



HAL
open science

On-line Assessment of prediction risk for wind power production forecast

Pierre Pinson, Georges Kariniotakis

► **To cite this version:**

Pierre Pinson, Georges Kariniotakis. On-line Assessment of prediction risk for wind power production forecast. EWEC 2003 - European Wind Energy Conference, Jun 2003, Madrid, Spain. hal-00528803

HAL Id: hal-00528803

<https://minesparis-psl.hal.science/hal-00528803>

Submitted on 6 Feb 2018

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.

On-line Assessment of Prediction Risk for Wind Power Production Forecasts.

P. Pinson,

G. Kariniotakis*

Ecole des Mines de Paris,
Center for Energy Studies,
France.

* georges.kariniotakis@ensmp.fr, tel: +33-493957501, Ecole des Mines de Paris, Centre d'Energétique, BP 207, 06904 Sophia-Antipolis, France.

Abstract—The paper introduces a new methodology for assessing on-line the prediction risk of short-term wind power forecasts. The first stage of this methodology consists in computing confidence intervals with a confidence level defined by the end-user. The resampling approach is used, which in contrast to existing methods for wind forecasting, does not need to make a restrictive hypothesis on the distribution of the errors. To account for the nonlinearity of the power curve and the cut-off effect, the errors are classified using appropriate fuzzy sets. The confidence intervals are then fine-tuned to reduce their width in case of stable weather conditions. For this purpose, an appropriate index, named as *MRI*, is defined reflecting the spread of the available Numerical Weather Predictions. A linear relation between that index and the resulting prediction error is shown. The second part of the methodology is to use the *MRI*-index itself as a preventive on-line tool to derive signals for the operator on the meteorological risk, i.e. the probabilities for the occurrence of high prediction errors depending on the weather stability. Evaluation results of this methodology over a one-year period on the case study of Ireland are given, where the output of several wind farms is predicted using a dynamic Fuzzy-Neural Networks based model. The proposed methodology is generic and can be applied to all kinds of wind power prediction models.

Keywords—Wind power, short-term forecasting, confidence intervals, prediction risk, on-line software, adaptive fuzzy-neural networks, numerical weather predictions.

I. INTRODUCTION

NOWADAYS, wind park installations in Europe exceed 23 GW, while the motivated by the Kyoto Protocol targets of the E.U. for 12% energy demand covered by renewables by year 2010, are translated to 21% electricity generation by renewables. To achieve these targets, wind power in the Member States should arise up to 45-60 GW. Such a large-scale integration of wind generation causes several difficulties in the management of a power system. Often, high level of spinning reserve is allocated to account for the intermittent profile of wind production, thus reducing the benefits from the use of wind energy. Predictions of wind power production up to 48 hours ahead contribute to a secure and economic power system operation.

Apart from spot forecasts of the wind parks output in the next hours, of major importance is to provide tools for assessing on-line the accuracy of these forecasts. Tools for on-line evaluation of the prediction risk are expected to play a major role in trading wind power in a liberalized electricity market since they can prevent or reduce penalties in situations of poor prediction accuracy. In practice today, uncertainty is given in the form of confidence intervals or error bands around the spot wind power predictions.

Typical confidence interval methods, developed for models like neural networks [1]-[5], are based on the assumption that the prediction errors follow a Gaussian distribution. This however is often not the case for wind power predictions, where error distributions appear some skewness, while the confidence intervals are not symmetric around the spot prediction due to the shape of the wind turbines power curve. Moreover, the level of predicted wind speed introduces some nonlinearity to the estimation of the intervals; i.e. at the cut-off speed, the lower power interval may switch to zero due to the cut-off effect. The limits introduced by the wind farm power curve (min, max power) are taken into account by the method proposed in [6], which is based on modelling errors

using a β -distribution, the parameters of which, have to be estimated by a post-processing algorithm. This approach however is applicable only to “physical”-type of models since such models estimate power using an explicit wind park power curve, which is not necessarily the case for statistical or artificial intelligence based models as the ones considered here [7].

In [8], [9] wind speed errors are classified as a function of look-ahead time and then they are transformed to power prediction errors using the wind turbine power curve vs. wind speed. This method however is also limited for application to physical models rather than statistical ones since it requires local wind speed predictions (at the level of the wind farm), while it does not provide uncertainty as a function of a pre-specified confidence level. The wind speed errors are estimations provided by the Numerical Weather Prediction (NWP) model. As a consequence, this method does not take into account the modelling error itself that might be due to the spatial refinement of weather predictions or to the power curve used. On the other hand, wind speed measurements are required, which, in general, might not be available on-line.

This paper proposes a methodology for assessing on-line the uncertainty of wind power predictions by the joint use of appropriately defined confidence intervals and prediction risk indices.

Initially, a generic approach is developed for the estimation of confidence intervals that can be applied to both “physical” and “statistical” wind power forecasting models. This is due to the fact that no hypothesis is made about the distribution of the prediction errors. The method accommodates both modelling errors and errors based on the NWPs. It uses past wind power data, which are often available on-line by a Supervisory Control and Data Acquisition (SCADA) system, as well as NWPs, which are the basic input to all models.

Generally, when confidence higher than 80% is required, the

intervals are quite wide. This can lead to conservative or costly managing strategies of the predicted wind power (i.e. allocation of high spinning reserve). Given that confidence intervals are estimations of the uncertainty based on the past performance of the model, the second objective of this work is to develop additional preventive tools able to assess on-line the prediction risk as a function of the forecasted weather situation. The aim is then to provide comprehensive information to the operators so that they are able to adjust the risk they are going to undertake when managing the predicted wind power, i.e. take low risk when forecasted weather situation is unstable.

The paper presents detailed results from the application of the method on the case study of Ireland, where the aim is to predict the output of several wind farms for 48 hours ahead using on-line measurements and predictions from Hirlam NWP system. Evaluation is based on one-year of data.

II. UNCERTAINTY IN WIND POWER PREDICTIONS

Let us define the prediction error for the look-ahead time $t+k$ as following:

$$e_{t+k/t} \triangleq P_{t+k} - \hat{P}_{t+k/t} \quad (1)$$

where $\hat{P}_{t+k/t}$ is the forecast for look-ahead time $t+k$ produced by the model at time origin t , and P_{t+k} is the measured wind power. This error can vary between -100% and 100% of the nominal wind park power. For a non-bounded prediction model it can take values even outside that range. The possible error of the prediction model, defined as “error margin”, depends on the level of measured wind power. Fig. 1 represents graphically the error margin as a function of the wind park characteristic curve.

For wind speeds below cut-in speed, the error margin is maximal since the model can predict a production up to the nominal wind park power. On the contrary, for higher wind speeds the model will show a positive error margin, i.e. the generated power is likely to be greater than the one proposed by the prediction model. Close to the cut-off wind speed the uncertainty is again maximal since the model can switch from a positive error margin to a negative one, or the inverse.




Fig. 1. The error margin as a function of the wind park power curve.

The observed prediction error itself is in general the result of three factors; a modelling error e_{mod} , an error due to the accuracy of the input meteorological predictions e_{NWP} and finally, a stochastic component linked to the process itself e_s :

$$e_{t+k/t} = f(e_{mod}, e_{NWP}, e_s) \quad (2)$$

III. METHODOLOGY FOR ESTIMATING THE UNCERTAINTY OF WIND POWER PREDICTIONS

In this Section, a methodology is developed that permits to assess on-line the prediction risk. It includes:

- (i) Development of confidence intervals for the spot power prediction. The approach is based on the resampling method, which is applied on samples of errors. Errors are classified using fuzzy sets to account for the level of power and the risk for cut-off events.
- (ii) Development of an on-line prediction risk index based on ensembles of NWPs. This Index permits to derive rules for assessing the probability of high or extreme prediction errors due to unstable weather situations.
- (iii) Dynamic fine-tuning of the size of the intervals depending on the weather stability. This permits to avoid excessive risk or to take preventive actions in situations where high errors are expected.

A. Error pre-processing based on fuzzy sets modelling

The first step before computing confidence intervals is to collect the prediction errors the model made in the past. The intervals that are going to be computed will rely on the most recent information on the model’s performance. For this, a window in the past (a certain number of hours) is defined and used as a sliding window for storing the errors. The size of this window defines the size of the errors sample. A separate sample is developed for each look-ahead time of the prediction horizon (i.e. for 1-hour ahead, 2-hour ahead, and so on). This is because prediction errors depend on the look-ahead time as illustrated in Fig. 2. The collected errors are the most recent ones at a given time: when the actual measured wind power is known, that value is compared with all the past predictions made for that time (from 48 hours to 1 hour ago).




Fig. 2. The distribution of the prediction errors varies as a function of the prediction horizon (left picture: 1-hour ahead prediction error distribution, right picture: 24-hour ahead prediction error distribution). Results are produced using data for a single wind farm in Ireland.

The power prediction errors depend on the errors involved in the prediction of wind speed by the NWP system [8-10]. Due to its shape, the wind park power curve is able to amplify (between cut-in and rated speed) or to reduce (below cut-in speed or between rated and cut-off speed) the uncertainty introduced by the NWPs. To account for this effect, the wind power curve is divided into three ranges of power: low, medium and high, which are characterised by fuzzy sets. The prediction errors are classified then as a function of these three ranges (Fig. 3). Hence, the confidence interval estimation is carried out using the error samples corresponding to the power class of the predicted power.

In a similar way, in order to deal with the risk due to the cut-off event, the universe of discourse of wind speeds is divided into two ranges corresponding to a “no cut-off risk zone” for low wind




Fig. 3. Splitting the power curve into three power class fuzzy sets and into two cut-off risk zones.

speeds, and to a “cut-off risk zone” for wind speeds close or higher than cut-off. An appropriate fuzzy set is associated to each zone as shown in Fig. 3. The cut-off risk is considered in priority compared to the power class splitting.

B. Confidence interval estimation by the resampling approach.

Here is a formal definition of confidence intervals: the interval computed from the sample data which, were the study repeated multiple times, would contain the true effect CL% of the time, CL being the confidence level.

A given set of observations (the sample) is a part of a whole population and can be seen as representative. The aim of methods like resampling is to have a better idea of the population distribution by going through the sample a high number of times. This evaluation of the population distribution can serve to estimate a mean, a variance, etc. No assumption is made concerning the distribution: that is the main reason why resampling is preferred to other methods like methods based on simple standard deviation for instance.

Let us consider a sample containing N observations of a mean m for a given stochastic process. The procedure to compute from this sample the 95% confidence interval for that mean involves the following steps:

- (i) N values are selected randomly and with replacement out of the original sample in order to create a new sample;
- (ii) the new sample is sorted in ascending order;
- (iii) the 2.5% lowest and 97.5% highest value of that set are determined.




Fig. 4. One step of the resampling process.

This procedure represents one step of the resampling process (Fig. 4). Indeed, these three actions are to be repeated a large number of times to re-create the population again and again. One always gets a new sample that is close to the original one, and the whole population distribution is not really simulated by this way.

But, by calculating the mean of respectively the 2.5% lowest and the 97.5% highest value of these randomly created samples, good estimates of the confidence limits for the mean m can be computed [1], [11], [12].

In the case of wind power forecasting, the resampling method is applied by considering error samples defined as a function of the look-ahead time, the power range and the wind speed range. These error samples are treated one after the other using the procedure described above, assuming that the prediction error the model makes is the mean of a distribution and that we would like to compute confidence intervals for that mean.

IV. INFLUENCE OF NUMERICAL WEATHER PREDICTIONS

Not only the power prediction model can be responsible for low quality forecasts, but also the numerical weather prediction system (i.e. due to low weather stability). Indeed, an unstable atmospheric situation can lead to very poor numerical weather predictions and thus to worthless wind power ones. In contrast, when the atmospheric situation is stable, one can expect more accurate wind power predictions from the model.

In general it is very difficult to validate the accuracy of numerical weather predictions since wind speed measurements are often not available on-line. This was the situation in the case-study examined in this paper, where only wind power measurements are available in parallel to Hirlam NWP of wind speed at four levels. In order to study the relevance of the NWPs, a dynamic approach is developed based on their correlation to the measured power. The aim is to estimate the probability of situations where Hirlam fails to predict local conditions for a certain period of time (i.e. due to local weather situations). For this purpose, cross-correlations between wind power and Hirlam wind speeds were estimated using a sliding window of 100 hours. The result is plotted in Fig. 5 for a period of 2000 hours. The distribution of the obtained correlations is shown in Fig. 6 for each level. The range of the values is between $\{-0.4, 0.92\}$. Low




Fig. 5. Dips below 0.6 of the cross-correlation show situations where poor accuracy of the wind power predictions should be expected.




Fig. 6. Distribution of the cross-correlation coefficients for various Hirlam levels.

values indicate situations where low reliability of Hirlam forecasts might be expected. The frequency of these periods is however limited since the distributions are centred on the 0.8 correlation value. An additional conclusion is also that the most correlated Hirlam level varies in time. This is an indication that a performing prediction model should consider more than one level as input.

In the following Paragraphs, the information included in the NWP is exploited to develop tools for on-line estimation of the uncertainty in power predictions.

A. Ensemble forecasts for the assessment of weather stability

Meteorological Centres are able to produce different scenarios of Numerical Weather Predictions by perturbing the initial conditions of the forecasting model or by using different NWP models. These scenarios are called ensemble forecasts and permit to evaluate the stability of the weather regime [13]. However, for wind power applications only one forecast for the next 48 hours is often made available (or purchased) at a given time. For instance, Hirlam gives a unique 48-hour ahead forecast every 6 hours. Nevertheless, for a given hour, several predictions can be available from different time origins in the past (-6, -12, -18... hours). In a stable and well-predicted weather situation it is expected that these predictions will not differ significantly. Weather stability can be assessed by comparing all the available forecasts for the considered period.




Fig. 7. Stable (left picture) and unstable (right picture) weather situations.

In order to have a general evaluation of that stability, 4 sets of predictions of various ages (0, 6, 12 & 18 hours) for the following 24 hours are compared. Figure 7 gives the examples of a stable atmospheric situation (left picture, the forecasts are quite close) and of an unstable one (right picture, spread forecasts).

B. Development of a norm to assess the weather stability

There are several possibilities to measure the spread of the various weather forecasts. In [10] the standard deviation of the forecasts for each time-step is mentioned as an example.

Our aim here is to evaluate the global atmospheric situation. This is why a unique representative index is defined for the following N_h hours instead of indexes for every look-ahead time.

In order to calculate the distance between two sets of forecasts, we propose a 2-norm between the N_h -valued vectors containing the predicted wind speed for the N_h following hours. Lets define as $\overline{ws}_{t-\alpha} \hat{=} [ws_{t+1/t-\alpha}, \dots, ws_{t+k/t-\alpha}, \dots, ws_{t+N_h/t-\alpha}]^T$ the N_f available sets of wind speed forecasts (NWPs), with α being the age of each set. The values for α can be 0, 6, 12, etc, for the case of Hirlam. The distance between the predictions of ages α_i and α_j is given by:

$$d(\overline{ws}_{t-\alpha_i}, \overline{ws}_{t-\alpha_j}) \hat{=} \left(\frac{1}{N_h} \sum_{k=1}^{N_h} (ws_{t+k/t-\alpha_i} - ws_{t+k/t-\alpha_j})^2 \right)^{\frac{1}{2}} \quad (3)$$

Then, an index, called hereafter “meteo-risk” *MRI*-index, is defined to measure the spread of the weather forecasts at a given time. It uses the most recent forecast as a reference and reflects the variability of the older forecasts:

$$MRI \hat{=} \sum_{i=1}^{N_f-1} p_i \cdot d(\overline{ws}_{t-\alpha_0}, \overline{ws}_{t-\alpha_i}) \quad (4)$$

with p_i ($i = 1, \dots, N_f-2$) being appropriate weights defined so that:

$$\begin{cases} p_i > p_{i+1}, & i=1, \dots, N_f-2 \\ \sum_{i=1}^{N_f-1} p_i = 1 \end{cases} \quad (5)$$

The use of the weights p_i permits to give more importance on the recent information we get from the weather predictions.

C. Relation between weather stability and wind power prediction error

In the frame of the case study of the paper, the horizon N_h for the calculation of the *MRI*-index is set to 24 hours. Since Hirlam forecasts are provided every 6 hours, there are four sets ($N_f=4$) of wind speed predictions covering the period. However, the same methodology can be applied to seven available sets of Hirlam forecasts on a 6-hour period for instance.

Figure 8 shows the distribution of the values of the *MRI*-index as obtained using the data of one of the considered wind farms in Ireland (see Section V). It can be concluded that most of the times the weather regime is quite stable (low *MRI*-index values) and that there are only few occurrences of really spread forecasts. Figure 7 shows the patterns of a “stable” and an “unstable” weather situation with *MRI*-index values of 0.3 and 2.9 respectively.




Fig. 8. The distribution of weather situations as expressed by the *MRI*-index.

For the same case study we collect wind power prediction errors as obtained by a fuzzy-neural network (F-NN) model [7] for a period covering one year. For the same period the *MRI*-index is estimated. By binning the data, calculating the average error for each bin, and comparing these averages to the global prediction error of the model, the representative points in Fig. 9 are obtained. The prediction error increases linearly with the *MRI*-index: the tighter the Hirlam predictions are, the more accurate the wind power prediction model is. A linear fitting gives the solid curve shown in Figure 9.

Making this assumption would mean that the prediction error of the model made follows an affine empiric relation:

$$e = e_0 + s \cdot MRI \quad (6)$$




Fig. 9. Prediction errors vs. *MRI*-index over a one year dataset: there is a linear relation between the prediction error and the *MRI*-index.

which is composed by a basic part of the error e_0 and by a NWP-dependent error, the latest being a direct consequence of the prediction model sensibility to the weather stability. The slope s of the linear fitting model represents that sensibility.

D. Use of the *MRI*-index to adjust on-line the prediction risk

1) Fine-tuning of the confidence intervals

The relation (6) indicates that when the *MRI*-index is low, the model is expected to be more accurate. In that case one would be ready to accept tighter confidence intervals for the predictions. The aim here is to use Eqn. (6) to define a scale factor for the confidence intervals depending on the value of the *MRI*-index. This scale factor can be applied to either enlarge or narrow the intervals width in the following N_h hours. For instance, when the meteorological index equals 0.5, the size of the intervals for the following 24 hours is reduced by almost 20%. The strategy chosen here is to only narrow the intervals when the *MRI*-index permits to do so. It can be seen from Fig. 8 that most (~65%) of the time, the *MRI*-index allows to decrease the interval size.

2) On-line use of the *MRI*-index for monitoring prediction risk

The relation that was drawn above permits to define rules concerning the expected prediction error depending on the *MRI*-index value. For that purpose, we bin the data and calculate the cumulative distribution function of the prediction errors for each bin (Fig. 10).




Fig. 10. Cumulative distribution functions of the prediction error depending on the *MRI*-index bin.

These curves give the probability with which, an error larger than a defined threshold occurs, depending on the value of the *MRI*-index. For instance, if at a certain time, the index takes a value

between 0 and 1, there will be a probability of 36% that an error $e_{+k/t}$ larger than the mean prediction error $\bar{e}_{+k/t}$ occurs. However, if at that same time the value of the index is within the [1.75-2] bin, the probability for the same kind of error is twice larger (75%):

$$\text{IF } MRI \in [1.75-2] \text{ THEN } e_{+k/t} > \bar{e}_{+k/t} \text{ with 75\% probability. } (7)$$

Table I gives the probability for errors to be larger than 1, 1.5 and 2 times the average error depending on the range of the *MRI*-index. The Table is estimated for the case of 5 wind farms in Ireland. Based on such a Table, several rules similar to the one given by (7) can be derived.

Table I also provides information on the probability of extreme prediction errors to happen (as extreme are defined errors larger than twice the mean prediction error of the model). Actually, for these 5 wind farms in Ireland, when the *MRI*-index takes low values (between 0 and 1) an extreme prediction error is unlikely to happen, and that is not the case if this one is within the bin [1.75-2] (17% probability of occurrence). On the other hand, if $MRI > 1.75$, an error of at least 50% of the average prediction error is expected.

In an on-line environment these rules permit to derive signals or alarms for the end-user of the wind power prediction model, informing that large prediction errors might occur. Then, the operator can consider such a signal for:

- taking preventive actions (i.e. increase spinning reserve),
- considering the lower interval, rather than the spot prediction of power, when trading, in order to avoid penalties, etc.

In the developed software, the *MRI*-index is monitored as shown in Fig. 14. The passage of the index to high zones triggers messages to the operator on the expected risk. The rules that are implemented evolve over time thanks to the storage of *MRI*-index values and prediction errors.

Probability of occurrence of errors larger than n times the mean prediction error	Ranges of <i>MRI</i> -index				
	Bin [0, 1]	Bin [1, 1.25]	Bin [1.25, 1.5]	Bin [1.5, 1.75]	Bin [1.75, 2]
$n = 1$	36	51	62	67	75
$n = 1.5$	10	22	28	30	38
$n = 2$	3	8	10	12	17

Table I. Rules for the occurrence of larger errors depending on the value of the *MRI*-index.

V. RESULTS

Results are presented for five real wind farms in Ireland (WF-A to WF-E) with a total installed power of a few tens of MW. The prediction model is the Adaptive Fuzzy-NN model described in [7]. The available time series cover a period of almost two years from which 6600 hours were used for training (learning set), 1000 hours for cross-validation and one year for testing the performance of the model. The results presented here are on the testing set. Concerning the computation of confidence intervals, 12 days of prediction errors are stored in the samples. The desired confidence level is set to 85%.

Figure 11 depicts an episode with the wind power predictions for the next 43 hours compared to the real values for WF-E. The 85% confidence intervals are built with the method described above. In order to illustrate the fine-tuning of the intervals, Fig. 12 gives the example of a weather situation classified as stable with respect to the “meteo-risk” index. For the first 24 look-ahead times the resampling on past errors produces quite broad intervals,




Fig. 11. Wind power prediction with 85% confidence intervals.

but their size is reduced by almost 20% afterwards thanks to the consideration of the weather situation.

The resampling approach that is used to design the confidence intervals displays a dynamic behaviour thanks to the updating of the sample of errors and also thanks to the fuzzy set modelling of the power curve. Figure 13 shows the evolution of the 6-hour ahead interval width during 1000 hours of the evaluation period. An interesting point is that these intervals are non-symmetric because they follow the concepts we described about the error margin – for instance the prediction error is more likely to be negative than positive when forecasting low power output.

Moreover, the interval size varies a lot over time. For the example of Fig. 13, the width of the 6-hour ahead confidence interval ranges from 10% to 70% of the wind farm nominal power.




Fig. 12. Wind power prediction with the resampling and fine-tuned confidence intervals. The intervals are narrowed for the first 24 hours due to a low MRI-index value.




Fig. 13. Evolution of the 85% confidence 6-hour ahead interval width over a 1000-hour period for the resampling intervals and the fine-tuned ones.

Depending on the weather stability described by the *MRI*-index, the intervals are tuned: when the weather situation is considered as stable (low index values) they are narrowed up to 30%. The dashed curves in Fig. 13 correspond to the final intervals.

The performance of the confidence interval estimation for the five wind farms in Ireland is summarized in Tab. II, i.e. the observed confidence at the end of the evaluation period (one year) for both resampling and fine-tuned intervals, as well as the percentage of reduced intervals and their average reduction due to the use of the *MRI*-index.

Wind farm	Observed confidence of resampling intervals (%) (1)	Observed confidence of fine-tuned intervals (%) (2)	No of times (%) intervals are reduced (3)	Average width reduction (%) (4)
WF-A	84.87	82.93	65.10	11.32
WF-B	82.38	80.31	66.60	8.85
WF-C	81.05	80.53	68.78	7.28
WF-D	83.60	82.16	63.98	7.28
WF-E	84.29	82.85	63.52	8.23

Table II. Observed confidence for the two types of intervals after one year of simulation and effects of the *MRI*-index on the interval reduction.

One can see that the consideration of the weather stability permits to narrow the intervals most of the times (~65% in column 3), and the average reduction is up to 11% of their initial size (WF-A). The corresponding confidence loss is not significant (column 2).




Fig. 14. Average width of the 6-hour ahead confidence intervals for WF-A and for various specified confidences (75, 80, 85, 90 & 95%).

The interval size reduction is illustrated in Fig. 14, where the average interval width (for the 6-hour ahead ones) is plotted against the observed confidence. The various interval sizes are quite important, but one has to keep in mind that this study deals with single wind farms and so the level of prediction error is higher than for the case of regional or national wind power predictions where there is some spatial smoothing effect.

Finally, Figure 15 gives a picture of an episode with the evolution of the *MRI*-index and the related prediction error: the period of low “meteo-risk” (time in the evaluation set between 200 and 450) matches with a period of very low prediction error and periods with unstable weather situations (time between 80 and 180 for instance) correspond to higher prediction errors. These features, as well as the mean prediction error of the model, can be given on-line to the operator and help him to seize the prediction risk due to the meteorological situation.




Fig. 15. Parallel evolution of the MRI-index and of the normalized prediction error.

VI. CONCLUSIONS

A generic methodology for assessing on-line the prediction risk of short-term wind power forecasts is presented. At a first stage, confidence intervals based on the resampling approach are derived, taking into account the prediction horizon, the power class and the cut-off risk. At a second stage a new meteorological risk (*MRI*) index was introduced to evaluate the weather stability. The *MRI*-index can be used either to fine-tune the confidence intervals or to give signals to the end-user on the probability of outliers.

The developed methodology was tested over a one-year evaluation dataset for five wind farms located in Ireland. The results are encouraging and consist a first step in the development of on-line tools that can be used in a complementary way to the prediction model itself.

The developed prediction risk method is implemented in the form of on-line module and integrated in the Armines Wind Power Prediction System (AWPPS). In the future this module will be also integrated and available through the ANEMOS advanced prediction platform (<http://anemos.cma.fr>). Apart from the prediction risk module, the core F-NN wind power forecasting models of AWPPS are integrated in the More-Care Energy Management System and installed for on-line operation in Ireland and other sites such as Crete, Madeira, etc.

VII. ACKNOWLEDGMENT

This work was performed in the frame of the ANEMOS Project (ENK-CT2002-00665) funded in part by the European Commission. The authors gratefully acknowledge ESB National Grid for providing data for the realization of the study.

VIII. REFERENCES

- Alves da Silva A. P., Moulin L. S., "Confidence intervals for neural network based short-term load forecasting", *IEEE Trans. on Power Systems*, Vol. 15, No. 4, Nov. 2000.
- Carney J. G., Cunningham P., Bhagwan U., "Confidence and prediction intervals for neural network ensembles", in *Proc. of the International Joint Conf. on Neural Networks 1999*, paper 2090, 1999.
- De Veaux R. et al, "Prediction intervals for neural networks via nonlinear regression", *American Statistical Association*, vol. 40, n°4, Nov. 1998, pp. 273-282
- Dybowski R., Gant V., "Confidence intervals and prediction intervals for feed-forward neural networks", in *Clinical Applications of Artificial Neural Networks*, Ed.: Dybowski-Gant, Cambridge University Press, 2000.
- Heskes T., "Practical confidence and prediction intervals", in *Advances in Neural Information Processing Systems 9*, Ed.: Mozer-Jordan-Petsche, Cambridge, 1997, MIT Press, pp. 176-182.
- Luig A., Bofinger S., Beyer H. G., "Analysis of confidence intervals for the prediction of regional wind power output", in *Proc. of the 2001 European Wind Energy Association Conference, EWEC'01*, Copenhagen, Denmark, pp.725-728, 2-6 July 2001.
- Kariniotakis G., Mayer D., "An advanced on-line wind resource prediction system for the optimal management of wind parks", in *Proc. of the 2002 MedPower Conference*, Athens (Greece), 4-6 Nov. 2002.
- Lange M., Waldl H.-P., "Assessing the uncertainty of wind power predictions with regard to specific weather situations", in *Proc. of the 2001 European Wind Energy Association Conference, EWEC'01*, Copenhagen, Denmark, pp.695-698, 2-6 July 2001.
- Lange M., Waldl H.-P., "Accuracy of short term wind power predictions depending on meteorological conditions", in *Proc. of the 2002 Global Windpower Conference*, Paris, France, 2-5 April 2002.
- Landberg L., "Poor-man's ensemble forecasting for error estimation", *Proc. of the 2002 AWEA Portland Conf.*, Portland, 2-5 June 2002.
- Masters M., *Neural, Novel & Hybrid Algorithms for Time-Series Prediction*, J. Wiley & sons, 1995.
- Simon J. L., *Resampling : The new statistics*, 2nd edition, Oct. 1997, [Online]. Available:<http://www.resample.com/content/text/index.html>
- Roulston M. S. et al, "Value of the ensemble prediction system for forecasting wind energy production", in *Proc. of the 2001 European Wind Energy Association Conference, EWEC'01*, Copenhagen, pp.717-720, 2-6 July 2001.