy logic based model,
- (iv) fuzzy neural network based on geographically distributed wind data.

The forecast horizons were specified to be the same as for load forecasting, namely 8 hours ahead at 20 minute intervals, and 24 or 48 hours ahead at 1 hour intervals

Wind speed time series data were collected on Crete and subsequently analysed to remove gaps in the data arising from faults due to the power supply. The resulting time series comprised just short of 120 days of data (2877 hourly values) – see figure 9. Since it is considered that this data still covers too short a time period to provide sufficiently long learning and evaluation data sets, some of the models have also been tested against a 364-day database of 1-minute wind speeds from Shetland (processed into 8736 hourly values).

An alternative meteorological-based approach to wind speed forecasting is represented in the project JOR3-CT95-0008 *Implementing short term prediction at utilities*, which includes a case study for the island of Crete. Within this project site-tailored (using WAsP) wind speed forecasts were developed from the Danish Meteorological Institute’s HIRLAM numerical weather prediction model and made available on the Internet. These forecasts have not been used within CARE to date since there are doubts surrounding their future availability and efficient “tailoring” of the data to future wind farm sites, however, the approach shows promise for incorporation in CARE adapted to larger systems.

3.2. Linear autoregressive (AR) wind speed forecast model

The on-line linear autoregressive (AR) wind speed forecasting module requires the wind speed time series (typically 1 minute to 1 hour timestep average values) from each wind farm. The subsequent conversion of wind speed to wind power requires the application of a filter comprising the mean power curve (see, for example, figure 10) for each type of wind turbine in the wind farm and an input of the number of operating wind turbines of each type on the farm. It is hoped that in the future this filter can be improved by the use of model output statistics (MOS) to improve the measured wind speed/power output correlation for each individual wind farm, taking account of factors such as array effects and degradation of performance with time.

Results are presented below from evaluation using historic wind speed time series data, which has been checked and corrected for missing or extreme values.

The wind speed prediction for the next timestep is evaluated from an **autoregressive (AR)** model of the input wind speed time series, thus:

$$y'(t) = a_1 y(t-1) + a_2 y(t-2) + \dots + a_p y(t-p) + X(t) \quad (4)$$

where the a_i , $i = 1, \dots, p$ are p non-seasonal parameters to be estimated δ and $X(t)$ is a white noise term..

An initial assessment of the most suitable autoregressive (AR) model parameters was made using the Shetland "wind year" data. This data comprises 364 days of 1-minute data for two distinct sites, which has been converted into 10-minute and hourly mean data for use in the evaluation.

A full parametric study was beyond the scope of the CARE project, but a number of indicative results are presented below. The main parameters which have been varied are the model order, p , and the number of input data values.

Different parameters have been compared on the basis of the root mean square (RMS) error for different look ahead intervals, k , defined as:

$$RMS(k) = \sqrt{\frac{1}{N-K} \sum_{t=1}^{N-K} \left[\text{Predicted_wind_speed}(t+k) - \text{Measured_wind_speed}(t+k) \right]^2} \quad (5)$$

Figure 11 demonstrates the performance of an AR(6) model with the Shetland (Scroo) hourly data for look ahead times up to 48 hours ahead, for models using 200, 600, and 1200 previous data points for characterisation, compared with persistence. All the graphs show an improvement over persistence. There is a further improvement on increasing the number of previous data points from 200 to 600 (except for very short look ahead times), but extending this window to 1200 data points is inconclusive. For short data inputs (200 values), it can be further shown that an AR(6) model out-performs an AR(12) model for look ahead times up to 18 hours ahead, after which the AR(12) model is superior. However, this effect is much less with longer data input time series (when the AR(6) model is possibly superior for the first two hours, after which there is little to choose between the two models).

Attempting to predict 10-minute wind speed data values for 8 hours ahead shows similar trends (figure 12). For short look ahead times, the benefit of extending the window to encompass a greater number of data points is particularly noticeable. It seems likely that the critical factor is the time scale of significant events in the wind spectrum, rather than the actual number of data points, and some further work is required to clarify this.

Predictions for the hourly data from Crete are presented in figure 13. This compares reasonably with the data from Shetland (figure 11) except that the curve is generally less smooth and the shape appears more linear (rather than gently convex) with look ahead time. These differences are, at least in part, due to the shortness of the data set from Crete, but there may also be meteorological differences which would justify further investigation (for example, it seems that the winds are more constant on Crete and the persistence forecast itself is therefore a better forecast than on Shetland).

3.3. Fuzzy logic based models for wind speed forecasting

Two fuzzy models were tuned through different structure optimisation procedures. A FARX(4;0,0) model (4 rules) was obtained by a run of the structure optimisation algorithm using as objective function the forecasting performance for an horizon of 24 hours. The second FARX(8;2,2) model (24 rules) was obtained by an optimisation procedure piloted by the performance of the model for an horizon of 8 hours ahead.

The "long-term" model FARX(4;0,0) is optimised to have a good performance for a 24-hours horizon. It uses 4 past values of wind speed. The "short-term" model FARX(8;2,2) was optimised to have a good performance for an horizon of 8 hours ahead. It uses 8 past values of speed and two past values of each transformed polar co-ordinate of wind direction. As with the AR models, a clear benefit with respect to persistence is gained by using the models for horizons up to 18 hours ahead (see figure 14).

3.4. Wind speed forecasting within the CARE advice system

The wind forecasting module outputs the predicted wind speed values for each prediction step with two further values representing the extremes of a 95% confidence interval. An example of the performance to be expected for several days chosen at random is shown in figure 15. It is clear that, although performance of the module is considerably better than persistence, beyond 12 hours the width of the 95% confidence interval is too large to be useful and that considerable divergence from even this wide range is probable. It is possible that more meteorological based forecasting may be more successful for such long forecasting horizons.

3.5. Wind forecasting : conclusions and discussion

For forecast horizons up to 2 hours ahead, there is a small improvement in the performance of the linear AR model compared to persistence for the Shetland time series data (however, these results are not repeated for the shorter Crete time series data).

For forecast horizons up to 8-12 hours ahead, there is generally improved performance over persistence of up to 20% (RMS error).

For forecasting horizons greater than 12 hours, the very wide 95% confidence bands (which basically encompass all likely wind speed values) suggest that an alternative, meteorological-based approach may be more beneficial.

4. CONCLUSIONS

Sophisticated models have been developed to predict load demand and wind power generation for an island power network.

The best results for load demand forecasting have been obtained with fuzzy neural network based techniques, but the resulting module must be optimised for different load patterns. This may be a problem if the load pattern changes markedly over time, or the module is required to operate for a different network. A standard autoregressive approach has also been implemented, which performs slightly less well but is likely to be more robust to changes in load pattern and should be readily transportable.

Wind speed forecasts of up to 20% better than persistence have been obtained for medium term horizons. For forecast horizons less than 2 hours ahead, a marked difference has been found between the wind regimes of Shetland and Crete, which make persistence forecasts difficult to out-perform in the latter case. Beyond 12 hours ahead, the error bands on the forecasts make them increasingly redundant and an alternative meteorological-based approach is likely to be more fruitful. Data is just coming on-line and some further work is required to validate and optimise the conversion of wind speed to wind power for specific wind farms.

Nomenclature

AR	:	autoregressive
MAPE	:	mean absolute percentage error
p.c.	:	personal computer
RMS	:	root mean square error
SCADA	:	Supervision, Control, and Data Acquisition

References

Kariniotakis, G., Nogaret, E., Stavrakakis, G., 1996, 'A fuzzy and a neural network based wind power forecasting model', 1996 *European Union Wind Energy Conf.* 20-24 May 1996, Göteborg, Sweden, pp. 596-599.

Case	MAPE (%) Monday	MAPE (%) Friday	MAPE (%) Week
PERS [DAY-1] 1997 (1995)	10.60 (11.20)	3.67 (3.05)	5.90 (5.73)
PERS [DAY-7] 1997 (1995)	6.77 (8.50)	6.33 (6.26)	6.31 (6.53)
Year 1997			
Daily detrending only / AR(0,1) - seasonal 24 hours / 4 weeks			
	5.01	4.08	4.57
Full detrending / AR (0,1) – seasonal 24 hours / 4 weeks			
	5.12	4.25	4.54
Daily detrending only / AR(1,1) – seasonal 24 hours / 4 weeks			
	5.08	4.59	4.78
Year 1995			
Daily detrending only / AR(0,1) - seasonal 24 hours / 4 weeks			
	5.39	3.85	4.18
Full detrending / AR(0,1) – seasonal 24 hours / 4 weeks			
	5.94	4.03	4.43
No detrending / AR (0,1) – seasonal 24 hours / 4 weeks			
	11.90	3.35	6.00
Daily detrending only / AR(1,1) - seasonal 24 hours / 4 weeks			
	5.69	4.26	4.53

Table I : Initial results of AR parameter studies




Figure 1 : Variation in weekly mean load on Crete (1994-1997)




Figure 2 :Variation in hourly load for a randomly selected week




Figure 3 : Load forecasting by persistence [DAY-1] (dotted line) compared with persistence [DAY-7] (solid line) - Crete 1997




Figure 4 (a-c): Load detrending procedure




Figure 4 (d-f): Load detrending procedure




Figure 5 : Load forecasting AR model (0,1) - seasonal 24 hours - detrending diurnal only (solid line) compared with [DAY-1] (dotted line)




Figure 6 : Load forecasting AR model (1,1) - seasonal 24 hours - detrending diurnal only (solid line) compared with [DAY-1] (dotted line)




Figure 7 : Performance of the fuzzy model compared to the [DAY-1] predictor




Figure 8 : Load forecasting AR model (0,1) - seasonal 24 hours - predicted data up to 48 hours ahead - best prediction (solid line), 95% confidence intervals (dashed line), actual readings (dotted line)




Figure 9 : Measured wind speed data from Crete




Fig. 10 : Typical wind turbine power curve




Figure 11 : Wind speed prediction errors for varying look ahead time and number of previous data points (NX) using Shetland Scroo hourly wind data [model AR(6)]




Figure 12 : Wind speed prediction errors for varying look ahead time and number of previous data points (NX) using Shetland Scroo 10 minute wind data [model AR(6)]




Figure 13 : Wind speed prediction errors for varying look ahead time and number of previous data points (NX) using Crete hourly wind data [model AR(12), predictions every 12 hours, starting at data point 1201]




Figure 14 : Improvement obtained by two fuzzy models.




Figure 15 : Wind speed forecasting AR(12) model - best prediction (solid line), persistence prediction (dashed line), 95% confidence interval (dotted line), actual readings (stars)