A NOVEL APPROACH TO WIND FORECASTING IN THE UNITED KINGDOM AND IRELAND

P.R.J. CAMPBELL and K. ADAMSON

Faculty of Engineering
University of Ulster
Belfast, Northern Ireland

Abstract: Electricity markets throughout the world are adapting to allow the integration of alternative energy sources, in particular Wind. Countries such as the USA have already reached penetration with wind energy being integrated into power grids. This development has lead to wind energy producers having to provide forecasts of expected yield within a given time period. The ability to accurately predict short-term wind speed is to become a critical requirement for wind energy operators in the UK and Ireland over the coming year, due to new energy trading regulations. As a result it is necessary to identify a method by which such forecasts can be accurately and efficiently produced. Such forecasts have traditionally been compared to the Persistence method of forecasting as a bench mark for accuracy. This paper examines the statistical approaches of ARIMA, Moving Averages and compares their performance against both Persistence and a novel Multi-Layered Perceptron which is trained using the Generalised Delta Rule, demonstrating that such an MLP implementation can demonstrate significantly improved accuracy over these more traditional statistical approaches.

Keywords**:** Short-term wind speed prediction; artificial neural networks; time-series; wind energy

1. INTRODUCTION

The past decade has become a very intensive period for forecasting research and development [Bailey *et al*, 1999*]*. However, it is also clearly demonstrated that this work is being carried out from the perspective of the large utility. Whilst this is acceptable in the Danish, Swedish and American markets, were the majority of wind parks are owned and operated by such organizations, it does not hold true for the UK and Irish marketplaces where the majority of wind parks are under the operational control of smaller organizations, who do not have sufficient expertise or budgetary ability to pursue such ambitious projects.

As a result it is unusual to find either active research or application of such forecasting techniques for direct application to the problem [Waldl and Giebel, 2000]. It is especially rare to find application of the large scale NWP systems (such as Prediktor, ewind or WAsP), due solely to the cost involved in the purchase and running of such systems [Giebel *et al*, 1999]. Consequently literature has show a significant gap in the existing technology in the UK and Ireland for a cost efficient, accurate and undemanding forecasting methodology.

However, as wind energy continues to be integrated into the power grid in the UK and Ireland, it is likely that a framework similar to that currently in operation in the USA will be adopted [Makarov *et al,* 2003]. As a result it is imperative that wind energy producers and utility companies have a suitably efficient and accurate method for producing forecasts in the range of 0 to 12 hours. This does not preclude the importance of longer term forecasts, in the region of 13 to 48 hours, however, experience has shown that time boundaries will eventually shrink to the 0-5 hour period as full integration and energy trading is implemented.

It has already been shown, that much of the current focus on forecasting and prediction, over whatever horizon focuses on NWP (or NPM) systems. The major disadvantage of these systems is their complexity and consequent cost, in terms of financial purchase, training (of the system) and integration. Indeed current attempts to introduce such systems to the UK have been less than successful for well documented reasons.

To meet this requirement of a low cost methodology which is easily and efficiently adaptable this study will present an Artificial Neural Network model. The proposed model's design will be presented as will descriptions of training, testing and validation. It will also be shown that the proposed model compares favourably with other statistical approaches for wind speed prediction; by means of comparison with several well regarded approaches.

It is important to note the boundaries within which the models' use is intended. It is not within the scope of this study to produce a metrological modelling tool to rival the large scale geographical models described in the previous chapter. The model discussed is intended for use in producing accurate forecasts for individual wind parks in the UK and Ireland, across a range of horizon windows.

It is also important to note that the model will forecast wind speeds rather than energy yield. Much

debate has centred on this issue, however, investigation into the performance of a Neural Network to produce accurate power estimations, failed to produce a suitably accurate model, when compared to the accuracy achievable through estimation of wind speed and subsequent computation of power yield based on the estimated figure.

2 NEURAL NETWORK DESIGN

The model presented in this study is a multi layered perception model (MLP), more commonly known as a feed-forward network, and has been trained using the backpropagation (or generalized delta rule) method. The foundations of the backpropagation method for learning in neural networks were laid by Rumelhart *et al* [Rumelhart *et al*, 1986]. The following sections describe the operational design on such a model.

An artificial neural network is constructed of a number of simple processing units, neurons, which are grouped in layers. Each layer is identified by the index l = 0…..L. Layers 0 and L correspond to the Input and Output layers, with layers in between which are called Hidden Layers. Communication within the network is based on a layer structure with each neuron interacting only with neurons of neighbouring layers and not with neurons within the same layer. Each neuron has an activation value, a, associated with it and the network processes data by the exchange of activation levels between interconnected neurons. Figure 1 presents an illustration to aid the explanation of network design.

The output value of the i-th neuron within layer l is signified by x_i^l.. It is calculated using, $x_i^l = g(a_i^{(l)})$, where g() is the monotone increasing function.

The network input of neuron i in layer l is given by:

$$u_i^{(l)} = \left(\sum_{j=1}^{n^{(t)}} w_{ij}^{(l)} x_j^{(l-1)} \right) - \theta_i^{(l)} \qquad (1)$$

Where $w_{ij}^{(l)}$ is the weighting of neuron *j* in layer *l*-1connected to neuron *i* in layer *l*, $x_j^{(l-1)}$ is the output of neuron *j* in layer *l*. $\theta_i^{(l)}$ is the bias value which is subtracted from the summation of the activated weightings.

The calculation process of the model begins at the input layer and terminates at the output layer. Input *I* initialises the activation weightings of the neurons in the input layers. The activation level one layer is then passed to the next layer; it is at this point that the weightings (between the layers) are altered by the delta rule (backpropagation algorithm). The model learns (it learns the mappings between the input and output nodes) by an incremental process of changing the weightings between the layers, in order to minimise the variance between the actual and target (from the training data) values.

2.1 The learning process

The learning process can be thought of as two sub processes. First, a set of data are fed into the input layer. This activation is then propagated through the network model to the output layer (via the hidden layers). The second stage actually performs the process which can be considered learning. The output is compared to the target to assess the performance of the model. An error value is assigned to each neuron in the output layer, these error values are then propagated back through the neurons to the hidden layers. The corresponding weightings are then altered in order to reduce the error for the next pass of the same pattern. It is the Generalised Delta Rule which is used to perform this backpropagation process and reduce the weightings [Njau, 1994].

2.2 Network model design

To predict wind speed a 4 layer model (shown in Figure 1) was selected, which consisted of an Input layer 2 hidden layers and an output layer. A number of different network architectures were tested including 1, 2, 5 and 7 layer models. A common engineering approach previously employed by Kalogirou *et al* [Kalogirou *et al*, 1999] (consisting of 5 layers) for wind speed prediction was used as a benchmark however, the 4 layer design proved to be the most efficient and accurate, demonstrating a 17% in accuracy when compared to Kalogirou's approach and 24% increase when compared to a 2 layer architecture.

As with the overall network design a range of activation functions were tested, these included; Sigmoid, Gaussian, Hyperbolic Tangent and Hyperbolic Secant. After analysis it was discovered that the network operated most efficiently using the activation functions shown in Table .1.

Layer	Layer Type	No. of Nodes	Activation Function
1	Input	3	Linear
2	Hidden	2	Sigmoid
3	Hidden	2	Gaussian
4	Output	1	Sigmoid

Table .1. Model Activation Functions

Activation Functions

The sigmoid function acts as an output gate that can be opened (1) or closed (0). Since the function is continuous, it also possible for the gate to be partially opened (i.e. somewhere between 0 and 1). Models incorporating sigmoid transfer functions often help generalized learning characteristics and yield models with improved accuracy however the use of sigmoid transfer functions can also lead to longer training times.

The gaussian transfer function significantly alters the learning dynamics of a neural network model. While the sigmoid function acts as a gate (opened, closed or somewhere in-between) for a neuron’s output response, the gaussian function acts like a probabilistic output controller. Like the sigmoid function, the output response is normalized between 0 and 1, but the gaussian transfer function is more likely to produce the “in-between state”. It would be far less likely, for example, for the neuron’s output gate to be fully opened (i.e. an output of 1). Given a set of inputs to a neuron, the output will normally be some type of partial response. That is the output gate will open partially. Gaussian based networks tend to learn quicker than sigmoid counterparts, but can be prone to memorization.

Model Data

The input data, comprised wind data values collected at each test WTG unit. The initial input consisted of the hourly average wind speed and the time/date of the records. This allowed the Network to learn the patterns of wind speeds at different times of the day, week, month and year. The output of the network is the wind speed for a target WTG unit based on the time frame in question.
As the model makes use of the backpropagation technique it is necessary to have all training targets normalised between 0 and 1, as the output neuron is restricted to a signal of only 0 or 1 values. As a result all inputs have also been normalised to the 0 to 1 range. It is accepted that by conducting this normalising process the characteristics of the training process are enhanced [Mohandes et al, 1998].

A random inclusion method was also selected for the training process. In addition the normalisation process attempted to spread the values throughout the range (0-1), as it is possible that all normalised values could fall between a range of 0.1 - 0.2. Obviously spreading these values across a range will present a more accurate output.

The learning rate for the network was set at 0.01, obviously this selection slowed the speed of learning by the network, however, when tested with a higher learn rate (closer to 1) the network began to experience instabilities. The momentum was also set at a generally accepted standard of 0.8, changes in the momentum failed to yield any significant positive improvements in network performance [Campbell and Adamson, 2003].

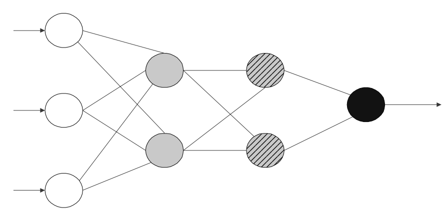

Figure 1 – 4-Layer Network

2.3 Training of the model

Selection of the training data for the model was a critical concern. The training data should present a typical dataset from which the model could learn, and should also provide a sufficiently extensive repository in order to include an array of wind conditions.

A 25 WTG (Wind Turbine Generator) park in County Donegal, Ireland was selected as the test site for the modelling process. This site was selected as it is one of the longest operational sites in Ireland and so provides a large dataset, containing a varied range of wind conditions. The dataset is collected centrally at the site and holds information on every aspect of site operation including performance and mechanical operation. This study focuses solely on the wind speed data collected by each WTG unit on a 10-minute interval.

The study focuses on predicting the wind speed rather than power output. From the dataset a subsection of data was selected, covering years 1996 to 2001. Data between the years 1996-2000 was used as the training and testing set for the neural network. While the data for 2001 was used as the verification set for the validation of the network model. Data is collected at ten minute intervals by each individual WTG, however, the model would eventually focus on a range of time windows

(Horizons). As a result the model was first trained to produce forecasts for an hourly average and then assessed for its ability to produce forecasts ranging from 10-minute intervals to monthly averages.

Tables 2 and 3 show examples of the training data used to train the model. As can be seen a selection of WTG units (WTG01, WTG10) were selected as the input values, in addition to a period input (month of year or 10 minute time slot) while WTG20 was used as the target for the modelling exercise. The use of units which are dispersed around the park (in terms of physical location) ensure that a park wide prediction is achieved.

Year	**Month**	**WTG1**	**WTG10**	**WTG20**
1996	1	6.9	7.25	6.83
1996	2	9.1	9.66	9.51
1996	3	6.98	7.5	7.13
1996	4	7.9	8.05	7.62
1996	5	7.22	7.22	7.22
1996	6	6.93	7.17	6.82
1996	7	5.48	5.48	5.48
1996	8	5.34	5.42	5.12
1996	9	6.71	6.83	6.51
1996	10	8.45	8.92	8.53
1996	11	7.65	8.55	7.76
1996	12	11.42	11.93	11.71
1997	1	7.15	7.27	6.94
...	...	...	...	...

Table 2 – A Sample of training data used

2.4 Network testing

Initially to test the network both hourly average and monthly average forecasts were predicted. The section below gives an example of the performance based on a monthly average forecast. Shorter term forecasts will be discussed in subsequent sections of this paper.

Monthly Predictions

Nine hundred randomly selected wind speeds were used for training and 300 wind speeds were used for validation purposes. These are presented in Table 4, and shown graphically in Figure 2. As can be seen the model produced a high level of accuracy in the prediction of monthly average wind speeds for the 12 months of 2001. The root mean square error (RMSE error) for the validation set was 0.11078. The network yielded a correlation of 0.983, and further demonstrated its ability by yielding a maximum percentage error within the validation set of 8%, a maximum error of 0.79.

TIMESLOT	**WTG1**	**WTG10**	**WTG20**
1	2.3	2.85	2.37
2	3.5	3.47	3.35
3	4.1	4.26	4.12
4	4.4	4.71	4.32
5	4.3	4.64	4.55
6	4.1	4.15	4.27
7	3.9	4.08	4.11
8	3.6	3.94	4.05
9	2.1	3.17	3.24
10	1.8	2.29	2.43
11	2	2.07	2.14
12	2.4	2.22	2.37
13	4.3	3.37	4.09

Table 3 – Sample of 10 min Training Data

This level of accuracy was continued in subsequent tests based at estimating the monthly speeds for 2002 and 2003. The results shown in Table 4 corresponded to error in real terms of 0.7 m/s, a value generally accepted as insignificant in terms of wind speed and its potential for energy generation, if we remember that energy yield is based on the cube of wind speed.

As the output shows, the ANN model shows significant ability and accuracy at producing prediction for the monthly average wind speeds. The application of the model on other UK and Irish based sites yielded similar results for monthly predictions. As a result the model was then selected for application to shorter term prediction models with the aim of reproducing the accuracy exhibited for horizons in the terms of minutes and hours.

Month	1	2	3	4	5	6	7	8	9	10	11	12
Actual	7.11	10.01	7.61	7.83	7.22	7.21	5.48	5.48	6.98	9.21	8.18	11.8
ANN	6.94	10.8	7.8	8.22	7.18	7.06	5.22	5.19	6.68	9.83	8.9	11.6

Table 4 – Model Output & Actual for 2001 test data

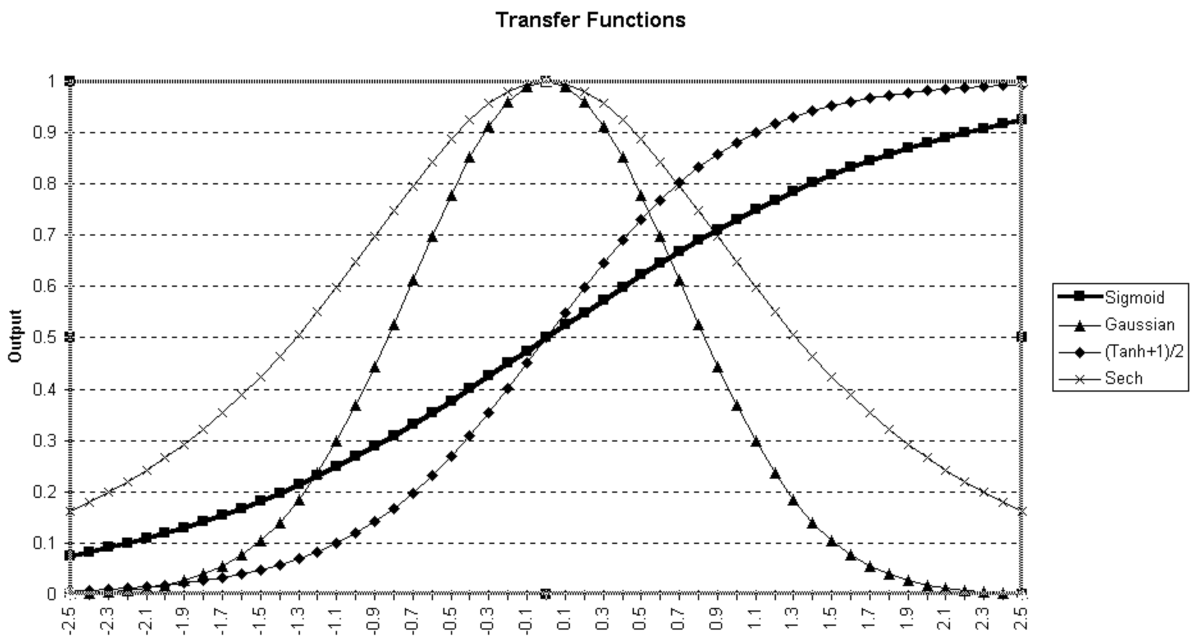

Figure 2 - Transfer (Activation) Functions

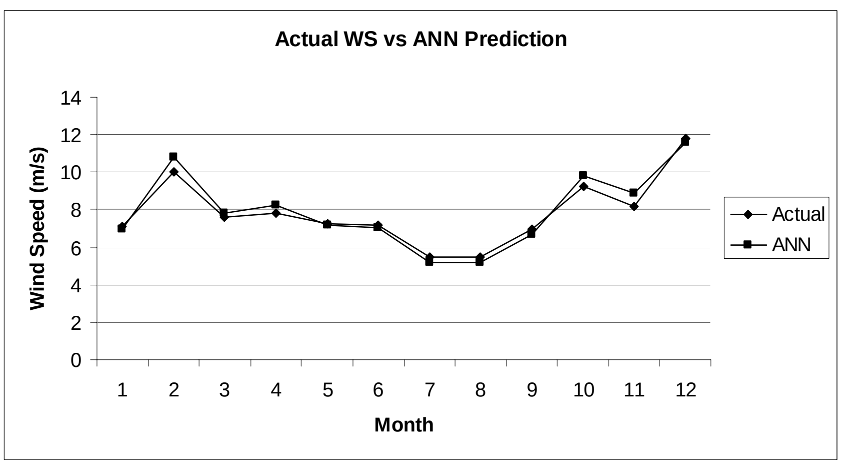

Figure 3 – Predicted monthly wind speeds compared to Actual for 2001

	Actual	Persistence	ARIMA 1	ARIMA 2	MA	ANN
Slot 1	2.3	2.3	2.3	2.3	2.3	2.45
Slot 2	3.5	2.3	2.32	2.89	3.5	3.37
Slot 3	4.1	2.3	3.66	4.17	3.3	3.97
Slot 4	4.4	2.3	3.65	3.91	4	4.35
Slot 5	4.3	2.3	4.18	4.15	4.26	4.31
Slot 6	4.1	2.3	3.82	3.67	4.26	4.27
RMSE	**-**	**1.64**	**0.61**	**0.37**	**0.37**	**0.12**

Table 5 – Results for One Hour Time Horizon in 10-minute steps

3 FORECAST EXPERIMENTATION

Having identified Neural Networks as a viable forecasting approach for prediction of monthly average wind speeds it is necessary to assess the models applicability to shorter horizons. Available literature shows that a number of statistical methods have been applied to the problem and so it was deemed appropriate to uses a selection of these methodologies for comparative purposes when attempting to assess the validity of the NN model.

4 METHODOLOGIES

This study examines four approaches to short term wind forecasting; Persistence, ARIMA, Moving Averages and Artificial Neural Networks. Each model presented is tested over a number of time periods (Forecast Horizons) and its performance compared to that of the Persistence approach. Each forecast has a RMSE (Root Mean Square Error) calculated and this figure is used to assess the accuracy of the model, a RMSE lower than persistence denotes a higher degree of accuracy.

4.1 Persistence

The most basic method for the short term forecasting of wind is the Persistence method. Expressed in its most simple terms this approach states that; *"the persistence forecast assumes that the conditions that existed at the beginning of the forecast period will continue or persist through to the end of the period;"* [Box and Jenkins, 1976]. This method is not only the simplest modelling approach but is also the most economical to implement, and surprisingly accurate for short term forecasting, of 1 to 5 hours. If we consider the time periods involved in short term forecasting we can see that generally the rate of change in contributing factors will normally not be sufficiently rapid to skew the Persistence model to any great degree. As a result the performance of any short term methodology should be measured by the extent it can improve on the persistence derived forecasts.

4.2 Box-Jenkins (ARIMA) approach

The ARIMA, Autoregressive Integrated Moving Average model developed by Box and Jenkins [Sfetsos, 2000], has a total of 3 components; autoregressive, integrated and moving average. In previous studies [Lin and Lee, 1996] the second component has been ignored, however, as will be seen in later sections, the inclusion of the integration component allows a more accurate model to be developed. An ARIMA(p,d,q), of order p and q, is represented as:

$$y_t = \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=1}^{q} \phi_j e_{t-j} + \varepsilon_t \qquad (2)$$

where ϕ_i are the autoregressive and ϕ_j the moving average components and ε is a random term. If the time series is not stationary in nature it can be differenced d times until achieved.

Two ARIMA models are used in this investigation, one as described by Sfetsos, which excludes the differencing component and employs an ARIMA(2,0,2) approach. The second model makes use of the differencing component and after testing a range of possible p,d,q combinations, employs an architecture of ARIMA(2,1,2).

4.3 Moving average

The moving average technique is one of the simplest and widely used forecasting methods, originally and still employed in stock market analysis. The Moving Average projects values in the forecast period, based on the average value of the variable over a specific number of preceding periods. A moving average provides trend information that a simple average (i.e. Mean) of all historical data would mask. The forecast value is based on the following formula:

$$F_{(t+1)} = \frac{1}{N} \sum_{j=1}^{N} A_{t-j+1} \qquad (3)$$

where N is the number of past periods to include in the Moving Average, A_j is the actual value at time j and F is the forecast value. Following individual testing of various MA models an N value of 3 was selected for the study.

4.4 Artificial neural network approach

The design and implementation of the Artificial Neural Network is discussed in section 4.2. The structure on the network is unchanged for any of the time horizons, however the period input is changed from months (as described in section 4.2.3 to 10-minute time slot identifier for short term forecasts).

However the network was retrained to produce predictions based on the required time horizons. The original features which indicated the selection of an MLP for this study still hold true, namely that an MLP is more accurate and robust than the simpler ANN approaches such as Linear Networks [Ranaweera *et al*, 1995], and also more economical to implement than the more complex approaches such as, Radial Basis Function networks [Campbell and Adamson, 2003].

4.5 Forecast horizons

As has previously been identified, there has been a dual need for forecasting, firstly to satisfy the integration and trading of power generated by wind sites, and secondly for the maintenance and operation of wind sites. Obviously each of these requirements will have differing forecast horizons. A trading forecast will be estimating in the 5-12 hour range, whilst for operation and maintenance the horizon would typically be in the region of 0-2 hours, such a need should not be confused with major overhaul of units which would require forecasts in the region of weeks or months.

To satisfy both of these conditions the methodologies outlined in the previous section will be compared based on their accuracy over a range of horizons, namely:

- Next hour in 10 minute intervals
- Next hour average wind speed
- Next 2 hours average wind speed
- Next 5 hours average wind speed
- Next 12 hours average wind speed

4.6 Forecast results

One hour horizon – 10-minute steps

Table 4.5 shows the results for each of the forecasting approaches employed based on a 10-minute time horizon. The Persistence approach uses the last value in the historical dataset as does ARIMA1, ARIMA2 and the Moving Average (MA) methodology. The ANN was trained using the original dataset as described earlier in this chapter. The MA approach is based on an N value of 3, while ARIMA 1 has a p,d,q of 2,1,2 and ARIMA2 has the corresponding values 2,0,2. The table shows the accuracy of each approach compared to the actual, for 1 hour ahead in 10 minute intervals. Such as forecast is useful to engineers for manual control of WTG units. Each time slot represents a 10-minute period.

As can clearly be seen each methodology with the exception of the Persistence model shows a high degree of ability in providing an accurate estimate of the wind spectra for the next 10 minute period. However the Multi Layered Perceptron produces the lowest error with a RMSE of 0.12.

One hour horizon

The table below contains the results for the application of the methodologies to the new time horizon of 1 hour. Once again Table 4.6 shows that the Persistence model dramatically under estimates the actual wind speeds encountered. This under estimation would have a much more import role to play when the actual power contained in the wind is calculated, if we consider that the power output from the park is calculated using the cube of the wind speed [Campbell and Adamson, 2004] we can see that such an under estimation can have significant consequences.

It is important to notice that the Multi Layered Perceptron approach shows an extremely high degree of accuracy, recording an error of only 0.003. It should however be noted that the model was originally developed to produce 1 hour average forecasts. As a result the accuracy for this horizon had been expected to be high. It is also interesting to note that based on this forecast horizon, ARIMA 1 is out performing ARIMA 2, conforming with and validating previous studies (discussed in Chapter 3) which stated that the differencing component had no effect on the model’s ability to forecast wind speed. However, it is also most interesting to see that the oldest and perhaps most simplistic method for forecasting, the Moving Average Approach, is the second most accurate method for prediction at this horizon level.

	Actual	**Persistence**	**ARIMA 1**	**ARIMA 2**	**MA**	**ANN**
1 Hr Ave	3.78	2.3	3.21	3.51	3.60	3.76
RMSE	**-**	**1.48**	**0.46**	**0.26**	**0.18**	**0.003**

Table 6: Results of One Hour Horizon Forecast

Two hour horizon

At this level of horizon the interested in forecast accuracy moves away from operational and mechanical requirements and becomes crucial for future integration and energy trading. As previously stated it is unlikely that the UK market will operate on such small time horizons initially, however, it is important for completeness to assess each methods ability at this level, should the time trading horizon be reduced at some future point.

Table 7 shows that the Persistence model continues to perform poorly, when compared the other methodologies. This fact is important as Persistence is still considered as the main benchmark for wind speed prediction up to 4 hours ahead. However, it is interesting to note that the Persistence error has reduced significantly for this longer forecast period. It is also interesting to note that ARIMA1 has outperformed ARIMA2 for the first time in these tests, suggesting that as the forecast window increases the differencing term incorporated in ARIMA1 (and excluded from ARIMA2) becomes a more important component. The Moving Average approach continues to show significant modelling power when applied to the problem of wind speed modelling. As can be seen it now outperforms the ANN model and is the most accurate model at forecasting wind speed.

Five hour horizon

According to literature it is at the point, attempting to forecast 5 hours ahead, that the methods which this study is employing (statistical models) should begin to show signs of failure. As we can see from the results for this 5 hour horizon, the accuracy of each of the methods has reduced. However, we can also clearly see that both the Moving Average and Neural Network approach still show low levels of error, 0.25 and 0.29.

The Persistence approach is once again the least accurate model with a RMSE of 1.43, an increase of error from the 2 hour forecast, supporting previous studies findings, which state that the accuracy of persistence is limited to less than 4 hours of a horizon. Once again the Moving Average approach demonstrates that it is the most accurate, showing the same level of accuracy at producing a 2 or 5 hour forecast.

It is interesting to note that the Moving Average approach has recorded the same level of accuracy for the 5 hour horizon as it did for the 2 hour horizon. We should also notice the performance of the ARIMA2 methodology which continues to show a higher degree of accuracy when compared to the ARIMA1 model, reinforcing the conclusion that the differencing term is critical for longer range horizons.

The 12 hour horizon

The 12 Hour forecast is the longest horizon forecast included in this study, and is well outside the acknowledged boundaries for statistical estimation of wind speed, as described in literature. As can be seen from Table 4.7 the Persistence model has continued its poor performance and unsurprisingly records its highest error, an RMSE of over 2.3.

Both ARIMA models have decreased significantly in accuracy at this horizon level, although the inclusion of the differencing term (in ARIMA1) ensures that the forecasting error stays with currently acceptable levels, at least in terms of commercial application under (current) UK guidelines.

The Moving Average and the Neural Network approaches continue to yield the most satisfactory results. However, both methods show signs of reduced accuracy at this range, although the reduction in accuracy is very small, especially for the Moving Average approach. Both models yield an error value of less than 0.50.

	Persistence	**ARIMA 1**	**ARIMA 2**	**MA**	**ANN**
2 Hour RMSE	1.07	0.37	0.44	0.25	0.30
5 Hour RMSE	1.43	0.28	0.35	0.25	0.29
12 Hour RMSE	2.31	0.92	1.37	0.28	0.39

Table 7: Results for 2,5 and 12 Hour Forecast Tests

5 CONCLUSIONS

This investigation set out to examine the intelligent prediction of wind spectra, and has presented a Multi Layered Perceptron which is capable of forecasting wind speeds for wind park sites in the UK and Ireland, for a variety of time horizons. It has also been clearly demonstrated that the current practice of implementing the Persistence approach as a low cost prediction tool is fool hardy and flawed. The results presented clearly show that of all possible low cost statistical approaches Persistence consistently performs least well and has in examined cases grossly underestimated the potential wind spectra.

While the results achieved with Neural Network technology would seem to contradict previous studies, it is be clear from examination of literature that each of the other modelling approaches performed as had previously been documented. Clearly showing that whilst, Persistence and ARIMA models struggle to adequately model wind spectra, Neural Network technology can be applied to the problem.

There are three possible reasons for the success of the Neural Network approach in this study; firstly, perhaps models developed in previous studies have not been sufficiently refined before being applied to the problem. Or secondly, that the wind energy spectra in the UK and Ireland are particularly suitable for modelling by Neural Networks. Perhaps the third possibility is the most satisfying explanation, that both of the possible reasons identified above have a part to play. That is to say that the UK and Irish spectra is more suitable to the NN approach and that the model presented has been rigorously researched, designed, implemented and trained [Fayyad, 1996].

Overall the study has produced results which validate and also extend current knowledge in the field of short term wind speed forecasting. The results demonstrate that both Moving Averages and Neural Networks can operate well within acceptable boundaries for horizons of up to 12 hours, based on UK/Irish data. Indeed both approaches have recorded RMSE values lower than some NWP systems, showing that for horizons of 12 hours or less, both of these statistical methods can not only compete but also outperform such large and expensive modelling systems. Section 2.4 also demonstrates that the Neural Network model is capable of producing highly accurate forecasts for horizons of up to 1 year, based on monthly averages, a horizon which was previously regarded as far outside of the approaches prediction ability.

The results of both the short term and longer term tests clearly demonstrate that Neural Networks can be applied to wind speed forecasting not just as a cost effective indicator, but as a viable alternative to larger, more complicated and more expensive (to implement) NWP systems, not only in the UK and Ireland but also world wide providing a affordable and accurate prediction tool.

If we consider the long term industry aim of developing a Decision Support System for Wind Energy, then we can clearly see that the model proposed in this paper has a significant role to play. Not only will it aid integration in the short-term but it can potentially form a core module of any DSS, by providing forecasts for resource planning and maintenance amongst other areas. As stated, a decision support system aims to aid management decision making. In order to perform this function it is necessary to have the maximum amount of knowledge available in the area. Consequently while the prediction model discussed enables management to gain better understanding of potential future conditions, it is equally important that they are equipped to make decisions based on past data and events.

Significant scope still exists for the refinement of ANN approaches, which could potentially involve multiple models each producing forecasts for a different time horizon. The varying types of ANN implementations should also be more fully investigated and tested to assess their applicability to the role of wind forecasting.

REFERENCES

Bailey, B., Brower, M.C. and Zack, J., 1999 ,"Short-Term Wind Forecasting", *In Proceedings of the European Wind Energy Conference*, pp. 1062-1065.

Giebel, G., *et al*, 1999, "Relative Performance of different Numerical Weather Prediction Models for Short Term Prediction of Wind Energy", *In Proceedings of the European Wind Energy Conference*, pp. 1078-1081.

Waldl, H.-P., Giebel, G., 2000, "The Quality of a 48-Hours Wind Power Forecast Using the German and Danish Weather Prediction Model", *Wind Power for the 21st Century, EUWEC Special Topic Conference*, Germany, 2000

Makarov, Y., *et al*, 2003 "California ISO Wind Generation Forecasting Service Design and Experience", *Proceedings of American Wind Energy Association Conference*

Rumelhart, D.E., Hinton and G.E., Williams, R.J., 1986 “Learning representations by back-propagating errors*”*, *Nature*, 323(9) pp 533-536

Njau, E.C., 1994, “An electronic system for predicting air temperature and wind speed patterns”, *Renewable Energy*, 4(7), pp 793-805.

Kalogirou, S., *et al. 1999,”*Wind Speed Prediction Using Artificial Neural Networks”, *In Proceedings of the European Symposium on Intelligent Techniques ESIT\rquote 99 on CD-ROM.*

Mohandes, A.M., Rehman, S. and Halawani, T.O., 1998, “A neural network approach for wind speed prediction”, *Renewable Energy*, Vol. 13, No. 3, pp. 345-354.

Campbell, P.R.J. and Adamson, K., 2003 “A Multi Layered Perceptron for Wind Speed Generation*", In Proceedings of the 12th International Conference on Applied Simulation & Modelling*.

Box J. and Jenkins G., 1976, “*Time series analysis, forecasting and control”*, Holden-Day.

Sfetsos, A., 2000, “A Comparison of various forecasting techniques applied to mean hourly wind speed time series”, *International Journal on Renewable Energy, 21*, pp. 23-35.

Lin, C. and Lee, C., 1996, “*Neural fuzzy systems”*, Prentice Hall.

Ranaweera, D., Hubele, N. and Papalexopoulos, A., 1995, “Application of radial basis function network model for short term load forecasting*”*, *IEE Proc Generation Transmission Distribution*, 142(1).

Campbell, P.R.J. and Adamson, K., 2003, “Estimation of Energy Yield from Wind Turbine Generators”, *In Proceedings of the International Conference on Power Systems, PowerCon.*

Campbell, P.R.J. and Adamson, K., 2004, "A Comparison of Forecasting Methodologies for Wind Speed Time Series", *In Proceedings of the 4th International Conference on Modelling, Simulation and Optimization*.

Fayyad, U., Piatetsky-Shapiro, G., Smyth, P. and Uthurusamy, R., (Editors), 1996, “*Advances in Knowledge Discovery and Data Mining*”, AAAI/MIT Press.

BIOGRPAHY

Piers Campbell is a Lecturer in the School of Computing and Information Engineering at the University of Ulster. He was awarded his PhD in 2004 having previously completed a BSc (Hons) in Computing Science and Master’s of Research, all from the University of Ulster. Piers’ research interests include; Artificial Intelligence, Neural Networks, Knowledge Engineering, Intelligent Systems, Simulation and Modelling, Time Series Analysis, Knowledge Based Systems, Renewable Energy, and SCADA Technology.

Kenny Adamson is Reader in the School of Computing and Mathematics at the University of Ulster. Kenny was sponsored by Short Brothers (now Bombardier) to undertake his primary degree in Electrical and Electronic Engineering at the Queen’s University of Belfast. Dr. Adamson has founded four interdisciplinary research groups and has published over one hundred reviewed research papers. He is a regular Programme Committee member and chairman for various International conferences and also a regular reviewer for several International journals.