An Auto-Associative Residual Processing and K-means Clustering Approach for Anemometer Health Assessment

David Siegel[1] and Jay Lee[1]

[1]*NSF I/UCRC for Intelligent Maintenance Systems (IMS), University of Cincinnati, Ohio, 45220, United States of America*

siegeldn@mail.uc.edu
jay.lee@uc.edu

ABSTRACT

This paper presents a health assessment methodology, as well as specific residual processing and figure of merit algorithms for anemometers in two different configurations. The methodology and algorithms are applied to data sets provided by the Prognostics and Health Management Society 2011 Data Challenge. The two configurations consist of the "paired" data set in which two anemometers are positioned at the same height, and the "shear" data set which includes an array of anemometers at different heights. Various wind speed statistics, wind direction, and ambient temperature information are provided, in which the objective is to classify the anemometer health status during a set of samples from a 5 day period. The proposed health assessment methodology consists of a set of data processing steps that include: data filtering and pre-processing, a residual or difference calculation, and a k-means clustering based figure of merit calculation. The residual processing for the paired data set was performed using a straightforward difference calculation, while the shear data set utilized an additional set of algorithm processing steps to calculate a weighted residual value for each anemometer. The residual processing algorithm for the shear data set used a set of auto-associative neural network models to learn the underlying correlation relationship between the anemometer sensors and to calculate a weighted residual value for each of the anemometer wind speed measurements. A figure of merit value based on the mean value of the smaller of the two clusters for the wind speed residual is used to determine the health status of each anemometer. Overall, the proposed methodology and algorithms show promise, in that the results from this approach resulted in the top score for the PHM 2011 Data Challenge Competition. Using different clustering algorithms or density estimation methods for the figure of merit calculation is being considered for future work.

1. INTRODUCTION

One of the fundamental requirements for data interpretation, model development, and system monitoring is the need to have properly working and calibrated sensory data (Venkatasubramanian, Rengaswamy, Yin, K., & Kavuri, 2003). Considering the importance of properly working sensors, there is considerable research in the area of sensor fault detection and diagnosis with a diverse set of applications ranging from automotive (Capriglione, Liguroi, Pianese, & Pietrosanto, 2003), aerospace (Patton, 1991), to nuclear power plants (Hines & Garvey, 2006). The wind energy in particular, is quite reliant on obtaining accurate sensor measurements of wind speed, since this ultimately is one of the inputs used to estimate the energy production for a given site (Petersen, Mortensen, Landberg, Hujstrup, & Frank, 1998). During feasibility studies of potential wind turbine sites, anemometers placed on meteorological towers are used to provide information on the long term wind speed characteristics. Historical wind speed data is one of the inputs provided to sophisticated meteorological models that provide an estimation of the energy production for a given site. Errors in the wind speed measurements can have significant effects on the estimated energy production which could affect the return on investment for a given site or whether the site is financed (Murakami, Mochida, & Kato, 2003).

Recent work in the area of anemometer fault detection includes the work by Kusiak, Zheng, and Zhang (2011), which propose a virtual sensor method using a multilayer perceptron neural network. This study also discusses the use of a wavelet de-noising method for data pre-processing and a control chart based on the residuals calculated from the predicted and measured wind speed. A more classical statistical approach was discussed in the work by Beltran, Llombart, & Guerrero (2009), in which a metric was derived from the difference in the 10 minute wind speed average data between two anemometers in close proximity. In addition to this prior work, a recent study by Clark, Clay, Goglia, Hoopes, Jacobs, and Smith (2009) was done to investigate the root cause of NRG #40 anemometers reading slower than the actual wind speed. The study discussed statistical metrics, signatures of anemometers with excess measurement error, calibration methods, and a physical explanation of the potential failure mode that was believed to be the cause of the sensor measurement error.

The paper is organized in the following manner: after the introduction, Section 2 describes the problem statement for the 2011 Prognostics and Health Management Society Data Challenge. This is followed by an overview of the algorithms used for the shear and paired data sets in Section 3. More detailed descriptions of the filtering and data normalization methods are described in Section 4. The residual processing method using auto-associative neural network models is presented in Section 5. Section 6 describes the use of k-means clustering and the figure of merit health value. Lastly, conclusions and future work are discussed in Section 7 and Section 8 respectively.

2. PROBLEM STATEMENT

The 2011 Prognostics and Health Management Society (PHM Society 2011) presented a data challenge problem dealing with this increasingly important topic of anemometer fault detection. Two different types of data sets titled the "paired data set" and the "shear data set" was provided for developing and evaluating anemometer fault detection algorithms. The paired data set consisted of data collected from two anemometers at the same height. Statistics from the two wind speed sensors, a wind direction measurement, and ambient temperature reading were provided. The statistics were calculated from a 10 minute time period and consisted of the mean, standard deviation, maximum, and minimum for each parameter. Data from paired anemometers in a nominal healthy condition were provided in 12 training data sets that comprised of 25 days worth of data. The competition also provided 420 test data files, in which each file contained 5 days worth of data. These test files were used to test the accuracy of the developed algorithm by the contest participants, in which the actual healthy state was unknown to the contest participants. In the test files, either one of the anemometers could be in a healthy or degraded state; the objective was to provide a correct healthy classification for both anemometers (PHM Society 2011).

The shear data set differed from the paired anemometer data set, in that there were either 3 or 4 anemometers and each anemometer was at a different height on the meteorological tower. Height information was provided for each anemometer and statistics from each anemometer were provided after processing the wind speed measurements in a 10 minute data block. As with the paired data set, the wind direction statistics and the ambient temperature statistics were also provided. In total, 28 or 23 parameters were provided in each data shear data file, the difference in the number of parameters is due to certain sites only having 3 anemometers instead of 4 anemometers. A total of 7 training data sets that comprised of 25 days worth of data were provided for the shear data set; the training data sets provided data from anemometers in a healthy condition.

Test files were also provided in which the health condition of the anemometers were unknown to the contest participants. The test files consisted of 225 files, with each file representing 5 days worth of data. The objective in the shear data set was to determine whether the set of anemometers were all in a nominal healthy state or one or more of anemometers had a fault and were exhibiting excessive measurement error. Unlike the paired data set, it was not required to determine which anemometer was experiencing a fault. The requirements for the shear anemometer data set were to detect whether the system was in either a healthy or abnormal health state (PHM Society 2011).

3. ANEMOMETER HEALTH ASSESSMENT METHODOLOGY

The overall approach for assessing the health state of an anemometer consists of a series of algorithmic processing steps. These steps include data pre-processing, a residual calculation, and ultimately a decision on the health status based on a figure of merit metric. The health assessment algorithms developed for the shear and paired data sets follow that step by step processing methodology; however considering the unique aspects of both data sets, there are specific differences with regards to data normalization and the residual calculation. Section 3.1 and section 3.2 presents an overview of the methodology for assessing the health condition of the shear anemometers and paired anemometers respectively. More specific details of each processing module along with intermediate results from each step are shown in the subsequent sections to further illustrate the anemometer health assessment method.

3.1. Algorithm for Shear Data Set

A flow chart of the health assessment algorithm for the anemometer shear data is provided in Figure 1. The algorithm used in this study has a training and monitoring phase, in which the training phase is developed using shear anemometer data from a nominal healthy state. The initial step in the training process is to perform data filtering. The data filtering step is designed to remove instances in which icing could occur as well as to remove other data samples in which there could be erroneous readings in wind speed, temperature, or other sensor measurements. The data normalization step is a specific step designed for the shear data and is based on the wind profile power law (Peterson, & Hennessey Jr., 1977). This normalization procedure uses the power law equation to place each of the shear anemometer wind speed measurements at a common reference height.

The data normalization step reduces the variation due to elevation; however an auto-associative neural network is used to further model the relationship and correlation structure between the anemometer wind speed statistics. In this study, multiple baseline data sets were available for model training and this provided the opportunity to have multiple auto-associative neural network models. The training of the auto-associative neural network models completes the training process and the algorithm can then be deployed in its monitoring phase.

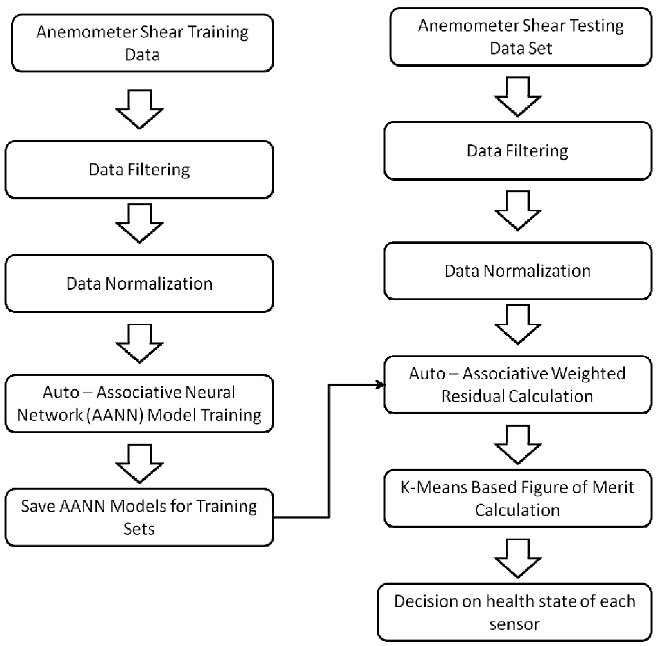

Figure 1. Algorithm Flow Chart for Shear Data Set

In the shear data set, a given monitored shear data set file consisted of either 3 or 4 anemometers and each data file comprised of 720 samples and a duration of 5 days. Thus, the data processing and health decision is performed on data from that 5 day period for a monitored shear anemometer set. The initial step for the monitored shear anemometers consist of performing the same data filtering and normalization that were used in the training set. A weighted residual calculation is performed using the auto-associative neural network models; a weighted approach is used in order to favor results from training models that more accurately predict the anemometer wind speed statistics. The residuals for the mean wind speed for each anemometer are then further processed in a k-means figure of merit calculation. The motivation for using a k-means clustering method is that prior literature suggested that the anemometers display a bimodal behavior in one of its failure modes and experience slowdown for a certain range of directions and wind speeds (Hale, Fusina, & Brower, 2011). Thus, the residuals might be quite small in a particular speed or direction regime and could be potentially quite higher in a different regime subset. A figure of merit calculation is performed for each anemometer, and a decision on the health status for each anemometer is made on whether the figure of merit value exceeds the threshold.

3.2. Algorithm for Paired Data Set

The health assessment algorithm used for the paired data set can be considered a subset of the one used for the shear data set; in that the algorithm used for the paired data set does not require the additional data normalization or auto-associative neural network based residual processing. The flow chart in Figure 2 shows the processing steps for the anemometer paired data, in which the initial step includes a data filtering step to remove data instances when icing takes place as well as other erroneous samples. Considering that the paired data set consists of anemometers at the same height, it is not necessary to use the wind profile power law for normalizing the data to a common reference height.

Although it is conceivable to train an auto-associative neural network for the paired anemometers; the initial rationale was that this would be too complex of a modeling approach for this situation. A direct comparison between the wind speed mean values provides a simple but an effective way of inferring the health state of the wind speed sensor. The approach used in this study calculates the difference between the wind speed mean values, denoted as d_{12} and d_{21}, and uses the difference signals as a surrogate for the residual signal used in the shear data set. The same k-means figure of merit calculation used in the shear health assessment algorithm is than applied to the difference signal.

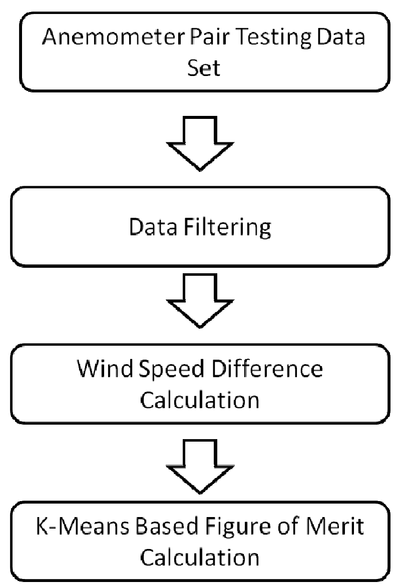

Figure 2. Algorithm Flow Chart for Paired Data Set

The primary focus of the health assessment algorithm for the paired data set was to provide a fault detection capability with the assumption that one of the anemometers was in a nominal healthy condition. However, during the course of a data file that represents a 5 day period, there potentially could be a cluster in which one anemometer was reading significantly lower than normal and another cluster where the other anemometer was reading significantly lower than normal. In this scenario, both figures of merit values could exceed the threshold and both anemometers would be reported as being in a degraded health state.

4. DATA PRE-PROCESSING

The subsequent processing steps in the health assessment algorithms for the shear and paired data set depend on quality data inputs. In both instances, data filtering is performed to remove erroneous data samples and provide a more suitable data set for further processing. An additional data normalization step is performed for the shear data set in order to place all wind speed measurements at a common reference height. Sections 4.1 and 4.2 provide the more specific details regarding the data pre-processing.

4.1. Data Filtering

The filtering routine is done to remove samples in which icing could be occurring and also for filtering out samples in which there are erroneous senor values. For removing instances in which icing is occurring, there is a variety of parameters that could be used to infer this condition; the wind speed direction standard deviation statistic in particular is quite useful for filtering out icing events. Considering that various statistics are calculated for each 10 minute data block, a value of zero in the wind speed direction standard deviation would imply that there is no variation in the wind speed direction for a 10 minute time period. Physically this is not possible and this condition of no variation in the wind speed direction is one of the key parameters that can be used for filtering out samples in which icing could occur.

The filtering settings used for the shear and paired data sets are provided in Table 1 and Table 2 respectively. For a given sample for the paired data set, it would have to satisfy all the listed ranges shown for the wind speed means, ambient temperature, wind direction mean, and wind direction standard deviation. It was observed in both the paired and shear training data sets that instances in which the wind direction were quite low resulted in more sudden changes in wind speed mean values. This resulted in larger differences between anemometer wind speed readings during these more abrupt changes. Considering this aspect, the filtering routine includes logic for the wind direction mean parameter to remove these samples in which the wind direction is below 50 degrees. It was also observed in the training data sets that the initial samples in each data file contained erroneous sensor values, thus the filtering routine also removed the first 20 samples.

The filtering routine for the shear and paired data set is quite similar, the major differences include that the paired data set filtering routine includes the anemometer wind speed standard deviation parameter. A low wind speed standard deviation would imply very little variation in the wind speed mean for a 10 minute period, which could imply icing. However, it was noted that including the anemometer wind speed standard deviation for the shear data set filter removed too many samples in a few of the test files, thus this setting was only used for the paired data set filtering routine.

Parameter	Filter Settings
Anemometer Mean 1	0.5 m/s – 26 m/s
Anemometer Mean 2	0.5 m/s – 26 m/s
Anemometer Mean 3	0.5 m/s – 26 m/s
Ambient Temperature Mean	-40 °C - 120 °C
Wind Direction Mean	Greater than 50 degrees
Wind Direction standard deviation	Greater than 0 degrees

Table 1. Filtering Settings for Shear Data Set

Parameter	Filter Settings
Anemometer Mean 1	0.5 m/s – 26 m/s
Anemometer Mean 2	0.5 m/s – 26 m/s
Anemometer 1 Standard Deviation	Greater than 1 m/s
Anemometer 2 Standard Deviation	Greater than 1 m/s
Ambient Temperature Mean	-40 °C - 120 °C
Wind Direction Mean	Greater than 50 degrees
Wind Direction standard deviation	Greater than 0 degrees

Table 2. Filtering Settings for Paired Data Set

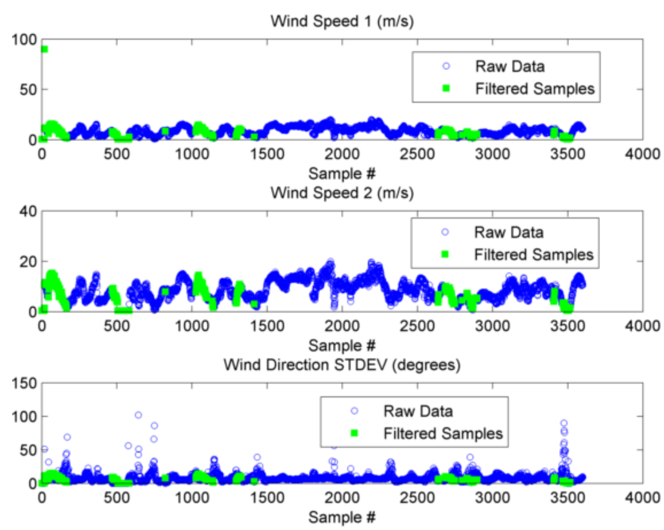

Figure 3. Raw and Filtered Out Samples – Training Set 6 for Shear Data Set

An example of how the filtering removes potentially icing events and erroneous data samples is shown in Figure 3. This example is from a shear training data set in which all 3 anemometer sensors are in a nominal healty condition; however, there is still a substantial amount of samples higlighted in green that are filtered out. The top plot highlights that the wind speed mean can have some extreme high or low values, as indicated by the outlier value near 100 and some of the values near or at zero. The middle plot shows the second anemometer wind speed mean reading and one can observe that there are several instances in which both anemometers are reading at or near zero. These near zero readings are likely due to icing. The wind direction standard deviation is shown in the bottom most plot and this parameter is also zero during these suspected icing samples. This example highlights that the filtering algorithm provides an adequate detection of icing and outlier samples.

4.2. Data Normalization

The data preprocessing for the shear data set includes an additional step of data normalization in order to compare the wind speed measurements at a common reference height. In prior work in the literature, the wind speed profile has been modeled as a logarithmic relationship and also by a power law model (Peterson et al, 1977). The use of the logarithmic equation includes an additional aerodynamic surface roughness parameter that depends on the site location; this was not provided in this study and thus only the power law equation was used for data normalization. The power law equation is described by Eq. (1), in which u_1 and z_1 are the wind speed and height at a known reference point and u_2 and z_2 are the wind speed and height at a location of interest. The exponent P is a constant that is based on prior experimental studies and regression fitting; a value of 1/7 is a common value for this constant and one that is used in this study (Hsu, Meindl, & Gilhousen, 1994).

$$\frac{u_2}{u_1} = \left(\frac{z_2}{z_1}\right)^P \tag{1}$$

For data normalization, each wind speed measurement is corrected to a height of 49 m. In Eq. (1), this would imply that z_2 is assigned a value of 49, while u_1 and z_1 are the known wind speed measurement and elevation for a given anemometer and u_2 is the corrected wind speed measurement at a height of 49 meters. An example of normalization process is illustrated in Figure 4 and Figure 5. Figure 4 is from the first shear training data set and is comparing the wind speed for anemometers 1 and 4. With regards to the numbering convention, anemometers 1-4 are sorted from the highest to the lowest height and in this example have a height of 59, 50, 30, and 10 meters respectively. As one can observe, there is significant differences in the raw wind speed values for anemometer 1 and 4, these two anemometers have the largest difference in elevation. Figure 5 shows the normalized wind speed mean values for anemometer 1 and 4 from the same training data set. From visual observation, one can observe that the differences in the normalized wind speed values are lower when compared with the raw data.

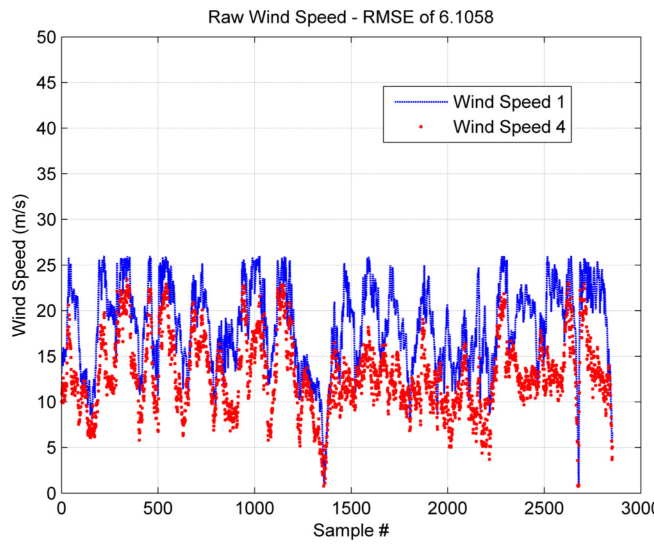

Figure 4. Shear Training Set 1 - Raw Wind Speed Mean Signals

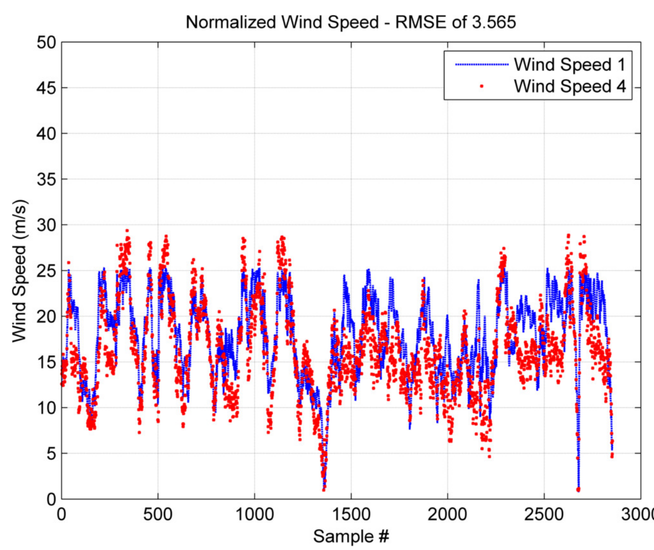

Figure 5. Shear Training Set 1 - Normalized Wind Speed Mean Signals

In order to quantify the differences in the wind speed measurements, the Root Mean Square Error (RMSE) is shown in each plot. The RMSE can be calculated between two anemometers by using Eq. (2), in which N is the number of samples in a data file and u_1 and u_2 are the wind speed mean values for the two anemometers considered in the calculation (Mohandes, Rehman, & Halawani, 1998).

The RMSE value for the normalized wind speed data in Figure 5 is much smaller than the RMSE value for the raw data in Figure 4; indicating that the normalization provided some measure of correcting for the different anemometer elevations.

$$\text{RMSE} = \left(\frac{1}{N} * \sum_{i=1}^{N} \left(u_1(i) - u_2(i) \right)^2 \right)^{1/2} \tag{2}$$

5. RESIDUAL BASED FEATURE EXTRACTION

5.1. Difference Signal

The residual based feature extraction for the paired data set does not involve any data normalization nor does it use auto-associative neural network models. The difference signals are used as a surrogate for the residuals and are defined by Eq. (3) and Eq. (4). They are simply the difference between the wind speed mean signals (u_1 and u_2) for the paired anemometers. The k-means based figure of merit calculation further processes these two difference signals to determine the health state of both anemometers.

$$d_{12} = u_1 - u_2 \tag{3}$$

$$d_{21} = u_2 - u_1 \tag{4}$$

An example of anemometers in a nominal healthy condition is illustrated in Figure 6, in which the wind speed mean values for each sample are matching very closely. The further processing of this data file by the difference signal and the figure of merit calculation resulted in this data file being classified in the healthy condition.

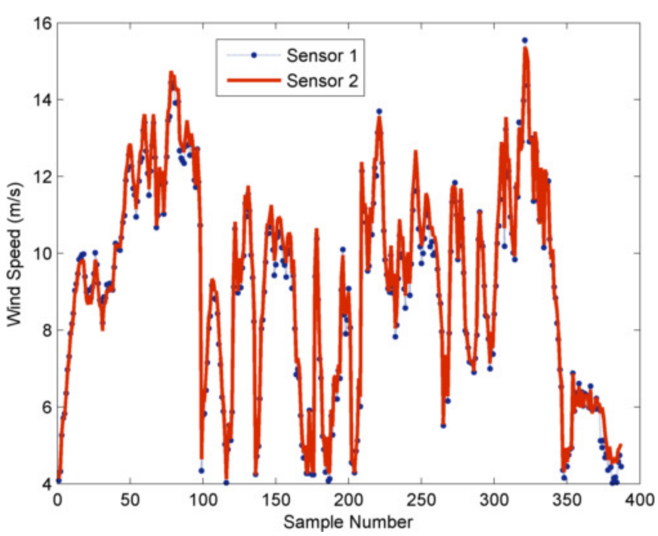

Figure 6. Wind Speed Mean Signals for Paired Test File 2 – Example of Anemometers in Nominal Healthy Condition

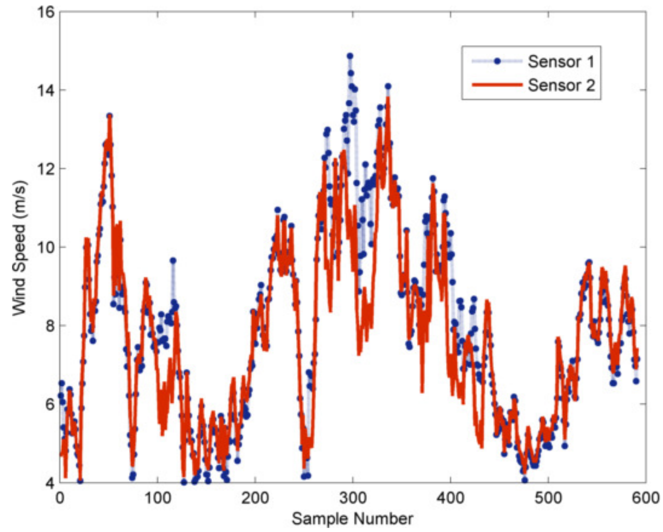

Figure 7. Wind Speed Mean Signals for Paired Test File 25 – Example of Second Anemometer Reading Slower

The signature that is exhibited when one of the paired anemometers is not working properly is highlighted in Figure 7. In this example, there are significant differences in the mean wind speed values for the two paired anemometers. However, these large differences are observed for only a portion of the samples. The observation that the signature only appears for a portion of the samples provides the motivation for clustering the difference signal and calculating a metric based on the cluster that contains information on the lagging sensor.

5.2. Auto-Associative Residual Processing

When a dynamic model of the system is not available a priori, the use of data driven health monitoring algorithms becomes a suitable alternative for monitoring the system health state (Schwabacher, 2005). Although there are various regression or distance from normal based metrics that are available, the use of auto-associative neural network (AANN) has some intriguing characteristics that make it particular suitable for this application. Its ability to learn non-linear correlation relationships and calculate residual values for each sensor provides a means to calculate a system health value. In addition, contribution plots for each sensor can also be used to provide diagnostic information (Thissen, Melssen, & Buydens, 2001). These attractive attributes of an auto-associative neural network have seen its usage for health monitoring span a diverse set of applications; from diesel engines (Antory, Kruger, Irwin, & McCullough, 2005) sensor health diagnostics and calibration (Xu, Hines, & Uhrig, 1999), to commercial aircraft engines (Hu, Qiu, & Iyer, 2007).

The theory and mathematics for the AANN were first described by Kramer (1991) and this method is effectively a way to perform non-linear principal component analysis.

Although principal component analysis (PCA) has been used in a variety of applications for process monitoring by using Hotellings' T^2 statistic and the residual square prediction error statistic (SPE); its assumption of the signals being linearly correlated is not satisfied in many engineering systems. An auto-associative neural network provides a similar framework, but has the ability to learn the non-linear correlation relationship among sensor variables. In this application, the underlying correlation relationship between the shear anemometer sensors is potentially non-linear; this is suggested by the power or logarithmic equations used to relate wind speed height and speed. The auto-associative neural network is applied after data normalization is done to correct for the wind speed height. However, the difference in the anemometer wind speed values in Figure 5 implies that the underlying relationship is not completely described by the power law. An auto-associative neural network can be used to further learn the sensor correlation relationship.

The AANN model structure consisted of 5 layers, an input layer, mapping layer, bottle neck layer, de-mapping layer, and an output layer as shown in Figure 8. One of its interesting aspects is that the network is trained with the same inputs and targets and thus the network is performing an identity mapping in which the output layer is providing an approximation of the inputs. The structure of the network used in this study follows the suggested configuration provided by Kramer (1991) and consists of 4 transfer functions. In sequential order, they consist of a tan-sigmoid transfer function, a linear transfer function, a tan-sigmoid transfer function, and a linear transfer function.

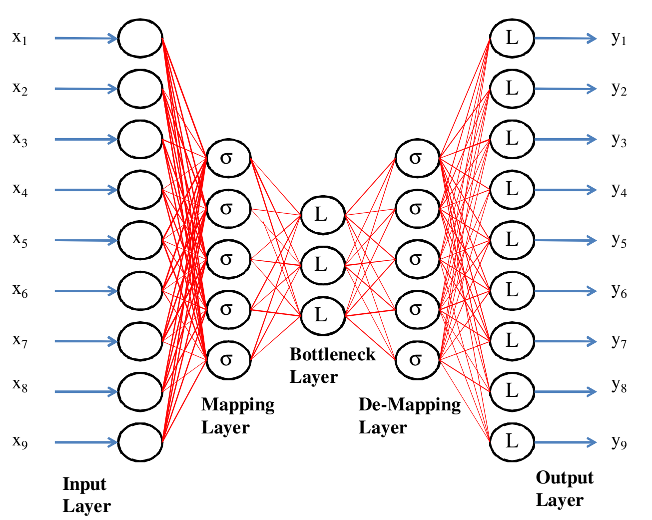

Figure 8. Auto-Associative Neural Network (9-5-3-5-9), σ is for tan-sigmoid transfer functions, L for linear transfer functions, x are inputs to the network and y are outputs of the network

Although the network structure uses the same transfer functions, there were some minor differences in the auto-associative neural network models for the 3 or 4 sensor shear anemometer configurations. The inputs for the network consisted of the wind speed mean, maximum, and minimum values for each anemometer; this provided 12 and 9 inputs for the 3 and 4 shear anemometer configurations respectively. The structure of the AANN model used in this study was configured so that the numbers of nodes in the mapping layer were the same as the number of nodes in the de-mapping layer. The number of mapping and de-mapping nodes consisted of 5 and 7 for the 3 and 4 anemometer configurations. In both configurations, the bottle neck layer consisted of 3 nodes. The number of bottleneck nodes represents the intrinsic dimension of the data in a similar sense to the number of principal components retained in linear PCA (Kramer, 1991).

As an additional extension of using the AANN models for anemometer health assessment, it was postulated that it might be advantageous to have an ensemble of training models. This could provide a way of giving more weight to training models that provide a more accurate sensor prediction for a given anemometer shear test data file. The rationale for considering this aspect is that there are several un-modeled sources of variation. Variation due to manufacturing, site topography, and installation, could potentially impact the AANN model accuracy.

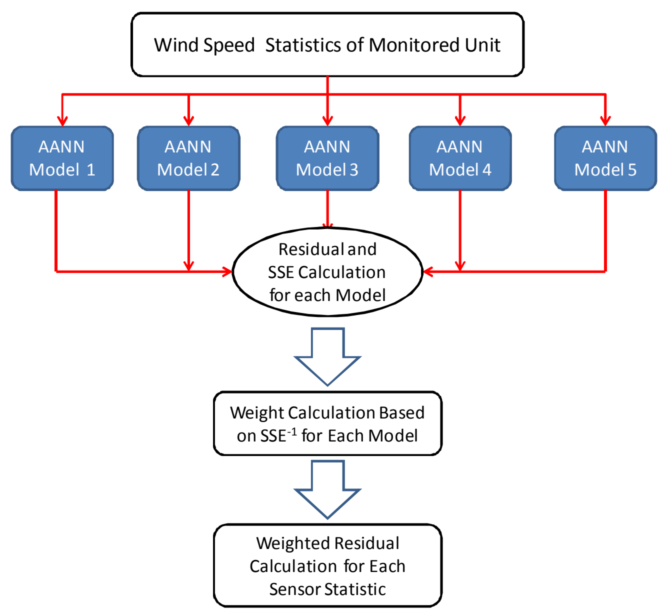

Figure 9. Flow Chart of Weighted Residual Calculation

Using a weighted approach allows one to weight training models that might more closely represent the test data set. This can reduce variances due to other factors and allows one to assume that the deviation from the model is due to anomalous anemometer sensor behavior. A conceptual diagram of the weighted residual approach is highlighted in

Figure 9 and the details of the calculation procedure are further described in this section.

The baseline files provided for the shear anemometer data set consisted of two training baseline files for the three anemometer configuration and five training baseline files for the 4 anemometer configuration. The weighted AANN residual approach consisted of having 7 trained AANN models for each of the baseline files, with 2 being assigned to the three anemometer configuration and 5 assigned to the 4 anemometer configuration. For a given test file, the residuals for each anemometer sensor statistic would be calculated for each model that matched the anemometer configuration for a given test file.

The residuals for each sensor statistic are weighted by a weight vector that is calculated from the sum of square error value as shown in Eq. (5) and Eq. (6). In this calculation, SSE_k is the sum of square error for k^{th} AANN model, and r_{ijk} is the residual based on the predicted AANN value and the actual sensor statistic value for the i^{th} data sample, the j^{th} sensor, and the k^{th} AANN model. In addition, N is the number of samples in the data file, and p and m is the number of input parameters and AANN models respectively. The weight for each model is calculated by taking the models SSE_k value and dividing that quantity by the summation of all the reciprocal SSE_k values. The weighted residual is then calculated by taking the weights for each model multiplied by the residuals as shown in Eq. (7). This provides a residual value for each sensor statistic that includes aspects from each training model, but provides more weight in training models that more closely match the test data set.

$$SSE_k = \sum_{j=1}^{p}\sum_{i=1}^{N}\left(r_{ijk}^{\ 2}\right) \quad \text{for } k = 1 \text{ to } m \tag{5}$$

$$w_k = \frac{(SSE_k)^{-1}}{\left(\sum_{k=1}^{m} SSE_k\right)^{-1}} \tag{6}$$

$$wr_{ij} = \sum_{k=1}^{m} r_{ijk} w_k \quad \text{for } i = 1 \text{ to } N \text{ and } j = 1 \text{ to } p \tag{7}$$

In order to evaluate the generalization of the AANN models and the weighted residual processing, the baseline data sets were randomly divided into a training set and a calibration set, in which the training set consisted of 70% of the available samples in a given baseline data file. An example of how well the predicted sensor statistic values match the actual values are shown in Figure 10 for the first shear baseline data file. In this example, the blue curve represents the weighted predicted value from the AANN models and the red samples are the actual wind speed mean values. A measure of the model fit can be assessed by the root mean square error value (RMSE). In this example, the RMSE value is significantly lower when the AANN models are used as opposed to the results in Figure 5 that were obtained with only data pre-processing and normalization.

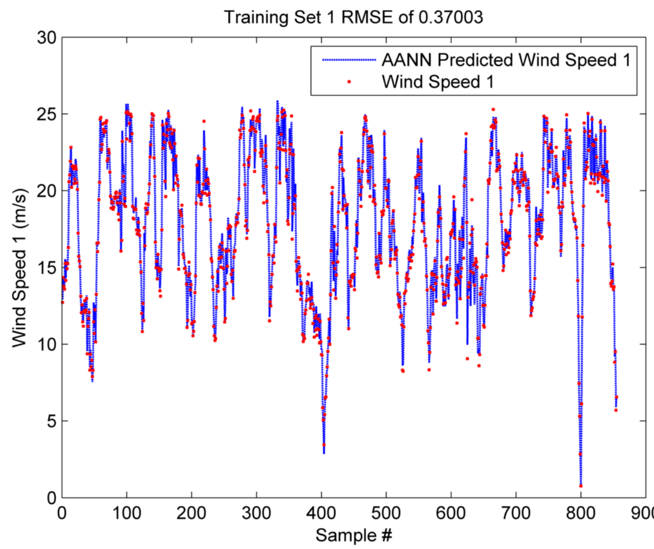

Figure 10. Shear Training Set - AANN Predicted and Measured Wind Speed Mean Values

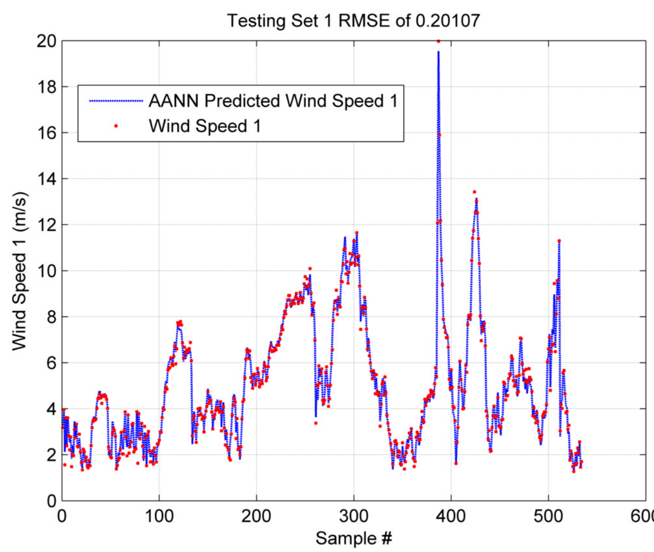

Figure 11. Wind Speed Predicted and Measured Value–Anemometers in Nominal Healthy Case (Shear Test File 1)

The trained AANN models and weighted residual processing method were then applied to the shear test files; example plots are shown in Figure 11 and Figure 12. In Figure 11, the results are for the first shear testing in which the predicted anemometer 1 wind speed mean and the actual anemometer 1 wind speed value are shown. Notice that the predicted and actual values match for the entire data set and the RMSE value is quite low. This is an example file in which the anemometers were considered to be in a healthy state.

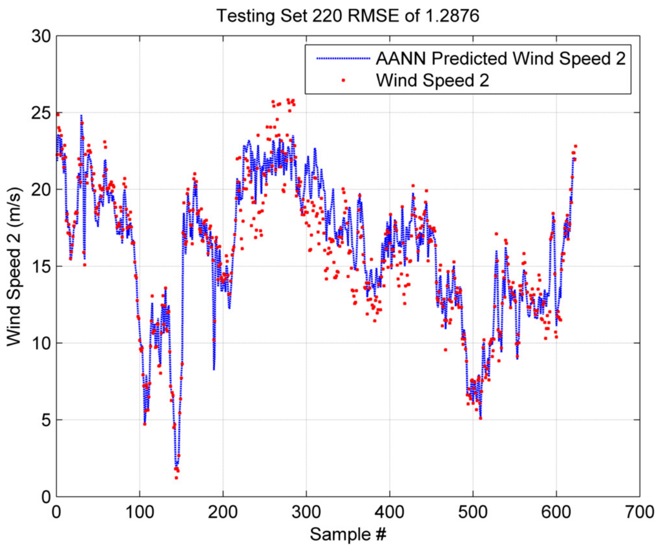

Figure 12. Wind Speed Predicted and Measured Value - Detected Faulty Case (Shear Test File 220)

An example of the anemometers in a degraded state is provided in Figure 12. In this example, there is a noticeable difference in the predicted and measured anemometer wind speed mean values for the second anemometer. The RMSE value for this case is also quite high. The bimodal fault signature is also observed, since the sensor is only lagging for a portion of the data samples. This highlights the motivation for clustering the residual signal, since the signature is only present for a particular subset of the operating conditions.

6. FIGURE OF MERIT

6.1. K-means Clustering

There are a variety of techniques used in data mining and artificial intelligence for clustering and density estimation (Jain, Murty, & Flynn, 1999). In this study, the k-means algorithm was used for partitioning the residual or difference wind speed mean values into two clusters. Density estimation using Gaussian mixture modeling was originally considered; however, the computation time became burdensome given the number of data files that had to be processed, and the k-means algorithm provided a more efficient way of determining the data clusters. The k-means clustering algorithm aims to partition the data set into a set of n clusters, where n is the number of clusters specified and its objective function is to minimize the within cluster sum of squares (Pollard, 1981). The algorithm is iterative in nature, in that it is initialized with a random set of centroids and through the iteration process updates the center locations in order to reduce the within cluster sum of squares distance. The interested reader is referred to the work by Hartigan and Wong (1979) for a more detailed description of the k-means algorithm.

Although the k-means clustering does not guarantee a global solution, 5 replications are used in this study in order to select the lowest local minimum that is obtained for the 5 replications. The mean value is calculated in each cluster and the minimum value of the two clusters is stored and denoted as the figure of merit value. There is an additional logic rule to prevent a small sample cluster from being included. If the sample size of one of the two clusters is below 60 samples, the mean of the other cluster is stored as the figure of merit value. A small cluster could be due to a small amount of outlier samples that potentially made it through the data filtering screening. The motivation for selecting the cluster with the minimum mean value is based on the prior literature that suggest that a degraded anemometer would be reading slower than normal (Clark et al, 2009).

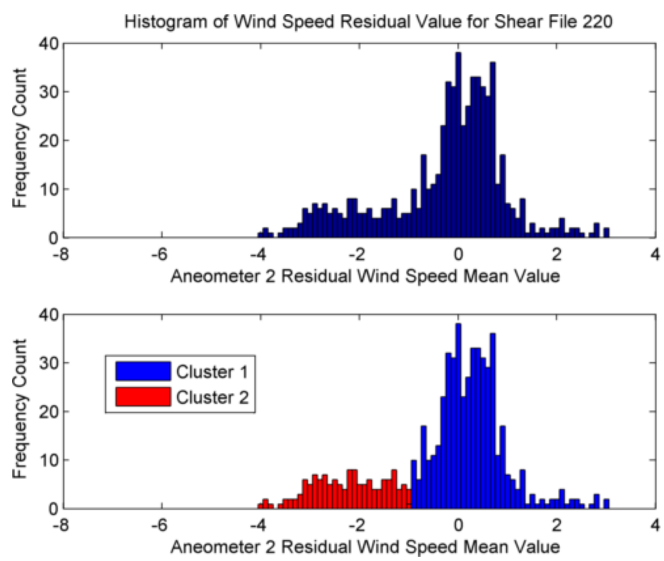

Figure 13. Wind Speed Residual Histogram and K-Means Clustering Result –File 220 Shear Data Set

An example of the k-means clustering result is provided in Figure 13. This result is from the residual wind speed signal for the system in a degraded health state. The histogram of the residual wind speed shows a bi-modal distribution in the top plot; the clustering result in the bottom graph indicates the two clusters that were determined using the k-means clustering. Considering that the figure of merit value is based on the mean value of the smaller of the two clusters, the k-means clustering provide a way of focusing on the samples when the anemometer is lagging. If one were to calculate statistics on the entire distribution without clustering, the algorithm would be less sensitive to the fault signature.

6.2. Figure of Merit Results

The previous section described how the figure of merit values were processed for both the shear and paired data sets; however the algorithm ultimately has to provide a decision statement on the health condition of each file. This required setting thresholds for the shear and paired figure of merit values. The literature suggests that an anemometer that is experiencing an increased level of friction and reading slower than normal could have an error of 1.5% to 3.0% and sometimes as high of a bias as 6% (Hale et al, 2011). The thresholds were based on selecting a value within that error range. The figure of merit thresholds for the shear anemometers were set at -0.35m/s for the three highest anemometers and a threshold of -0.5m/s for the anemometer at the lowest elevation. The anemometer at the lowest elevation was set with a more conservative threshold since it was believed that the AANN predicted values had more error for this anemometer. One should note that many of the shear files only had 3 anemometers, so in many instances only the first 3 thresholds are used. Considering that a fault is based on a lagging anemometer, a fault is declared if any of the figure of merit values are below its threshold and healthy otherwise. An example result for the figure of merit values is provided in Figure 14; this result is for the first anemometer for the shear data set. In this example, one can observe that the majority of the files are above the threshold. In total, 62 of the 255 shear anemometers were considered to be in a faulty state.

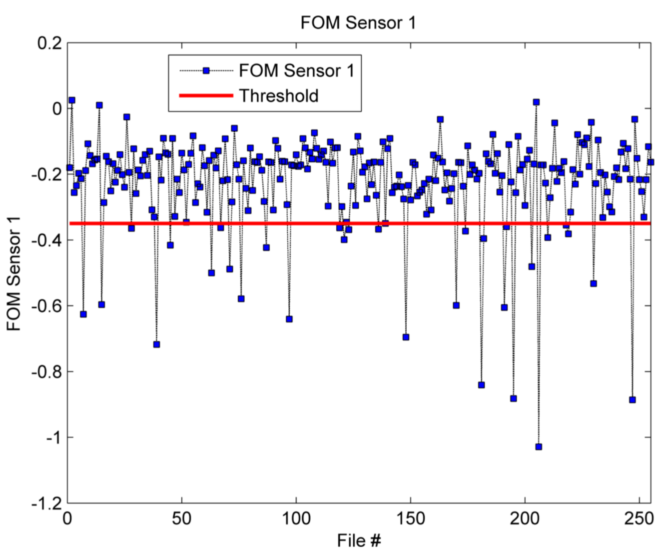

Figure 14. Figure of Merit Results for Shear Data Set

The figure of merit thresholds for both paired anemometers were set at -0.375m/s respectively. If the figure of merit value for a paired anemometer is below the threshold, that anemometer is considered in a failed condition and healthy otherwise. Although the algorithm is based on the difference signal and detecting degraded behavior for one of the anemometers; there were a few occurrences when the algorithm detected that both anemometers were in a failed state. This can occur if both anemometers are lagging but not in the same operating regime regarding wind speed or direction.

The figure of merit results for the paired data set is shown in Figure 15. The results show that the majority of files for the paired data set are detected in a healthy state. The paired health assessment algorithm detected 50 files with a degraded first anemometer, 43 files with a degraded second anemometer, and 325 files were classified as being in a healthy state. Only 2 files were detected as having both paired anemometers in a degraded state. This could be an indication that the algorithm was only suited for anemometer fault detection if there is at least one anemometer in a baseline state.

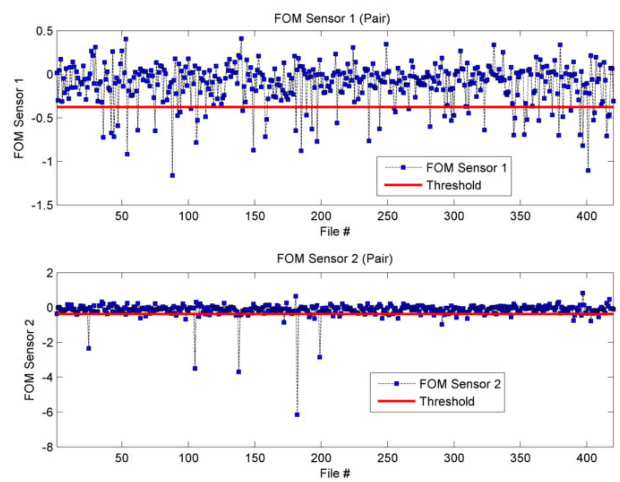

Figure 15. Figure of Merit Results for Paired Data Set

7. CONCLUSIONS

This paper introduced a health assessment methodology for assessing the condition of anemometers in two different configurations. The methodology consisted of a series of algorithmic processing steps from data filtering, to a residual calculation, to a k-means figure of merit health value. Although the algorithms for the shear and paired data sets were quite similar, the use of an auto-associative neural network and additional data normalization were performed by the shear health assessment algorithm. The algorithms for the paired and shear data sets resulted in the most accurate results for the Prognostics and Health Management Society 2011 Data Challenge. This highlights its potential merits for anemometer fault detection. In a general sense, this algorithm could be applied to many other sensor health monitoring applications. In particular, the use of auto-associative neural networks and the k-means clustering approach would be advantageous when redundant sensors are not available and the sensors can have an intermittent fault signature. The weighted residual processing using multiple baseline models is also useful for handling unit to unit variances since it weights training models that more accurately match the monitored unit. Extension of this health assessment algorithm can be approached in two directions; refinement for the specific case of anemometer fault detection, and also reconfiguring the algorithm for other applications.

8. SUGGESTIONS FOR FUTURE WORK

The proposed health assessment algorithm for the shear and paired anemometers provided encouraging results and there are several refinements considered for future work. The algorithm used for assessing the paired anemometers was tuned for the situation in which at least one of the two anemometers was in a healthy state. The inclusion of the wind speed variance information is being considered for future work in developing an algorithm that can detect that both paired anemometers are in a degraded state. This would provide a necessary extension to the proposed framework and would provide a way of detecting sensor problems without assuming that at least one of the reference measurements is in a healthy state.

Regarding refinements in the individual processing modules, a natural staring place would be the data filtering and normalization steps. These ultimately provide the inputs for all further processing, and improvement in removing samples due to icing or outlier values would likely aid the algorithms health monitoring accuracy. The use of an auto-associative neural network for calculating residuals can be compared with other residual processing methods, including the use of kernel principal component analysis methods as well as the traditional PCA methods. Also, the weighted residual processing method could be compared to selecting the top 1 or 2 models; the method for fusing the residual values is one area for further research. Although k-means was initially used for clustering the residual signal, density estimation using a mixture of Gaussians or other clustering techniques can be considered. In addition, evaluation of this proposed algorithm on other applications would further test its ability to generalize and work for other engineering systems.

ACKNOWLEDGEMENT

The work is supported by the NSF Industry/University Cooperative Research Center on Intelligent Maintenance systems (IMS) at the University of Cincinnati as well as its company members.

REFERENCES

Antory, D., Kruger, U., Irwin, G.W., & McCullough, G. (2005). Fault Diagnosis in Internal Combustion Engines using Nonlinear Multivariate Statistics. Proceedings of IME-Part I: Journal of Systems and Controls Engineering, vol. 219, pp. 243-258. doi: 10.1243/095965105X9614

Beltran, J., Llombart, A., & Guerrero, J.J. (2009). Detection of Nacelle Anemometers Faults in a Wind Farm. Proceedings of the International Conference on Renewable Energies and Power Quality, April 15-17, Valencia, Spain.

Capriglione, D., Liguroi, C., Pianese, C., & Pietrosanto, A. (2003). On-line Sensor Fault Detection, Isolation, and Accommodation in Automotive Engines. IEEE Transactions on Instrumentation and Measurement, vol. 52, pp. 1182-1189. doi: 10.1109/TIM.2003.815994

Clark, S.H., Clay, O., Goglia, J.A., Hoopes,T.R., Jacobs, L.T., Smith, R.P. (2009). Investigation of the NRG # 40 Anemometer Slowdown, Proceedings of the AWEA Windpower Conference and Exhibition, May 4-7, Chicago, IL.

Hale, E., Fusina, L., & Brower, M. (2011). Correction factors for NRG #40 Anemometers Potentially affected by Dry Friction Whip: Characterization, Analysis, and Validation. Wind Energy, doi: 10.1002/we.476

Hartigan, J.A., & Wong, M.A. (1979). Algorithm AS 136: A K-Means Clustering Algorithm, Journal of the Royal Statistical Society. Series C (Applied Statistics), vol. 28, pp. 100-108.

Hines, J.W., & Garvey, R.D. (2006). Development and Application of Fault Detectability Performance Metrics for Instrument Calibration Verification and Anomaly Detection. Journal of Pattern Recognition Research, vol. 1, pp. 2-15.

Hsu, S.A., Meindl, E.A., & Gilhousen, D.B. (1994). Determining the Power-Law Wind Profile Exponent under Near-Neutral Stability Conditions at Sea. Journal of Applied Meteorology, vol. 33, pp. 757-772.

Hu, X., Qiu, H., & Iyer, N. (2007). Multivariate change detection for time series data in aircraft engine fault diagnostics. IEEE Conference on Systems, Man, and Cybernetics, October 7-10, Montreal, Quebec. Doi: 10.1109/ICSMC.2007.4414131

Jain, A.K., Murty, M.N., & Flynn, P.J. (1999). Data Clustering: A Review. ACM Computing Surveys, vol. 31, pp. 264-323.

Kramer, M.A. (1991). Nonlinear Principal Component Analysis using Autoassociative Neural Networks. AIChe Journal, vol. 37, pp. 233-243. doi: 10.1002/aic.690370209

Kusiak, A., Zheng, H., & Zhang, Z. (2011). Virtual Wind Speed Sensor for Wind Turbines. Journal of Energy Engineering, vol. 137, pp. 59-69. doi: 10.1061/(ASCE)EY.1943-7897.0000035

Mohandes, M.A., Rehman, S., & Halawani, T.O. (1998). A Neural Networks Approach for Wind Speed Prediction. Renewable Energy, vol. 13, pp. 345-354. doi: 10.1016/S0960-1481(98)00001-9

Murakami, S., Mochida, A., & Kato, S. (2003). Development of Local Area Wind Prediction System for Selecting Suitable Site for Windmill. Journal of Wind Energy and Industrial Aerodynamics, vol. 22, pp. 679-688. doi: 10.1016/j.jweia.2003.09.040

Patton, R.J. (1991). Fault detection and diagnosis in aerospace systems using analytical redundancy, Computers & Control Engineering Journal, vol. 2, pp. 127-136.

Petersen, E.L., Mortensen, N.G., Landberg, L., Hujstrup, J., & Frank, H.P. (1998). Wind Power Meteorology Part II: Siting and Models. Wind Energy, vol. 1, pp. 55-72.

Peterson, E.W., & Hennessey Jr., J.P. (1977). On the Use of Power Laws for Estimates of Wind Power Potential. Journal of Applied Meteorology, vol. 17, pp. 390-394.

PHM Society 2011 Data challenge Competition, (2011). [http://www.phmsociety.org/competition/phm/11]

Pollard, D. (1981). Strong Consistency of K-Means Clustering. The Annals of Statistics, vol. 9, pp. 135-140. doi: 10.1214/aos/1176345339

Schwabacher, M. (2005). A Survey of Data Driven Prognostics. In AIAA Infotech@ Aerospace Conference, September 26-29, Arlington, VA.

Thissen, U., Melssen, W.J., & Buydens, L.M.C. (2001). Nonlinear Process Monitoring using bottle-neck neural networks. Analytica Chimica Acta, vol. 446, pp. 371-383. doi:10.1016/S0003-2670(01)01266-1

Venkatasubramanian, V., Rengaswamy, R., Yin, K., & Kavuri, S.N. (2003). A Review of Process Fault Detection and Diagnosis Part I: Quantitative Model-Based Methods. Computers and Chemical Engineering, vol. 27, pp. 293-311. doi: 10.1016/S0098-1354(02)00160-6

Xu, X., Hines, J.W., & Uhrig, R.E. (1999). Sensor Validation and Fault Detection using Neural Networks. In Proceedings of the Maintenance and Reliability Conference, May 10-12, Gatlinburg, TN.

David Siegel is currently a Ph.D. student in Mechanical Engineering at the University of Cincinnati and a research assistant for the Center for Intelligent Maintenance Systems. Related work experience in the field of prognostics and health management include internships at General Electric Aviation and at the U.S Army Research Lab. His current research focus is on component-level prognostic methods as well as health monitoring algorithms for systems operating under multiple operating regimes.

Jay Lee is an Ohio Eminent Scholar and L.W. Scott Alter Chair Professor at the University of Cincinnati and is founding director of National Science Foundation (NSF) Industry/University Cooperative Research Center (I/UCRC) on Intelligent Maintenance Systems which is a multi-campus NSF Center of Excellence between the University of Cincinnati (lead institution), the University of Michigan, and Missouri University of Science and Technology. His current research focuses on autonomic computing, embedded IT and smart prognostics technologies for industrial and healthcare systems, design of self-maintenance and self-healing systems, and dominant design tools for product and service innovation. He is also a Fellow of ASME, SME, as well as International Society of Engineering Asset Management (ISEAM).