

Bayesian Inferencing for Wind Resource Characterisation

Marcos S. Miranda, *Member, IEEE*, Gavin Shaddick, Rod W. Dunn, *Member, IEEE*, F. Li, *Member, IEEE* and Keith R. W. Bell

Abstract—The growing role of wind power in power systems has motivated R&D on methodologies to characterise the wind resource at sites for which no wind speed data is available. Applications such as feasibility assessment of prospective installations and system integration analysis of future scenarios, amongst others, can greatly benefit from such methodologies. This paper focuses on the inference of wind speeds for such potential sites using a Bayesian approach to characterise the spatial distribution of the resource. To test the approach, one year of wind speed data from four weather stations was modelled and used to derive inferences for a fifth site. The methodology used is described together with the model employed and simulation results are presented and compared to the data available for the fifth site. The results obtained indicate that Bayesian inference can be a useful tool in spatial characterisation of wind.

Index Terms—Bayesian inference, Markov chain Monte Carlo simulation, spatio-temporal modelling, statistical methods, wind power integration, wind resource characterisation.

I. INTRODUCTION

THE ever-increasing presence of wind power generation in modern power systems has driven much research effort into methodologies to characterise the wind resource and the impact of wind power on the power system. Such methodologies play an increasingly important role as they make it possible to assess incurred risks in terms of the system reliability against possible requirements for higher operational costs or further investments in the system infrastructure.

One area of particular interest is the modelling of wind speeds at locations where little or no information regarding the wind resource is available. In this paper, a methodology for characterisation of the wind speed at a given site is proposed. It uses a Bayesian approach to model the spatial correlation between different sites - weather stations for which wind speed data is available - and infer wind speeds for the new location.

Resource characterisation can be an important tool for assessing the integration of wind power into the power system

This work was supported by the UK EPSRC, through the Supergen initiative.

M. S. Miranda (e-mail: m.miranda@bath.ac.uk), R. W. Dunn (e-mail: r.w.dunn@bath.ac.uk), and F. Li (e-mail: eesfl@bath.ac.uk) are with the Department of Electrical Engineering, University of Bath, Bath, BA2 7AY, UK.

G. Shaddick (e-mail: g.shaddick@bath.ac.uk) is with the Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK.

K. R. W. Bell (e-mail: kbell@eee.strath.ac.uk) is with the Department of Electrical Engineering, University of Strathclyde, 16 Richmond Street, Glasgow, G1 1XQ, Scotland, UK.

and can be used, for example, in the statistical analysis of siting and sizing of wind power plants and its influence on generation capacity adequacy, transmission system security and capability, amongst others.

The methodology developed in this paper uses a Bayesian framework to characterise the known resource (weather stations data) and estimate the wind at the desired new location. The model is implemented in two (hierarchical) levels. At the first level, the wind data is described as the sum of a temporal and a spatial component, plus an unstructured (random) component. At the second level, the temporal and spatial components are modelled as a random walk and a multivariate normal distribution respectively.

This approach allows the temporal characterisation of the wind resource at the desired location while at the same time maintaining the spatial correlation of the given data, thus allowing the equivalent of a virtual weather station to be created. The model was implemented using the OpenBUGS (Bayesian inference Using Gibbs Sampling) software package [1] and used Markov chain Monte Carlo (MCMC) simulation to estimate the model parameters.

In the paper, the methodology is described and implemented using data from four MetOffice (UK) weather stations. Estimates are made for a fifth location (also a weather station) for which wind data is available. Simulation results are presented and compared with the data for the fifth location.

II. SPATIAL CHARACTERISATION OF WIND SPEEDS

Wind speed modelling and characterisation has gained significant attention over the last two decades, with considerably complex implementations becoming feasible as a result of the increase in computational processing speed. This broad subject area can be generally divided into two main areas: temporal and spatial characterisation.

Temporal prediction, sometimes referred to as “short-term prediction”, relates to the forecast of future wind speeds (or wind power), and its prediction horizon is greatly dependent on the application area [2]-[4]. It can range from seconds (e.g. wind turbine control, power quality management, mitigation of stress loads on turbine system), to hours and days (market operations, power system operational security, generation dispatch, maintenance scheduling).

Spatial modelling, on the other hand, mostly applies to the characterisation of the resource at sites where insufficient or no information is available. It can be a useful tool in feasibility

studies of prospective installations and system integration, assessment of the impact on system reliability, generation adequacy, and transmission system capability, amongst others.

Different approaches can be found in the literature regarding the problem of spatial modelling, many applied to climatologic variables (wind speed included) or pollutant concentration analysis. These can either employ physical models, which take the terrain into account, such as WAsP [5], or statistical models. Amongst the latter, we find the Maximum Likelihood method [6], the Kalman filter [7], Bayesian inference [8], [9] and the Measure-Correlate-Predict (MCP) method [10]. The latter is widely adopted by the wind industry and considered the standard methodology for feasibility studies of future projects. A drawback of the MCP technique is the requirement of data (at least 6-12 months) for the candidate site, which implies the installation of an anemometric tower, making it impractical for broad integration studies. The work in this paper tries to overcome this requirement by inferring the wind speeds at an unmeasured location.

III. THE BAYESIAN METHODOLOGY

One important feature of Bayesian statistics is how the probability associated to a variable in a given process is treated. In contrast to the frequentist approach, where probabilities are based on how often events occur in a (hypothetical) sample, in Bayesian inference, results are based on a combination of data and prior beliefs. This results in a posterior distribution, in which degrees of belief are attached to the variable [11].

Bayes' theorem (1) states that the probability of a variable A given the occurrence of another variable B is equal to the normalised likelihood of B given A ($P(B|A) / P(B)$) times the probability of A . In other words, the conditional probability of A on B can be found using the conditional probability of B on A , and the degree of belief on A , the prior $P(A)$.

$$P(A|B) = \frac{P(B|A) * P(A)}{P(B)} \quad (1)$$

A prior probability is a marginal probability, i.e. not conditional, defined on some knowledge of what value a variable may assume, but not necessarily based on observations of the variable, which is not always possible.

The posterior probability, on the other hand, is a conditional probability, which takes into account the information contained in the prior and the normalised likelihood, as shown above.

In many cases, it is not possible to derive the posterior distributions analytically, but samples may be generated using Markov chain Monte Carlo (MCMC) methods. The result is a distribution of samples from the posterior distribution rather than just a single parameter estimate. Empirical summary statistics can be calculated from this distribution and used to draw inferences about their true values, such as the expected value and associated uncertainty (credible intervals).

In MCMC, a 'burn-in' stage during which model parameters converge is followed by a period in which the

samples are stored to estimate the posterior distribution. Only these remaining values are used in the analysis.

A. Hierarchical Models and Expert Knowledge

One important aspect of the Bayesian approach is the ability to implement nested structures, also referred to as hierarchical models. This enables a more detailed description of the variables in the model and the introduction of expert knowledge at each of the levels, if such information is available. This feature also allows the use of data from various sources recorded at different resolutions.

This can be exemplified by considering a fictitious model of hourly electricity demand over the winter. The overall demand distribution may be modelled by a normal distribution, the mean of which could be time-dependent (time-series). Further, the time-series data available may be modelled by two additive functions, the domestic and the industrial consumption. The individual time-series models themselves may have any desired structure, e.g. an autoregressive process. The domestic consumption can further be made dependent on the average daily temperature, which in turn can be characterised using any desired formulation.

In the specific case of wind speed modelling, such feature allows the inclusion of, for instance, other climatologic variables or seasonal variations.

IV. MODEL DEVELOPMENT

A Bayesian hierarchical model was developed to model the spatial correlation structure in the data set and the individual sites temporal autocorrelation structure, thus allowing the characterisation of the wind speeds at a prospective site.

A. Input Data

The data used in the model were one year of hourly average wind speeds from four weather stations at the north of Scotland (Aultbea, Tain Range, Duirinish and Altnaharra). Wind speed estimates were calculated for the target site of Loch Glascarnoch, where a weather station is also present. A map with the location of the four reference stations and their respective distance to the estimated site is shown in Fig. 1.

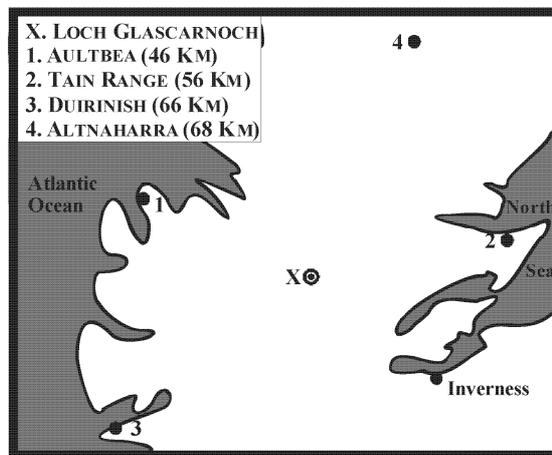


Fig. 1. Location of target and reference weather station sites.

A summary of the weather stations wind characteristics for the year studied (1997) is presented in Table I. The height above ground level at which measurements were taken follows the Met Office guidelines of standard exposure (over level, open terrain, at 10m above the ground). The estimated location is represented by S_X . On average, around 4% of the data is unavailable; however the methodology used allows the inclusion of such points as missing data.

TABLE I
SUMMARY STATISTICS OF WEATHER STATION DATA

Site	Summary Statistics				
	Missing Data (%)	Mean Speed (m/s)	Variance	Maximum Speed (m/s)	Prevailing Direction (0°- East)
S_X	7.91	4.48	11.90	20.09	150°
S_1	1.39	4.90	11.51	21.12	240°
S_2	4.11	4.58	8.32	18.54	240°
S_3	2.56	5.20	15.63	25.75	210°
S_4	7.81	3.97	9.80	20.60	210°

The correlation/covariance between the four input stations is shown in Table II. Since all stations had missing data points, the correlation coefficients shown in the table were calculated using only time-steps in which data was available for all stations (a total of 7662 points out of 8760).

TABLE II
CORRELATION/COVARIANCE BETWEEN WEATHER STATION DATA
(CORRELATION BELOW AND COVARIANCE ABOVE DIAGONAL)

		Covariance			
		Site	S_1	S_2	S_3
Correlation	S_1	--	6.4083	11.1908	7.0866
	S_2	0.6527	--	7.0320	6.8288
	S_3	0.8319	0.6148	--	7.7360
	S_4	0.6639	0.7544	0.6230	--

B. Model Structure

The hierarchical structure implemented in this work has two main modelling levels:

- Level 1: The wind speed data from the reference weather stations is defined as the sum of a temporal and a spatial component, plus an unstructured error component, which represents the information contained in the data that is not explained by the other two components. The data (observation) equation can be written as:

$$U_w(s,t) = \theta(t) + M(s) + \varepsilon_U(s,t) \quad (2)$$

where $\theta(t)$ and $M(s)$ are the temporal and spatial components, respectively. $\varepsilon_U(s,t) \sim N(0, \sigma_U(s)^2)$ is the unstructured error, $t=1, 2, 3, \dots, 8760$, and $s=1, 2, 3, 4$.

- Level 2: The temporal part is modelled as a first order random walk, with a zero-mean random term, $\varepsilon_\theta(3)$,

$$\theta(t) = \theta(t-1) + \varepsilon_\theta \quad (3)$$

and the spatial part is modelled as a multivariate normal

distribution (4), with mean $\mu_M(s)$, and variance σ_M^2 .

$$M(s) = MVN(\mu_M(s), \sigma_M^2) \quad (4)$$

In any Bayesian hierarchical model, there will be a final level of modelling which assigns ‘hyper-priors’ to the parameters in the higher levels of the model.

The temporal component in this model follows a first order random walk, which gives flexibility in that assumptions of stationarity do not have to be made.

The spatial component is incorporated into the covariance matrix (scaled by a spatial variance, σ_M^2) of the multivariate normal distribution described in (4). Correlations are proportional to the distance, d , between sites, $f(d) = \exp(-\phi \cdot d)$. Predictions can then be obtained for unmeasured locations conditional on the measured values. The prior distribution for the rate of decay of the correlation function, ϕ , was chosen to be a uniform distribution with limits corresponding to expected drops in correlation, ρ , over distance, which can be estimated using the relationship $\phi = -\log(\rho)/d$.

C. Non-normality correction

It is a well known characteristic of general wind speed series that its variation at a given site can be modelled using the Weibull distribution [12]. As described above, this is a potential problem for the hierarchical model adopted since it assumes normally distributed errors at all levels (temporal, spatial and unstructured components).

Therefore a transformation of the original wind speed data was required and the Box-Cox transformation was used as it presents a simple and straightforward procedure for non-normality correction [13], as detailed below.

The Box-Cox transformation $y^{(\lambda)}$, of a dataset y is defined as:

$$\begin{cases} y^{(\lambda)} = \frac{y^\lambda - 1}{\lambda}, & \text{for } \lambda \neq 0 \\ y^{(\lambda)} = \ln(y), & \text{for } \lambda = 0 \end{cases} \quad (5)$$

The choice of the value for the parameter λ can be made through an analysis of the log-likelihood function:

$$f(y, \lambda) = -\frac{n}{2} * \ln \left[\frac{\sum_{i=1}^n (y_i(\lambda) - \bar{y}(\lambda))^2}{n} \right] + (\lambda - 1) * \sum_{i=1}^n \ln(y_i) \quad (6)$$

The value of λ which maximises (6) is used in the Box-Cox transformation (5). A plot of the log-likelihood function using data from each of the four reference weather stations can be seen in Fig. 2. The function is calculated for values of λ between [-2,2] in steps of 0.1.

From Fig. 2, it can be seen that the value of λ which maximises the log-likelihood function for the four reference sites is within the range [0, 0.3]. For consistency, the same data transformation should be applied across all sites, therefore the value $\lambda = 0.2$ was chosen.

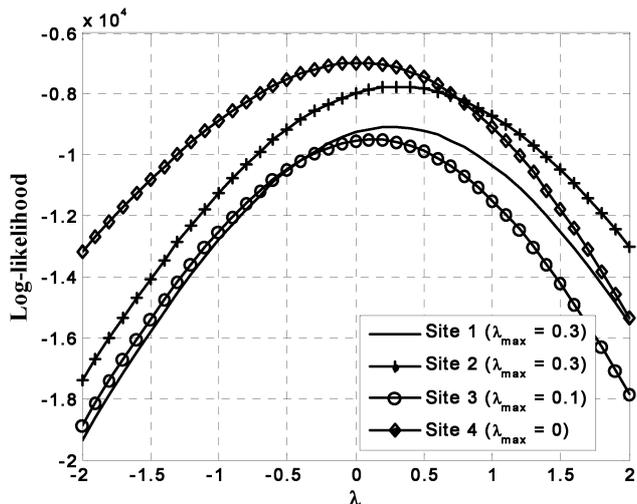


Fig. 2. Log-likelihood functions of wind speed data for the four reference weather stations as a function of λ in the Box-Cox transformation.

To validate the procedure used, Matlab's *boxcox* function was also used. This function performs a continuous assessment of the data on the variable λ , searching the resulting function maximum through an optimisation routine. This yielded the following values $\lambda_{MAT} = [0.2768, 0.3178, 0.1223, 0.0301]$ for the four sites, showing good agreement with the results obtained using discrete values for λ . The difference between the results was not significant enough to justify the use of the more complex calculation procedure, which may not be readily available to many.

V. SIMULATION RESULTS

In the analysis of MCMC simulation results, the initial outputs are usually discarded to allow for the convergence of model variables. The following results were taken for a run of 6,000 samples, the first 1,000 were discarded and only the remaining 5,000 used in the analysis. Convergence was assessed by visual examination of the "time series" plots of the samples.

The first stage in the model assessment was the analysis of the residuals in (2), i.e. the estimates of the unstructured error component, $\varepsilon_U(s,t)$, for the four reference stations, as depicted in Fig. 3. It is worth emphasising that the values plotted have not been transformed back (reverse of the Box-Cox transform). There is no noticeable trend/cycle in the residuals over time or discrepancies/patterns across the stations, the presence of which would have indicated unspecified components of the model and the possibility of non-stationarity.

A comparison between the data available from the estimated site and the model output is shown in Fig. 4. It is worth highlighting that the model output includes only the temporal and spatial components, not accounting for the unstructured error component, which will be present in the data. For better visualisation, the one-year estimate (8760 points) was split in two halves in this figure. The agreement between both curves is very good.

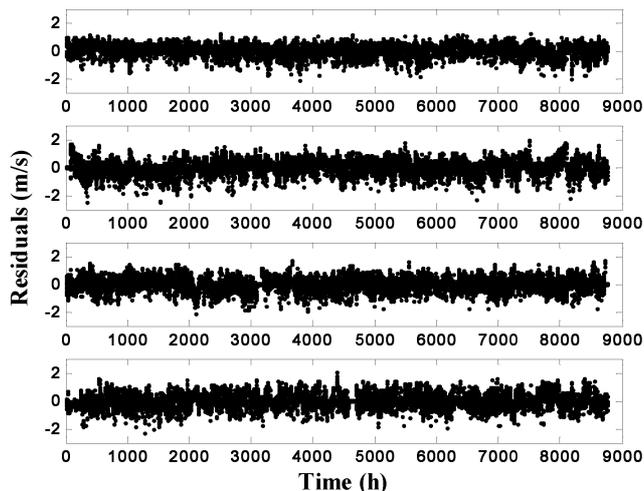


Fig. 3. Model residuals (unstructured error term) for transformed data. Reference stations 1 to 4, from top to bottom.

As explained in *Model Structure* (section IV.B), the unstructured error component was assumed to be normally distributed with zero mean, therefore the data and estimate average values should be very close. Indeed, the average wind speed calculated for the site was $\bar{U}_{w,data} = 4.48 m/s$, whereas the model output yielded $\bar{U}_{w,pred} = 4.49 m/s$.

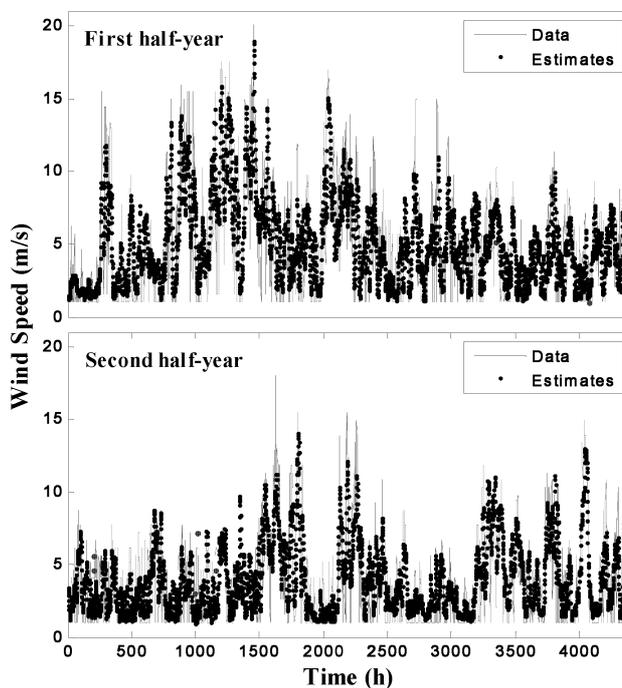


Fig. 4. Comparison between time series for target weather station wind data and model estimates.

Although predictions from the model lack the random error component, an estimate of its confidence intervals was carried out using the parameters found for the reference weather stations. In the model described in (2), the standard error,

$\sigma_U(s)$, of the random component, $\varepsilon_U(s,t)$, can be found for each weather station ($\sigma_U(s) = [0.468, 0.577, 0.535, 0.561]$, for the transformed data). Since these values are very similar, a simple average, $\hat{\sigma}_U = 0.535$, was used as an estimate of the standard deviation for the target site. The 95% confidence interval was then obtained using:

$$U_{w,conf}(t) = U_{w,pred}(t) \pm 1.96 * \hat{\sigma}_U \quad (7)$$

This calculation was performed before transforming back the data. A plot of the data, estimates and the 95% confidence interval for 10 days in early September (data points [6000, 6240]) is shown in Fig. 5.

A linear regression fit of the data against the estimates is shown in Fig. 6, together with the regression equation and the calculated R^2 value. From the plot it can be seen that the model tends to overestimate at lower s and underestimate at higher speeds. Since there is limited practical interest in accurately knowing the temporal behaviour of past wind speeds, a histogram of the data and estimates is also shown in Fig. 7. The frequency of occurrence of each wind speed range shows a good degree of matching between the data and the estimates.

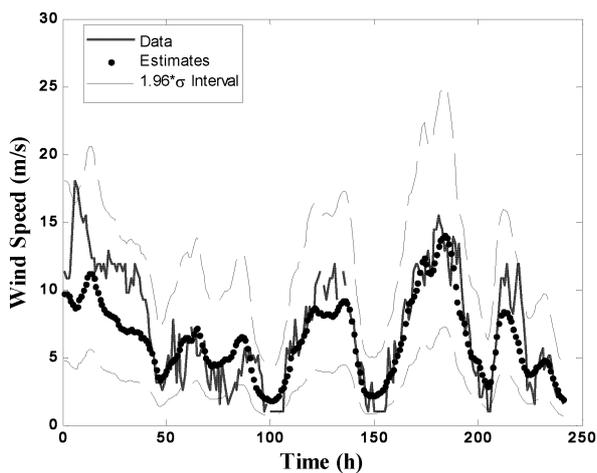


Fig. 5. Zoom of time-series plot for 10 days, with 95% ($\pm 1.96 * \hat{\sigma}_U$) confidence intervals included.

Posterior estimates of the unstructured component of the model are available at each of the site locations in terms of the variances of $\varepsilon_U(s)$. The standard deviation (after transforming back the data) for these estimates are very similar across the four reference sites, $\sigma_U(s) = [1.316, 1.694, 1.821, 1.70]$, indicating consistent model fit over each of the sites. Likewise, such samples can also be obtained for the target site, in the form of the difference between the available data and the estimates. The resulting distribution can be seen in Fig. 8. As expected, the unstructured error mean is close to zero, and its standard deviation is well within the range of those found for the reference stations.

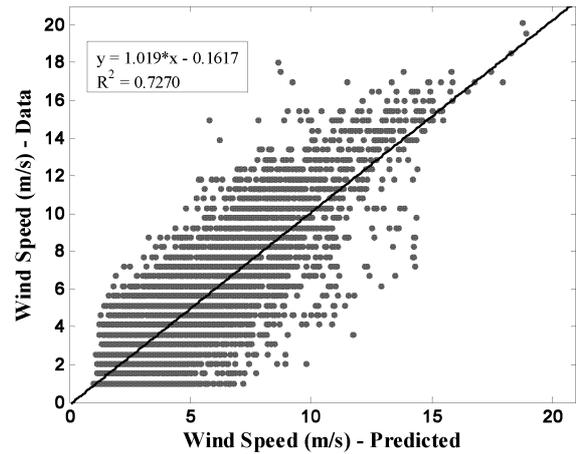


Fig. 6. Linear regression plot of wind speed data against the estimated values.

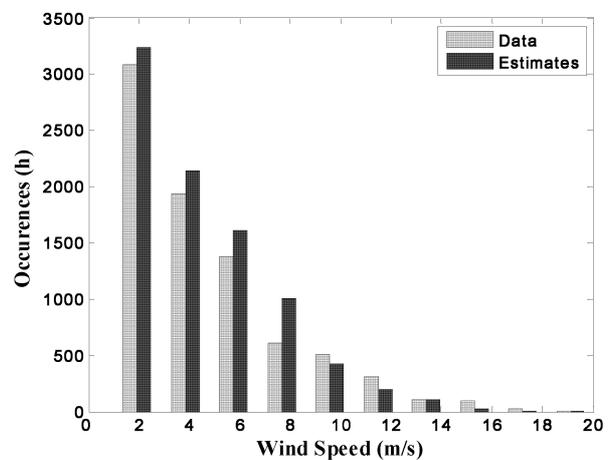


Fig. 7. Histogram comparing wind speed data and model estimates.

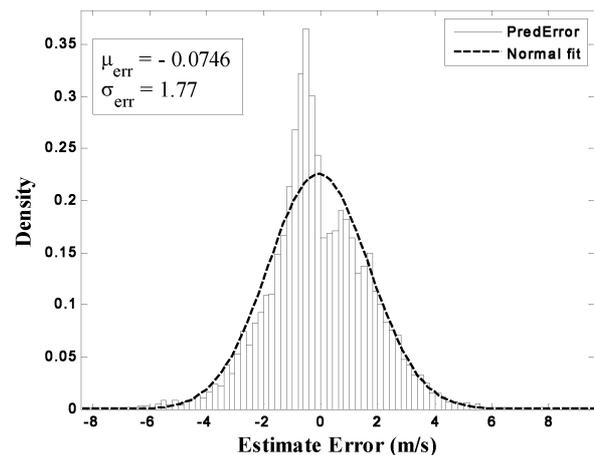


Fig. 8. Histogram and normal distribution fit of the random component samples ($U_{w,data} - U_{w,pred}$) at target site.

A. Discussion of Results

The results obtained were encouraging with respect to the application of Bayesian inferencing to spatial wind speed characterisation. The temporal and spatial components were

responsible for 35% and 5% of the uncertainty in the model, respectively. This greater contribution from the temporal component was expected, given the number of data points available relative to the spatial component (8760 time samples, against 4 spatial “points”). It is possible that the use of more and further dispersed reference stations could improve the contribution of the spatial component. Similarly, a more complex temporal structure could better represent the data, reducing the uncertainty associated with the unstructured component and thus improving the estimates. Another alternative would be the inclusion of other explanatory variables, such as atmospheric pressure, into the hierarchical structure of the model.

VI. CONCLUSIONS

This paper presents preliminary results of a Bayesian spatio-temporal model applied to the spatial characterisation of wind speeds at sites where no information about the wind resource is available. Aspects of the methodology are described and results are shown for wind speed estimates for a target site using one year of data from four neighbouring weather stations as the only input. The developed model performance was assessed by comparison with data available for the target site. The Bayesian framework presents great potential in this application area, as it provides great flexibility in the model development, allowing the inclusion of knowledge derived from the physical phenomena (use of additional variables) that may have influence over the variables of interest. As further work, the authors are now pursuing the integration of additional climatologic/orographic variables into the model to reduce the unstructured error component and improve overall performance.

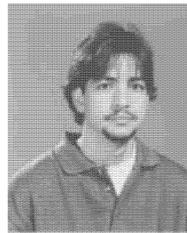
VII. ACKNOWLEDGMENT

The authors thank the UK Met Office and the British Atmospheric Data Centre for the data used in this research.

VIII. REFERENCES

- [1] OpenBUGS project website, <http://mathstat.helsinki.fi/openbugs>, April, 2006.
- [2] L. Landberg, "Short-term prediction of the power production from wind farms," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 80, pp. 207-220, 1999.
- [3] U. Focken, M. Lange, D. Heinemann, and H. P. Waldl, "Previento - Regional Wind Power Prediction with Risk Control," in *Proc. 2002 Global Windpower Conference*.
- [4] M. Milligan, M. Schwartz, and Y.-H. Wan, "Statistical Wind Power Forecasting Models: Results for U.S. Wind Farms," in *Proc. 2003 Windpower Conference*, Austin-TX.
- [5] WASP (Wind Atlas Analysis and Application Program) software website, <http://www.wasp.dk>, April, 2006.
- [6] K. V. Mardia, and C. R. Goodall, "Spatial-Temporal Analysis of Multivariate Environmental Monitoring Data," in *Multivariate Environmental Statistics*, vol. 6, G. P. Patil and C. R. Rao, Ed. New York/Amsterdam: Elsevier/North-Holland, 1993, pp. 347-385.
- [7] K. Mardia, C. Goodall, E. Redfern, and F. Alonso, "The kriged kalman filter," *Test*, vol. 7, pp. 217-276, 1998.
- [8] J. Haslett, and A. E. Raftery, "Space-time modelling with long-memory dependence: Assessing Ireland's wind power resource (with Discussion)," *Applied Statistics*, vol. 38, pp. 1-50, 1989.
- [9] G. Shaddick, and J. Wakefield, "Modelling daily multivariate pollutant data at multiple sites," *Applied Statistics*, vol. 51, pp.351-372, 2002.
- [10] T. Burton, D. Sharpe, N. Jenkins, and E. Bossanyi, *Wind Energy Handbook*. Chichester: Wiley, 2001, p. 642.
- [11] P. Congdon, *Bayesian Statistical Modelling*. Chichester: Wiley, 2001, p. 556.
- [12] Danish Wind Industry Association website, *Guided Tour on Wind Energy*, [Online]. Available: <http://www.windpower.org/en/tour/wres/weibull.htm>, April, 2006.
- [13] NIST/SEMATECH, *e-Handbook of Statistical Methods*. [Online]. Available: <http://www.itl.nist.gov/div898/handbook/pmc/section5/pmc5.2.htm>, April, 2006.

IX. BIOGRAPHIES



Marcos S Miranda (M'2006) was born in Belo Horizonte, Brazil in 1973. He graduated in Electrical Engineering in 1994 and obtained a Master's degree by Research in 1997, both from the Federal University of Minas Gerais (UFMG, Brazil).

He then worked as a Research Assistant at the Brazilian Wind Energy Centre. In 1998 he started his PhD at CREST (Centre for Renewable Energy Systems Technology) at Loughborough University (UK). Since the end of 2003 he has been with the Dept. of Electrical Engineering, University of Bath (UK) as a Research Associate. His areas of interest include renewable energy sources, power systems, dynamic modelling and control and energy efficiency.



Gavin Shaddick was born in Devon, England in 1970. He graduated from the University of Warwick in 1991 after studying Mathematics and Statistics and obtained an MSc in Statistics from University College, London. After working at the London School of Hygiene and Tropical Medicine and Imperial College, he started his PhD at Imperial in temporal and spatial statistical methods in epidemiological in 1997. Since January 2001, he has been a lecturer in statistics in the Department of Mathematical Sciences at the University of Bath. Gavin Shaddick's current funded research includes high resolution spatial modelling of pollutants in the EU and the health effects of air pollution episodes.



Rod W. Dunn (M'1996) was born in Glasgow, Scotland in the UK on 29th July 1959. He graduated from the University of Bath, UK, in Electrical and Electronic Engineering in 1981. He started his PhD on novel control algorithms and embedded computer to implement the control for magnetically levitated urban transport vehicles.

In 1985 he became a Research Assistant within the Dept. of Electrical and Electronic Engineering at the University of Bath. He joined the academic staff in early 1986. Currently Dr Dunn is a Senior Lecturer within the Dept. of Electronic and Electrical Engineering at the University of Bath. Rod Dunn's current funded research includes electrical energy storage, reliability of future power systems and power system operation and planning techniques.

Furong Li (M' 1997) was born in Shannxi, China. She received her B.Eng. in Electrical Engineering from Hohai University, China in 1990, and her Ph.D. in 1997. She took up a lectureship at the University of Bath in 1997. Her major research interest is in the area of economic operation and planning of power systems.

Keith R. W. Bell received his B.Eng. in electrical engineering and electronics and Ph.D. on "Artificial Intelligence and Uncertainty in Power System Operation" from the University of Bath in the U.K. in 1990 and 1995 respectively. After a spell with Ansaldo Trasporti in Naples, he worked as a Research Associate at UMIST before taking up a post with the National Grid Company. Currently he is a Senior Lecturer within the Department of Electronic and Electrical Engineering at the University of Strathclyde.