

Probabilistic Multi-Layer Perceptrons for Wind Farm Condition Monitoring

Filippo Fiocchi¹, Domna Ladopoulou², and Petros Dellaportas^{2,3}

¹*Department of Computer Science, University College London, London, United Kingdom*

²*Department of Statistical Science, University College London, London, United Kingdom*

³*Department of Statistics, Athens University of Economics and Business, Athens, Greece*

Abstract. We provide a condition monitoring system for wind farms, based on normal behaviour modelling using a probabilistic multi-layer perceptron with transfer learning via fine-tuning. The model predicts the output power of the wind turbine under normal behaviour based on features retrieved from supervisory control and data acquisition (SCADA) systems. Its advantages are that (i) it can be trained with SCADA data of at least a few years, (ii) it can incorporate all SCADA data of all wind turbines in a wind farm as features, (iii) it assumes that the output power follows a normal density with heteroscedastic variance and (iv) it can predict the output of one wind turbine by borrowing strength from the data of all other wind turbines in a farm. Probabilistic guidelines for condition monitoring are given via a CUSUM control chart. We illustrate the performance of our model in a real SCADA data example which provides evidence that it outperforms other probabilistic prediction models.

Keywords: CUSUM control chart, fine-tuning, heteroskedasticity, normal behaviour modelling, transfer learning.

1 Introduction

Due to the environmental and sustainability benefits, many energy producers have shifted their attention towards wind energy; one of the cleanest and fastest-growing sources of renewable energy. However, the severe weather conditions and the often remote locations of wind farms lead to both high operational failure rates and high maintenance costs that account for 20 – 30% of the total cost related to power generation (Santolamazza et al., 2021). There is therefore a need to reduce the high operational and maintenance costs of wind turbines via a cost-effective and high-precision condition monitoring system. Such a system is crucial for an early fault detection and thus minimization of the prolonged wind turbine downtime.

While specific monitoring sensors can be employed to detect incoming faults, wind farm operators are sceptical of employing expensive technology without direct economic justification (Zaher et al., 2009). Most wind farms have already been equipped with supervisory control and data acquisition (SCADA) systems. These systems record the operational status of individual wind turbines and all their components, such as wind speed, power, rotor speed and blade pitch angles. They produce large amounts of data, which are collected at a high frequency, for example, one Hz, and recorded as a 10-minute averaged interval. Since these systems are already installed, monitoring performance and predicting early faults by extracting features from SCADA data is a widely researched cost-effective measure of failure prevention.

A typical approach to the problem is to produce a model that predicts the output power of a wind turbine under normal behaviour based on features retrieved from SCADA data; see, for example, Rogers et al. (2020). Then, one could monitor the

wind turbine operation by using the model to inspect whether the observed and model-predicted output powers deviate in some statistical sense. Although this is a sensible, well-known approach, practical and effective condition monitoring of a wind farm requires a modelling perspective with four particular characteristics, outlined below.

First, to accommodate weather fluctuations, the model needs to be able to scale well with the data size since it requires to be trained for at least one calendar year which is 52,560 10-minute intervals. Second, it should be able to extract all the wealth of information available in the SCADA system which is expressed in a few dozen features per wind turbine per time interval. Third, the model needs to provide probabilistic predictions in the form of a predictive density and not only point estimates of output power so that a proper probabilistic assessment of the deviation of the observed from the expected output power can be performed. The heteroscedastic nature of the predictions is also necessary to properly model the wind turbine output power because low and high wind speeds produce low variance power distributions. This immediately requires that the predictive density has sufficiently good coverage probabilities in out-of-sample data scenarios. Last but not least, the model should be able to deal with the peculiarities of wind farm SCADA data recording in the following sense. It is common that historical SCADA data may have many missing data in a particular wind turbine because it was out-of-order for a long period of time or it has been recently installed. A good model should be able to predict the output power of this wind turbine with an inferential procedure that borrows strength from the features and output power of all other wind turbines in the wind farm.

We develop a condition monitoring system based on a model that has all the four necessary characteristics described above

based on a probabilistic multi-layer perceptron (PMLP) with transfer learning via fine-tuning. We assume that the output power follows a normal distribution with input-dependent mean and input-dependent variance which form our predictive density for each 10-minute interval. For the probabilistic condition monitoring we propose the use of CUSUM control charts. We illustrate its performance in a real data application by comparing it with two other probabilistic models: a Gaussian process and a Bayesian neural network. We found that our model predicts better with respect to root mean square error, mean absolute error and maximum calibration error. We demonstrate that by transfer learning we can use information from all turbines in a wind farm to improve the prediction of the output power of a single wind turbine. We also provide a real data example in which our proposed condition monitoring system expressed via a CUSUM control chart reveals an early warning in a particular wind turbine failure.

The rest of this paper is organized as follows. Section 2 contains related work, Section 3 presents the proposed methodology, Section 4 presents the empirical application, Section 5 discusses further developments and Section 6 concludes with our key findings.

2 Related Work

Our proposed modelling perspective is based on the notion of normal behaviour modelling, which attempts to produce, conditional on SCADA input features, predictions of the output power of a wind turbine under healthy conditions and then diagnose possible anomalous performance by comparing the predicted and the observed output powers. This is an unsupervised learning strategy which differs from another strand of the literature that uses supervised learning and exploits fault instances obtained from operational and events files from a SCADA system. Thus, in this Section, we will only present related work from articles on modelling normal behaviour. The more relevant to our work papers that use probabilistic predictions will be presented in subsection 2.2.

Normal behaviour modelling studies vary in their selection of the input features used to identify faults. Predictions are performed to the output power (Kusiak and Li, 2011; Li et al., 2014; Pandit et al., 2020; Rogers et al., 2020; Schlechtingen et al., 2013), the generator's temperature (Guo et al., 2012; Schlechtingen and Santos, 2011; Zaher et al., 2009), the gearbox bearings' temperature (Kusiak and Verma, 2012; Schlechtingen and Santos, 2011; Xiang et al., 2021; Zaher et al., 2009; Zhang and Wang, 2014), or multiple output variables (Bangalore and Tjernberg, 2015; Kong et al., 2020; Schlechtingen and Santos, 2011).

2.1 Non-probabilistic normal behaviour modelling

A series of articles (Bangalore and Tjernberg, 2015; Kusiak and Verma, 2012; Schlechtingen and Santos, 2011; Xiang et al., 2021; Zaher et al., 2009; Zhang and Wang, 2014), proposed deep neural networks to model operational characteristics

such as the temperature of the gearbox bearings, cooling oil and winding temperature of the wind turbine. Autoencoders have been proposed by Zhao et al. (2018) and Yang and Zhang (2021) whereas generative adversarial networks combined with autoencoders have been used by Jin et al. (2021). Other methods include (Guo et al., 2012), who predicted the generator's temperature with a nonlinear state estimate technique, and (Schlechtingen et al., 2013), who proposed an adaptive neuro-fuzzy interference system to predict the output power.

2.2 Probabilistic normal behaviour modelling

Due to their intrinsic probabilistic nature, Gaussian processes are common candidates for probabilistic normal behaviour modelling, see for example Li et al. (2014); Pandit et al. (2020); Rogers et al. (2020). The large impediment to their widespread application, namely their computational complexity which increases cubically with the data sample size, is usually being dealt with variational inference with inducing points (Titsias, 2009). However, different input features increase the hyperparameter dimension since proper implementation requires different kernels in each dimension, which creates both convergence and kernel identification problems. Finally, although the issue of heteroscedasticity can be solved by assuming that the error process also follows a Gaussian process (Lázaro-Gredilla and Titsias, 2011), it is extremely difficult to resolve all three issues above simultaneously. Finally, although multi-task Gaussian processes is a modelling perspective that simultaneously exploits the data of all wind turbines, the computational complexity increases quickly and treating missing values from one or more wind turbines is not possible. The works of Le et al. (2005); Lázaro-Gredilla and Titsias (2011) have incorporated noise heteroscedasticity whereas the trade-off between heteroscedasticity and computational complexity has been considered by Rogers et al. (2020). Recently, Deng et al. (2024) addressed the necessity of a probabilistic setting for power curve estimation by incorporating prior distributions on the weight parameters of a long short term memory (LSTM) Bayesian neural network. They improve their model by using a temporal convolutional neural network for temporal feature learning and an embedding layer to map discrete features (e.g., integer years) to dense vectors. The predictions are obtained via Monte Carlo sampling. Finally, Zhang et al. (2020) considered the problem of predicting output power conditional on future wind speed and direction forecasts and adopted a modelling approach that closely resembles ours. They employed an improved deep mixture density network by transforming the output energy to the interval $(0, 1)$ and obtained probabilistic power predictions by assuming that it follows a mixture of Beta distributions. As is similarly emphasized in our study, they highlighted the fact that informed decision-making is one of the key benefits of a probabilistic approach to wind power prediction.

3 Proposed Methodology

We present a modelling perspective guided by the need to provide a condition monitoring system that has immediate practical applications in any wind farm. We first describe a fully connected deep neural network that predicts both the mean and the variance of output power at a particular 10-minute interval. The computational power of the neural network, together with the stochastic heteroscedastic output power, allows the incorporation of many features obtained from large SCADA datasets that achieve good training based on as many as possible environmental and operational conditions and a statistically sound monitoring system based on both the mean and the variance of the power output. The final ingredient in our model that is hugely important and necessary in realistic condition monitoring systems is the incorporation of transfer learning which is simultaneously trained in all wind turbines in a wind farm. Finally, we present a realisation of our proposed probabilistic condition monitoring system based on a CUSUM control chart. The overall flowchart of the proposed probabilistic condition monitoring system is presented in Fig. 1.

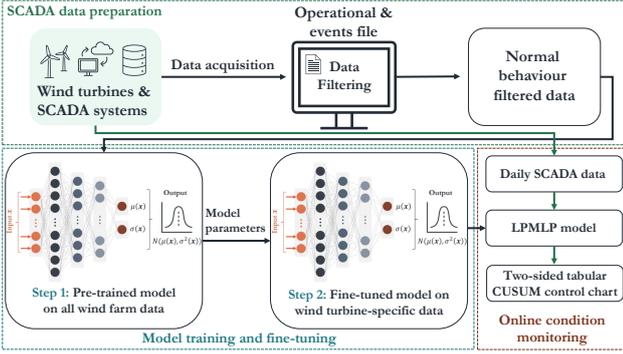


Figure 1. Overall flowchart of the proposed probabilistic condition monitoring system. LPMLP: Large probabilistic multi-layer perceptron presented in Section 3.2.

3.1 Probabilistic multi-layer perceptron

To model a heteroscedastic output noise we assume that for each 10-minute interval the power output of a wind turbine y conditioned on the SCADA input features $\mathbf{x} \in \mathcal{R}^{d_0}$ follows a normal distribution with unknown mean and variance that both depend on \mathbf{x} , so $y \sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$. We will approximate $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ by a probabilistic multi-layer perceptron (PMLP) with branching prediction heads. We now describe the architecture of our proposed model.

Given an input vector $\mathbf{x} = (x_1, x_2, \dots, x_{d_0})$ that represents the d_0 SCADA input features at a particular time interval, the predicted PMLP power output $\hat{y} \in \mathcal{R}$ follows a normal density with parameters $\mu(\mathbf{x})$ and $\sigma^2(\mathbf{x})$ which are continuous functions of the inputs \mathbf{x} . In fact, our proposed PMLP models $\mu(\mathbf{x})$ as a linear piecewise function of \mathbf{x} and $\sigma(\mathbf{x})$ as a nonlinear piecewise function of \mathbf{x} .

For a number of k hidden layers with widths d_1, d_2, \dots, d_k and input and output dimensions d_0 and $d_{k+1} = 1$ respectively, we adopt the linear functions $g_i : \mathcal{R}^{d_{i-1}} \rightarrow \mathcal{R}^{d_i}$ for $i = 1, \dots, k$ defined as

$$g_i(\mathbf{x}) = \mathbf{W}_i \mathbf{x} + \mathbf{b}_i,$$

where $\mathbf{W}_i \in \mathcal{R}^{d_i \times d_{i-1}}$, $\mathbf{x} \in \mathcal{R}^{d_{i-1}}$ and $\mathbf{b}_i \in \mathcal{R}^{d_i}$. Furthermore, we adopt the Rectified Linear Unit function $ReLU_i : \mathcal{R}^{d_i} \rightarrow \mathcal{R}^{d_i}$ defined as

$$ReLU_i(\mathbf{x}) = (\max\{0, x_1\}, \max\{0, x_2\}, \dots, \max\{0, x_{d_i}\}),$$

and the Softplus function $S : \mathcal{R} \rightarrow \mathcal{R}$ defined as

$$S(x) = \log(1 + e^x) + \delta \quad (1)$$

for a small $\delta > 0$ used for numerical stability needed when x becomes very small. These three functions are the basic ingredients of our PMLP, but we note that the functions $ReLU$ and S can be replaced by any other appropriate alternative.

We now define the branching structure of the PMLP. We use two sets of functions g and $ReLU$, corresponding to predictions of $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ and indexed by the superscripts μ and σ respectively. Different functions g have different sets of parameters $\mathbf{W}_i, \mathbf{b}_i$ where by different functions $ReLU$ is meant that they operate on different dimensions d_i . We further assume that there exist two corresponding numbers of hidden layers k_μ and k_σ with widths $d_1^\mu, d_2^\mu, \dots, d_{k_\mu}^\mu$ and $d_1^\sigma, d_2^\sigma, \dots, d_{k_\sigma}^\sigma$. Then, the output \hat{y} of the PMLP with k_μ and k_σ hidden layers is represented as

$$\begin{aligned} \mu(\mathbf{x}) &= g_{k_\mu+1}^\mu \circ ReLU_{k_\mu}^\mu \circ g_{k_\mu}^\mu \circ \dots \circ g_2^\mu \circ ReLU_1^\mu \circ g_1^\mu(\mathbf{x}) \\ \sigma(\mathbf{x}) &= S \circ g_{k_\sigma+1}^\sigma \circ ReLU_{k_\sigma}^\sigma \circ g_{k_\sigma}^\sigma \circ \dots \circ g_2^\sigma \circ ReLU_1^\sigma \circ g_1^\sigma(\mathbf{x}) \\ \hat{y} &\sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x})) \end{aligned} \quad (2)$$

where \circ denotes the function decomposition operator.

The branching mode of the PMLP is achieved by setting a layer k^* such that $1 \leq k^* \leq \min\{k_\mu, k_\sigma\}$ and setting

$$\begin{aligned} ReLU_i^\mu(\mathbf{x}) &= ReLU_i^\sigma(\mathbf{x}) \\ g_i^\mu(\mathbf{x}) &= g_i^\sigma(\mathbf{x}) \\ d_i^\mu &= d_i^\sigma \end{aligned}$$

for all $1 \leq i \leq k^*$. Thus, before the k^* -th layer there is a common deep neural network; after the k^* -th layer, the network branches into two paths for predicting $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ with each path having its own sequence of hidden layers with potentially different widths and depths. Note that the total sum of hidden layers is $k_{total} = k_\mu + k_\sigma - k^*$. A visual representation that corresponds to the choice of the PMLP architecture given $k_\mu = k_\sigma = 4$, $k^* = 3$, $d_0 = 34$, $d_1^\mu = d_1^\sigma = 100$, $d_2^\mu = d_2^\sigma = 80$, $d_3^\mu = d_3^\sigma = 40$, $d_4^\mu = d_4^\sigma = 20$ that was used in the application of Section 4 is depicted in Fig. 2. We call this particular PMLP architecture A1.

The training of the PMLP is achieved by solving the following empirical risk minimisation problem. Given n data points $(\mathbf{x}_i, y_i) \in \mathcal{R}^{d_0} \times \mathcal{R}$, $i = 1, 2, \dots, n$, and number and widths of the hidden layers, find \hat{y}_i that is represented with an PMLP and minimises $\min_{\mathbf{W}, \mathbf{b}} \mathcal{L}_1$ where $\mathbf{W} = \{\mathbf{W}_i\}_{i=1}^{k_{total}}$,

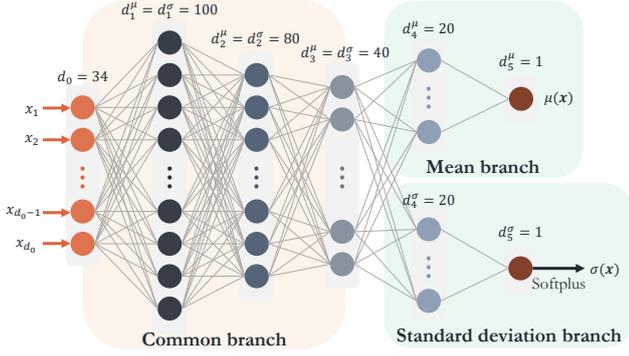


Figure 2. PMLP architecture A1 used in the application of Section 4 with shared initial layers and independent branches for mean and standard deviation.

$\mathbf{b} = \{\mathbf{b}_i\}_{i=1}^{k_{\text{total}}}$ denoting all parameters, and \mathcal{L}_1 is a loss function. By assuming independence between \hat{y}_i , the product of normal densities $\prod_{i=1}^n \mathcal{N}(\mu(\mathbf{x}_i), \sigma^2(\mathbf{x}_i))$ can be viewed as a multivariate predictive density so a plausible loss function for our PMLP is the minus logarithmic score, see for example Gneiting and Raftery (2007), defined as

$$\mathcal{L}_1 = - \sum_{i=1}^n \log \left(\mathcal{N}(y_i | \mu(\mathbf{x}_i), \sigma^2(\mathbf{x}_i)) \right) \quad (3)$$

where $\mathcal{N}(y|\mu, \sigma)$ denotes the p.d.f. of a normal density with mean μ and variance σ^2 evaluated at y .

3.2 Transfer learning via fine-tuning

The realistic application of condition monitoring of a wind turbine should necessarily take into account the fact that the wind turbine is part of a wind farm. Thus, instead of training a model for each wind turbine separately, we could use the PMLP model of subsection 3.1 to train all wind turbines simultaneously. No particular care is needed other than treating the data so that the corresponding feature values of each wind turbine correspond to the power output of this wind turbine: our proposed PMLP treats the power output conditional on the corresponding features of each 10-minute period independently of the power outputs of the other 10-minute periods. Therefore, we can apply model (2) to all wind turbines of the wind farm by just denoting \mathbf{x} as the features from all wind turbines and \hat{y} as a vector power output for a particular 10-minute period. The loss function (3) is just the sum of log-normal densities of each wind turbine at each 10-minute interval. The mathematical formulation of this model is as follows.

Let \mathbf{x}_{ij} be the input feature vector of the j -th wind turbine at the i -th 10-minute interval with corresponding output power y_{ij} . Then, the model is defined as

$$\begin{aligned} \mu(\mathbf{x}_{ij}) &= g_{k_\mu+1}^\mu \circ \text{ReLU}_{k_\mu}^\mu \circ g_{k_\mu}^\mu \cdots g_2^\mu \circ \text{ReLU}_1^\mu \circ g_1^\mu(\mathbf{x}_{ij}) \\ \sigma(\mathbf{x}_{ij}) &= S \circ g_{k_\sigma+1}^\sigma \circ \text{ReLU}_{k_\sigma}^\sigma \circ g_{k_\sigma}^\sigma \cdots g_2^\sigma \circ \text{ReLU}_1^\sigma \circ g_1^\sigma(\mathbf{x}_{ij}) \\ \hat{y}_{ij} &\sim \mathcal{N}(\mu(\mathbf{x}_{ij}), \sigma^2(\mathbf{x}_{ij})). \end{aligned} \quad (4)$$

If we have data for n_j time intervals for wind turbine j for $j = 1, \dots, J$, the loss function of model (4) is

$$\mathcal{L}_2 = - \sum_{j=1}^J \sum_{i=1}^{n_j} \log \left(\mathcal{N}(y_{ij} | \mu(\mathbf{x}_{ij}), \sigma^2(\mathbf{x}_{ij})) \right). \quad (5)$$

Notice that the possibly different number of available data n_j from each wind turbine indicates that this model utilizes all available data from all wind turbines. This is of huge practical importance in realistic applications since SCADA data typically contain many missing data.

However, by training all data of wind farms in the same model, the amount of available data may become prohibitively vast to allow frequent re-training. The usual treatment of such huge datasets is to adopt some transfer learning techniques so that knowledge gained from pre-training the model based on all wind turbines data can be used to boost performance in predicting the power of a single wind turbine. We call such a PMLP model that uses pre-training based on loss function \mathcal{L}_2 and then predicts the power of only one wind turbine using fine-tuning and the loss function \mathcal{L}_1 a Large PMLP (LPMLP) model. When new data arrive, re-training of the PMLP model based on \mathcal{L}_2 is not anymore necessary. Moreover, additional important practical advantages arise. First, consider the very realistic scenario in which a wind turbine has fewer data points because, for example, it has been out-of-order for a long period of time or has been recently installed. Normal behaviour modelling of this wind turbine might be very hard or even impossible to achieve. Our LPMLP model is capable of producing a predictive density for such low-information wind turbines by using the data from all wind turbines in the wind farm. Second, transfer learning can be used to improve the predictive power of model (2) as we empirically show in Section 4.

Our proposed transfer learning is achieved via the following fine-tuning. We first train the LPMLP model (4) with the loss function (5). Then, we use the parameters of the pre-trained model as initial values, and we predict the mean and standard deviation of the output power of the wind turbine we are interested in. Thus, the training uses the loss function (3) that refers to one only wind turbine and is achieved very fast. Fig. 3 depicts the architecture of one LPMLP used in the application of Section 4 based on six wind turbines with $k_\mu = k_\sigma = 3$, $k^* = 3$, $d_0 = 34$, $d_1^\sigma = d_1^\mu = 300$, $d_2^\sigma = d_2^\mu = 200$, $d_3^\sigma = d_3^\mu = 100$. We call this particular architecture A2. To show the scalability difference, we note that for our model in Fig. 3 the input vectors correspond to all features of all wind turbines, and the model has 91,002 parameters whereas the PMLP model of Fig. 2 has only 16,508 parameters.

For the advantages of transfer learning in similar problems see, for example, Ma et al. (2024). Notable theoretical and empirical motivations for introducing transfer learning in the given probabilistic regression context include (i) regularisation: pre-training the LPMLP allows learning general features and patterns from a larger amount of data which can be especially beneficial in the case where the operating conditions of the wind farm are extremely diverse. Thus, it is easier for the

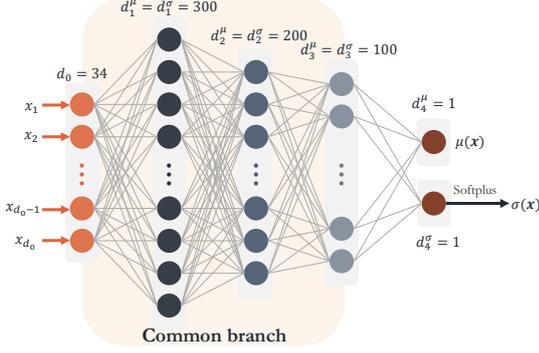


Figure 3. LPMLP architecture A2 used for both pre-training on wind farm data and fine-tuning on wind turbine-specific data used in the application of Section 4 with shared initial layers.

model to deal with conditions that have not been observed in the particular wind turbine under inspection; (ii) data sparsity: pre-training on the whole data of a wind farm dataset can help alleviate data sparsity by providing additional data for learning more general but not yet observed representations. For example, see Liao et al. (2023) for a discussion of limited data availability for newly built wind farms; and (iii) improved convergence: pre-training the LMPL facilitates faster convergence during fine-tuning on a single wind turbine.

3.3 Condition Monitoring

The reason that we propose the PMLP and LPMLP models which have a probabilistic flavour by predicting both the mean and the variance of the output power is to construct a reasonable, probability-based, data-driven condition monitoring system. This system has the usual ingredients that consist of first training the model during a healthy wind turbine operation and then observing both SCADA and output power data. Anomalous operational behaviour is reported when the observed output power departs, in a probabilistic sense, from the expected predicted normal density under the trained model. Our condition monitoring system is based on testing model adequacy by testing the hypothesis that the observed output power observations have been generated by our trained model. This requires, of course, some evidence that our model has sufficiently good coverage probability so that in an out-of-sample healthy wind turbine environment the observed power output indeed follow the normal density indicated by the trained model. As it will be evident in our application study, our proposed PMLP and LPMLP models do indeed have good coverage probabilities in an out-of-sample large empirical exercise.

Assume that the model training phase has ended and we are now in the live condition monitoring phase. To emphasize that the data is observed sequentially in time, we will be now using a subscript t rather than i for the time intervals. Assume that we observe T consecutive power outputs y_t , $t = 1, 2, \dots, T$, and we have predictions from our model $N(\mu(\mathbf{x}_t), \sigma^2(\mathbf{x}_t))$. Under the hypothesis of model adequacy, or no-fault, we define $v_t = (y_t - \mu(\mathbf{x}_t))/\sigma(\mathbf{x}_t)$ and test whether the sample

v_t , $t = 1, 2, \dots, T$ comes from a $N(0, 1)$ distribution. For example, we could perform a test on a daily basis for 10-minute intervals where $T = 144$.

We propose a two-sided tabular CUSUM control chart (Montgomery, 2009) defined as

$$S_H(t) = \max \{0, v_t - k + S_H(t - 1)\}$$

and

$$S_L(t) = \max \{0, -k - v_t + S_L(t - 1)\},$$

where $S_H(0) = S_L(0) = 0$. The reference (or allowance or slack) value of k is traditionally specified with past experience or from simulations based on supervised settings. A common choice is about halfway between the target mean zero and the out-of-control value that we are interested in detecting quickly. Therefore, if we are interested in detecting a deviance of one standard deviation a plausible choice is $k = 1/2$; see Montgomery (2009). Recent advances propose an adaptive adjustment of its value (Wu et al., 2009). Note that $S_H(t)$ and $S_L(t)$ are viewed as cumulative deviations from zero that are greater than k that are not allowed to take negative values. If either of them exceeds a decision interval I , the process is considered to be out-of-control. A reasonable choice of I is five times the process standard deviation (Montgomery, 2009) which is one in our case, so we propose $I = 5$. For $t = 1, \dots, T$ the test statistics are defined as $A_t^H = \max \{S_H(i), i = 1, \dots, t\}$ and $A_t^L = \max \{S_L(i), i = 1, \dots, T\}$. The null hypothesis of normal operation in the wind turbine is rejected (with a one-sided test) if $A_n = \max (A_t : 1 \leq t \leq n)$ exceeds its upper critical value, where $A_t = \max (A_t^H, A_t^L)$.

4 Experiments

4.1 Data Acquisition and Filtering

The application utilises 10-minute SCADA and events data from the six Senvion MM92 wind turbines at Kelmarsh wind farm in the UK (Plumley, 2022). The dataset spans from 3rd January 2016 to 1st July 2021 and comprises over 1.7 million data points containing 110 variables including date-time, wind speed, bearing temperature, and power output. The data includes 10-minute averages, standard deviations, minimum and maximum values of all measured variables. The recording of wind speed summary statistics started at 25th September 2017 so we considered only data entries from this date onwards. An operational status and events file was utilised for data filtering to ensure consistent modelling of normal behaviour. These files provide valuable insights into the operational conditions of the wind turbines, covering a range of scenarios from technical failures to operational or environmental standbys and warnings. To enable accurate behaviour modelling, it was essential to remove out-of-control condition data records and base our model training solely on filtered data. The data elimination process is illustrated in Fig. 4 and 5 where plots of output powers against wind speeds are depicted for each wind turbine before and after the filtering.

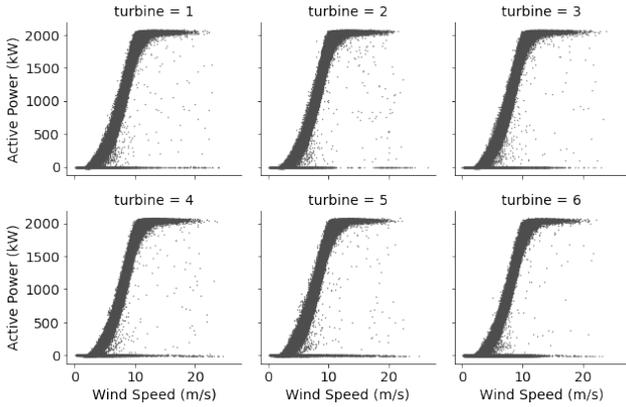


Figure 4. Real data wind power curves for all six wind turbines at Kelmarsh wind farm before removing standbys and warnings existing in the operational status and event file.

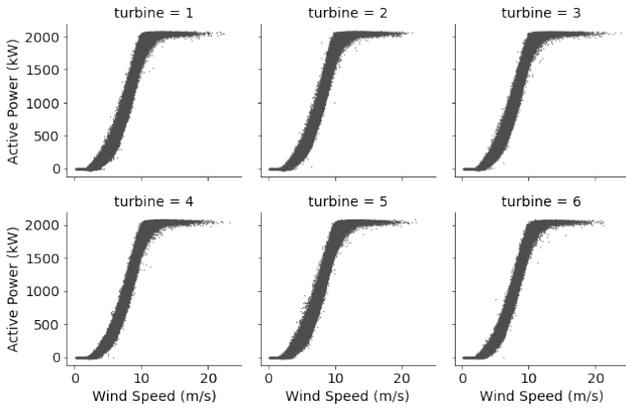


Figure 5. Real data wind power curves for all six wind turbines at Kelmarsh wind farm after removing standbys and warnings using the operational status and events file provided by the data provider Plumley (2022).

All our analyses used as input features 34 operational and environmental variables are listed in Table 1. The final filtered dataset across the six wind turbines comprises 966,936 data points, with 189,381 of them belonging to one single wind turbine used in the application Section 4.2.

4.2 Application

For all our experiments, we used a randomly selected 80 – 20 data ratio for the train and test data, respectively. When a validation set was required, the train-validation data ratio was 72 – 8. The PMLP and LPMLP models were trained using the Adam optimiser with a learning rate of 0.001, a batch size of 32, and 100 and 500 epochs respectively. This was done for both A1 and A2 architectures. Additionally, the LPMLP models were fine-tuned for 50 epochs using learning rates of 10^{-3} and 10^{-4} for A1 and A2 respectively. The Softplus activation function (1) used a $\delta = 0.001$. Early stopping on a validation set was employed during all training stages. All results refer to the out-of-sample prediction of the output

Table 1. SCADA operational and environmental variables used as input features in the proposed models. Avg: Average; Stdev: Standard deviation; Min: Minimum; Max: Maximum; temp: temperature.

| <i>Feature description</i> | <i>Feature description</i> |
|----------------------------|-------------------------------|
| Avg. wind speed | Avg. front bearing temp. |
| Stdev. wind speed | Stdev. front bearing temp. |
| Min. wind speed | Min. front bearing temp. |
| Max. wind speed | Max. front bearing temp. |
| Avg. rear bearing temp. | Avg. rotor bearing temp. |
| Stdev. rear bearing temp. | Avg. stator1 temp. |
| Min. rear bearing temp. | Avg. nacelle ambient temp. |
| Max. rear bearing temp. | Avg. nacelle temp. |
| Avg. transformer temp. | Avg. gear oil temp. |
| Avg. gear oil inlet temp. | Avg. drive train acceleration |
| Avg. top box temp. | Avg. hub temp. |
| Avg. conv. ambient temp. | Avg. transformer cell temp. |
| Avg. motor axis1 temp. | Avg. motor axis2 temp. |
| Avg. CPU temp. | Avg. blade angle pitch A |
| Avg. blade angle pitch B | Avg. blade angle pitch C |
| Avg. gear oil inlet press | Avg. gear oil pump press |
| Tower acceleration x | Tower acceleration y |

power of only one randomly chosen wind turbine of the wind farm.

To evaluate the proposed probabilistic PMLP and LPMLP modelling frameworks, we compared their performance to other probabilistic methods. In particular, we used a Gaussian process regression and a Bayesian neural network with stochastic output. The Gaussian process inference was based on variational inference with inducing points (Snelson and Ghahramani, 2005) using a radial basis function kernel and 100 pseudo-inputs. This approximate inference procedure was used because the large sample size required low computational complexity methodologies, a problem quite useful in Gaussian process literature; see, for example, Titsias (2009). The Bayesian Neural Network was defined as a set of fully connected layers with mean field normal hierarchical priors with mean and standard deviation having $\mathcal{N}(0, 0.1)$ and $\mathcal{N}(S(0.001), 0.1)$ hyperpriors respectively; here, S is the Softplus function defined in (1) with $\delta = 0$. It was trained using the Flipout method which re-models the stochastic variational inference implementation strategy as a weight perturbation (Wen et al., 2018) over 100 epochs using early stopping with the loss given by (3).

The results shown in Table 2 are based on 151,504 and 37,877 training and test 10-minutes intervals respectively. Note that the models have been trained with SCADA data which have the size of nearly three years of wind turbines operation. We base our comparison with the usual root mean square and mean absolute error defined as $\text{RMSE} = (n^{-1} \sum_{i=1}^n (y_i - \mu(\mathbf{x}_i))^2)^{1/2}$ and $\text{MAE} = n^{-1} \sum_{i=1}^n |y_i - \mu(\mathbf{x}_i)|$ respectively. Although these metrics provide a good indication on how well our model predicts the output power, it is based only on the predicted means $\mu(\mathbf{x}_i)$. An issue of more importance to our condition monitoring methodology is the out-of-sample coverage probabilities that are defined as the probabilities that a confidence interval region will include the

true power output. These are empirically estimated in our test data and shown in Fig. 6 where the calibration error, defined as the difference between observed and theoretical coverage probabilities are plotted for all six models and for twenty different intervals. The maximum calibrated error, which is the maximum observed deviation, is reported in Table 2.

Notice that the architectures A1 and A2 of the PMLP and LPMLP models are different. We incrementally built these architectures by inspecting loss functions in training and testing data. For our empirical exercise, it is of particular interest to investigate whether the scaling of the architecture or the transfer learning is the actual cause of the model performance improvement. To illustrate this, Table 2 and Fig. 6 also contains results of the PMLP model with architecture A2 and of the LPMLP model with architecture A1 without pre-training. By controlling the network architecture, there is evidence that pre-training improves both RMSE and MAE metrics.

Table 2. Out-of-sample performance metrics. PMLP: probabilistic multi-layer perceptron; LPMLP: large probabilistic multi-layer perceptron; RMSE: root mean square error; MAE: mean absolute error; MCE: maximum calibration error. Best performance is indicated in bold.

| Method | RMSE | MAE | MCE |
|-----------------------------------|--------------|--------------|-------------|
| <i>Sparse Gaussian process</i> | 45.17 | 30.44 | 17.79 |
| <i>Bayesian neural network</i> | 30.06 | 16.47 | 1.83 |
| <i>PMLP with architecture A1</i> | 28.19 | 14.99 | 1.12 |
| <i>LPMLP with architecture A1</i> | 25.38 | 13.62 | 1.91 |
| <i>PMLP with architecture A2</i> | 26.09 | 14.17 | 1.04 |
| <i>LPMLP with architecture A2</i> | 23.73 | 12.67 | 1.77 |

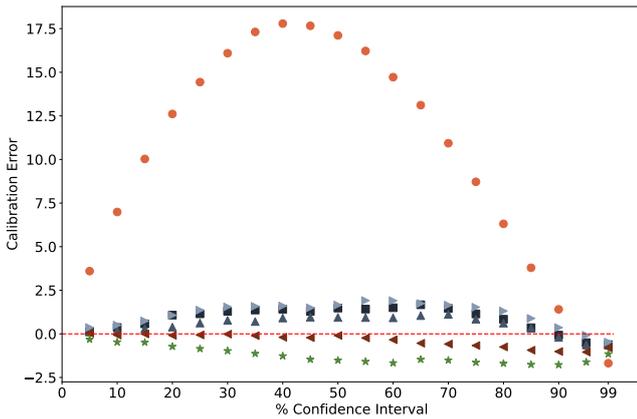


Figure 6. Calibration error of twenty binned confidence levels including the 95% and 99% confidence intervals. ● Sparse Gaussian process (RBF kernel); ▲ PMLP with architecture A1; ▼ PMLP with architecture A2; ■ Bayesian neural network; ► LPMLP with architecture A1; ★ LPMLP with architecture A2.

We illustrate our proposed unsupervised method for condition monitoring using two different examples, one in which the examined wind turbine is operating normally and one in which the process is observed to be out-of-control; the results

of these examples are presented in the CUSUM control charts of Fig. 7 and 8 respectively.

In both cases, observations obtained in daily intervals of $T = 144$ 10-minute intervals are being monitored using CUSUM control charts. Fig. 7 presents a typical situation of a good operation of a wind turbine.

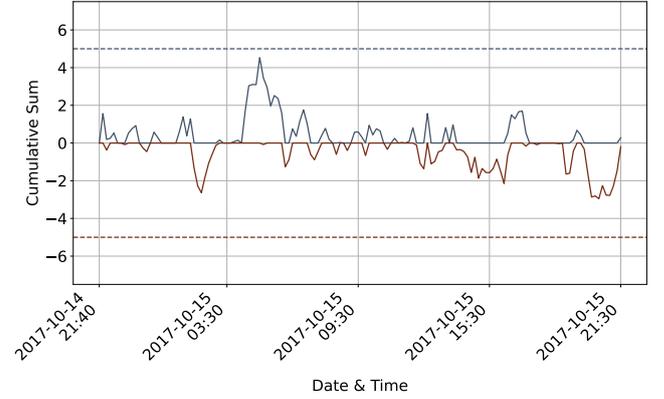


Figure 7. CUSUM control chart for 24 hours during normal wind turbine behaviour. In the chart (—), (—), (---), and (---) represent S_H , S_L , I and $-I$.

In contrast, Fig. 8 the CUSUM control chart indicates that the process is out-of-control at 07:20AM on the 17th of January 2018. By inspecting the operational status and event file of this wind turbine we identified that there is a triggered alarm on the 18th of January 2018 at 04:22:23AM about a forced outage concerning a fault. It is, therefore, evident that our monitor system would have identified the fault 21 hours prior to the occurrence.

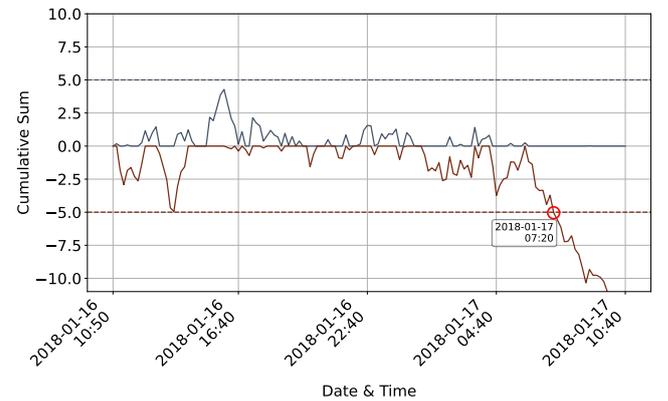


Figure 8. CUSUM control chart for 24 hours prior to an operational fault started on the 18th of January 2018 at 04:22:23AM. In the chart (—), (—), (---), and (---) represent S_H , S_L , I and $-I$. The process starts to be out-of-control on 17th of January 2018 at 07:20AM.

5 Discussion

When the condition monitoring system identifies malfunction, such as the one in Fig. 8, the engineers in the field would like

to identify which component of the wind turbine is causing the problem. To achieve this, we could extend our method to monitor different variables of the SCADA inputs, such as temperatures or pressures across various parts of the wind turbine. An obvious way to achieve this is to replicate our methodology by replacing the output power with another variable. A complete condition monitoring system would thus be a collection of CUSUM control charts with different outputs modelled with different LPMLP models.

A further methodological direction would be to replace the fine-tuning of LPMLP which is based on parameter initialisation; see, for example, Shwartz-Ziv et al. (2022). We could use the parameter estimates of a pre-trained model to construct informative prior distributions for transfer learning. Then, highly informative posteriors will be available for the model that predicts the power of one single wind turbine. This is, in essence, a Bayesian neural network with informative priors that replace the vague priors we used in our application.

6 Conclusion

We presented a condition monitoring system that can be immediately applied to a wind farm. We took particular care to accommodate issues that are routinely met in pragmatic wind farm operations. In particular, all SCADA data from all wind turbines collected in many years can be used to train our model. Our monitoring system has the ability to incorporate wind turbines that have many missing data. It is based on a model that produces heterogeneous predictive densities that are well suited to wind turbine data; its probabilistic nature provides a scientifically sound monitoring system that produces automatic monitoring of a wind farm in the form of a control chart.

For the data we used, our probabilistic model outperforms other probabilistic methods in terms of RMSE and MAE and has good coverage probabilities in large out-of-sample empirical exercises. Although it is an unsupervised monitoring system, we were able to illustrate its immediate applicability by inspecting the fault events file after an anomaly is detected.

Future research involves investigating ways to incorporate a further step that will identify which operational characteristic of the wind turbine caused the anomaly detected by the control chart.

Acknowledgments

We would like to thank Dr Evangelos Morfiadakis, CEO at International Wind Engineering, for introducing us to the problem and providing valuable insight; and Dr Charlie Plumley, Senior Performance Engineer at Glennmont Partners, for helpful discussions.

References

- Bangalore, P. and Tjernberg, L. B. (2015). An artificial neural network approach for early fault detection of gearbox bearings. *IEEE Transactions on Smart Grid*, 6(2):980–987.
- Deng, Z.-G., Zhang, X., Li, Z., Yang, J., Lv, X., Wu, Q., and Zhu, B. (2024). Probabilistic prediction of wind power based on improved Bayesian neural network. *Frontiers in Energy Research*, 11.
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.
- Guo, P., Infield, D., and Yang, X. (2012). Wind turbine generator condition-monitoring using temperature trend analysis. *IEEE Transactions on Sustainable Energy*, 3(1):124–133.
- Jin, X., Xu, Z., and Wang, Q. (2021). Condition monitoring of wind turbine generators using scada data analysis. *IEEE Transactions on Sustainable Energy*, 12(1):202–210.
- Kong, Z., Tang, B., Deng, L., Liu, W., and Han, Y. (2020). Condition monitoring of wind turbines based on spatio-temporal fusion of scada data by convolutional neural networks and gated recurrent units. *Renewable Energy*, 146:760–768.
- Kusiak, A. and Li, W. (2011). The prediction and diagnosis of wind turbine faults. *Renewable Energy*, 36(1):16–23.
- Kusiak, A. and Verma, A. (2012). Analyzing bearing faults in wind turbines: A data-mining approach. *Renewable Energy*, 48:110–116.
- Lázaro-Gredilla, M. and Titsias, M. K. (2011). Variational heteroscedastic gaussian process regression. In *ICML*, pages 841–848.
- Le, Q. V., Smola, A. J., and Canu, S. (2005). Heteroscedastic gaussian process regression. In *Proceedings of the 22nd international conference on Machine learning*, pages 489–496.
- Li, J., Xiao, L., Li, H., and Li, R. (2014). Normal behavior models for the condition assessment of wind turbine generator systems. *Electric Power Components and Systems*, 42(11):1201–1212.
- Liao, Q., Cao, D., Chen, Z., Blaabjerg, F., and Hu, W. (2023). Probabilistic wind power forecasting for newly-built wind farms based on multi-task gaussian process method. *Renewable Energy*, 217:119054.
- Lázaro-Gredilla, M. and Titsias, M. K. (2011). Variational heteroscedastic gaussian process regression. In *ICML*, pages 841–848.
- Ma, Y., Chen, S., Ermon, S., and Lobell, D. B. (2024). Transfer learning in environmental remote sensing. *Remote Sensing of Environment*, 301:113924.

- Montgomery, D. C. (2009). *Statistical quality control*, volume 7. Wiley New York.
- Pandit, R., Infield, D., and Kolios, A. (2020). Gaussian process power curve models incorporating wind turbine operational variables. *Energy Reports*, 6:1658–1669.
- Plumley, C. (2022). Kelmarsh wind farm data (0.0.3). *Zenodo*.
- Rogers, T. J., Gardner, P. A., Dervilis, N., Worden, K., Maguire, E., Papatheou, E., and Cross, E. J. (2020). Probabilistic modelling of wind turbine power curves with application of heteroscedastic gaussian process regression. *Renewable Energy*, 148:1124–1136.
- Santolamazza, A., Dadi, D., and Introna, V. (2021). A data-mining approach for wind turbine fault detection based on scada data analysis using artificial neural networks. *Energies*, 14(7):1845.
- Schlechtingen, M. and Santos, I. (2011). Comparative analysis of neural network and regression based condition monitoring approaches for wind turbine fault detection. *Mechanical Systems and Signal Processing*, 25(5):1849–1875.
- Schlechtingen, M., Santos, I., and Achiche, S. (2013). Using Data-Mining Approaches for Wind Turbine Power Curve Monitoring: A Comparative study. *IEEE Transactions on Sustainable Energy*, 4(3):671–679.
- Shwartz-Ziv, R., Goldblum, M., Souri, H., Kapoor, S., Zhu, C., LeCun, Y., and Wilson, A. G. (2022). Pre-train your loss: Easy bayesian transfer learning with informative priors. *Advances in Neural Information Processing Systems*, 35:27706–27715.
- Snelson, E. and Ghahramani, Z. (2005). Sparse gaussian processes using pseudo-inputs. In Weiss, Y., Schölkopf, B., and Platt, J., editors, *Advances in Neural Information Processing Systems*, volume 18. MIT Press.
- Titsias, M. (2009). Variational learning of inducing variables in sparse gaussian processes. In *Artificial intelligence and statistics*, pages 567–574. PMLR.
- Wen, Y., Vicol, P., Ba, J., Tran, D., and Grosse, R. (2018). Flipout: Efficient pseudo-independent weight perturbations on mini-batches. In *International Conference on Learning Representations*.
- Wu, Z., Jiao, R., Yang, M., Liu, Y., and Wang, Z. (2009). An enhanced adaptive cusum control chart. *Iie Transactions*, 41:642–653.
- Xiang, L., Wang, P., Yang, X., Hu, A., and Su, H. (2021). Fault detection of wind turbine based on scada data analysis using cnn and lstm with an attention mechanism. *Measurement*, 175:109094.
- Yang, L. and Zhang, Z. (2021). Wind turbine gearbox failure Detection Based on SCADA Data: A Deep Learning-Based approach. *IEEE transactions on instrumentation and measurement*, 70:1–11.
- Zaher, A., McArthur, S., Infield, D., and Patel, Y. (2009). Online wind turbine fault detection through automated scada data analysis. *Wind Energy*, 12(6):574–593.
- Zhang, H., Liu, Y., Yan, J., Han, S., Li, L., and Quan, L. (2020). Improved deep mixture density network for regional wind power probabilistic forecasting. *IEEE transactions on power systems*, 35(4):2549–2560.
- Zhang, Z. and Wang, K. (2014). Wind turbine fault detection based on scada data analysis using ann. *Advances in Manufacturing*, 2(1):70–78.
- Zhao, H., Liu, H., Hu, W., and Yan, X. (2018). Anomaly detection and fault analysis of wind turbine components based on deep learning network. *Renewable Energy*, 127:825–834.