
Proceedings Paper:
Lin, W., Worden, K., Maguire, A.E. et al. (2020) Towards population-based structural health monitoring, part VII: EOV fields – environmental mapping. In: Dilworth, B.J. and Mains, M., (eds.) Topics in Modal Analysis & Testing. Proceedings of the 38th IMAC, A Conference and Exposition on Structural Dynamics 2020. 38th International Modal Analysis Conference, 10-13 Feb 2020, Houston, TX, USA. Conference Proceedings of the Society for Experimental Mechanics Series, 8. Springer, pp. 297-304. ISBN: 9783030477165. ISSN: 2191-5644.

https://doi.org/10.1007/978-3-030-47717-2_31

Towards population-based structural health monitoring, Part VII: EOV Fields: environmental mapping

W. Lin[1], K. Worden[1], A.E. Maguire[2] E.J. Cross[1]

[1]Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK
[2]Vattenfall Research & Development, New Renewables, The Tun Building, Holyrood Road, Edinburgh, EH8 8AE, UK

Abstract

In a population-based structural health monitoring setting, data from one structure in a population, where the health state is known, may be used to make inferences about the health state in any nominally-identical structure. Any deviation from the learned 'healthy response' potentially indicates damage. However, as in standard applications of structural health monitoring, the healthy response from different structures also varies with the changes in environmental conditions across the population. This paper investigates the modelling of the change in environment across a population of structures located in one geographical region, such as a wind turbine farm. A data-driven mapping method (based on Gaussian process regression) will be introduced that aims to quantify and normalise variation stemming from the environment, such that the remaining response is only sensitive to damage or performance anomalies. The way in which environmental maps are constructed and implemented is demonstrated via a case study from an offshore wind farm. The ideas introduced here will constitute a *field* in the framework of population-based structural health monitoring presented elsewhere in the conference.

Key words: Population-Based Structural Health Monitoring (PBSHM), Environmental and Operational Variation (EOV), Gaussian process (GP)

1 Introduction

Population-based structural health monitoring ((PB)SHM) extends the subject of SHM from an individual structure to a group (or population) of structures. One of the motivations behind this is to address the issue caused by the lack of damage state data. Many would agree that supervised learning is one of the most well-suited approaches to damage identification tasks at the higher levels of Rytter's hierarchy, i.e. damage classification, assessment and prediction [1]. To this end, labelled data for various damage cases are crucial in inferring detailed diagnostic information about a structure. However, it is often difficult to obtain these data, especially for high-value structures such as wind turbines and aircraft. A possible solution can be provided by the population-based approach as it allows the damage state data from one structure to aid inference on damage states for all structures in the population [2]. In this paper, the discussion is focussed on the knowledge transfer between nominally-identical structures, referred to as a homogeneous population [3]. An almost ideal example is given by an offshore wind farm, based on which, the analysis in this paper is carried out.

The data collected from operating structures, unlike that from highly-controlled laboratory experiments, are subject to complex environmental and operational variations (EOVs). The influence on measured responses due to EOVs can be comparable to, or sometimes more significant than, that due to damage [4]. Since it is almost impossible to find a damage-sensitive response that is not heavily influenced by EOVs [5], understanding and normalisation of the effect due to EOVs become necessary steps to implement SHM in reality.

Much effort has gone into studying how EOVs may affect the responses of aerospace and civil structures. Some EOVs directly affect how a structure vibrates. This is demonstrated by the flutter or buffeting in long-span bridges as a result of wind-induced vibrations [4, 6]. There are also factors that tend to alter the material properties and boundary conditions of a structure, which then lead to a different dynamic response. For example, [4] summarised that thermal-induced vibrations in bridges may stem mainly from changes in stiffness and boundary conditions (e.g. thermal contraction). In the case of wind turbines, various EOVs exhibit a combined effect on the dynamic responses of components. An example discussed in [7] showed how modal frequencies negatively correlated with temperature (given a rotor speed within a specific range), and how this negative correlation could also be moderated by the wind-induced aerodynamic damping.

When an individual structure is assessed, as in traditional SHM, the changes in EOVs are often mainly thought of as temporal. For example, in the monitoring campaigns of bridges, the EOVs can be considered as the daily or seasonal variations in temperature, wind loading, and traffic loading [6]. However, in PBSHM, the spatial dimension also needs to be considered. As well as the temporal trends, the structures in a population can also be subject to different EOVs based on where they are located.

The additional difficulties in the wind farm case arise from the fact that, the *homogeneous* population of nominally-identical structures, may have significant differences in their responses as a result of their positions within the EOV *field*. An example can be seen in a wind farm; although the turbines are all located in the same geographical region, the wind conditions at different turbines can still differ due to wake effects. This issue adds to the complexity of the problem, and new methods are called for. In particular, this issue means that a more sophisticated approach to the *transfer* of information between the structures may be needed [3, 8]. If a population is *heterogeneous* [9, 10], the problem of *transfer* is compounded. As observed in [9, 8], individuals in a heterogeneous population may differ in terms of *topology*, *geometry* or structure, and all of these characteristics are sensitive to EOVs. For example, geometrical dimensions may change with thermal expansion and material properties can be sensitive to temperature and humidity (as noted above). It is not so obvious how EOVs might induce topology change; however, this is possible, e.g. a bridge may change its topology if an expansion joint closes due to temperature increase, and a new ground node (boundary condition) appears in the representation [9]. In summary, the Venn diagram from [8], on sensitivity of transfer, is modified to that in Figure 1.

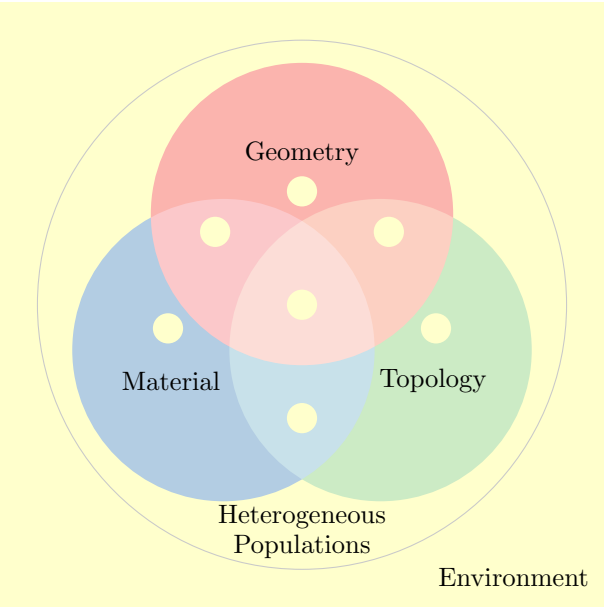

Fig. 1: Categories of heterogeneous population within population-based SHM [8], amended for EOV effects.

The spatial distribution of environmental conditions across a wind farm has been the subject of numerous studies in the field of wind turbine aerodynamics. A large number of computer experiments have been carried out to model the process when turbine rotors extract kinetic energy from the wind, referred to as wake modelling. These physics-based models typically consist of two parts: the first part of the model characterises the aerodynamics of turbine rotors [11, 12, 13], and the second part simulates the entire turbulent wind field and provides the input blade

loads to all rotors in the modelled farm [14, 13]. As a result, the models simulate how much power each turbine can generate, given a specific wind farm configuration, and in other words, how the wind is disturbed due to the existence of operating turbines. However, such models are primarily developed with the purpose of improving the rotor aerodynamic efficiency or optimising wind farm configuration. In this context, a model is usually considered as sufficient so long as it is able to reproduce most of the physical phenomena, i.e. the simulated environment can produce an effect similar to the real wind field over a period of time. In contrast, models with higher accuracies are required for the purpose of SHM, where ideally the model would be able to predict the exact time and location of a physical phenomenon. Such high level in accuracy is difficult to achieve through physics-based models alone (if at all) without a high computational cost. On this account, a data-driven approach to modelling the wind farm environment has been considered here.

This paper aims to understand, through data-based modelling, how EOVs vary temporally and spatially across a population. It is the first step towards understanding the effect of spatially changing EOVs on dynamic responses, after which a new method to remove the EOV-induced trends can be attempted. The work in this paper is based on data from an operating wind farm, Lillgrund, with a brief introduction given in the next section. Section Two will provide a detailed description of the field-mapping approach and the associated data-driven algorithm that are used to model the spatially changing EOVs. The section that follows will present and discuss the results from the environmental mapping model. The paper will finish with conclusions and future implications.

1.1 Lillgrund Wind Farm

The Lillgrund offshore wind farm is located in the middle of the Öresund region between Denmark and Sweden. It consists of 48 turbines rated at 2.3MW, with the farm arrangement seen in Figure 2. The turbines are spaced slightly further away (4.3D) in the prevailing wind direction of 225°-255°compared to the spacing in the perpendicular direction (3.3D). Such spacings are generally considered dense, because the layout was originally designed for smaller turbines [15]. This compact layout was preserved so as to facilitate the investigation into how turbine wakes can affect power production.

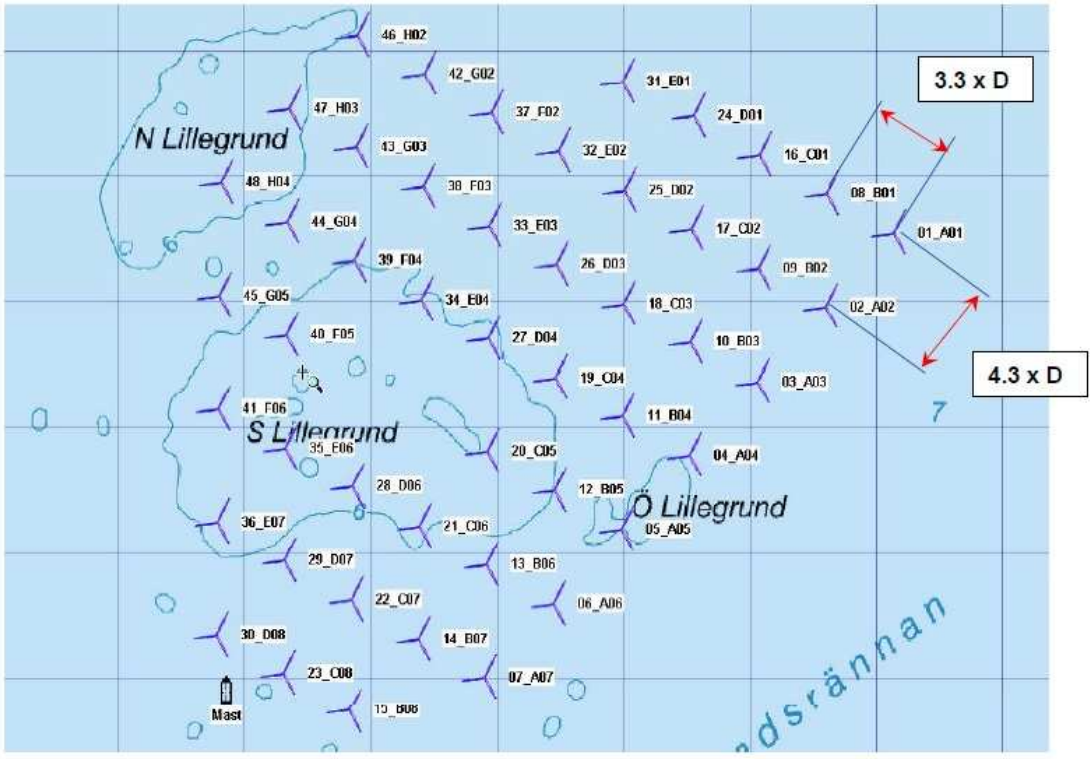

Fig. 2: Turbine arrangement in Lillgrund wind farm [15]

The data used in this paper were collected during a year of operation from the Supervisory Control and Data Acquisition (SCADA) system in Lillgrund. The SCADA data are presented as a statistical summary, including the mean, maximum, minimum, and standard deviation, of the actual recorded data every ten minutes. In the analysis that follows, the ten-minute mean values are used as an indication of the actual measured quantities.

2 Environmental Mapping

2.1 Motivation

According to the unique mechanism in wind turbines, most operational factors are driven by the environmental conditions. In particular, turbine rotor speed is directly proportional to wind speed across a specified operational range, and nacelles should be facing the incoming wind direction by default. The strong correlation between environmental and operational factors indicates that, in the first instance, the environment can be viewed as the dominant driving force behind the variations in an EOV field. As such, the method presented here is called an environmental mapping as it creates a spatial map of the environment across a farm as an indication of the changing EOV field across a population.

As well as being an important step in understanding the effect of EOVs on vibration signals, the environmental map can be applied in two other ways: for SHM purposes, this map can be used as the environmental input to individual turbine models in order to predict the normal, yet different, responses across a farm; another possibility lies in adapting it to a model that optimises the overall power production of the farm by means of wake steering or induction control. Taking into account these potential applications, wind loading becomes the most important aspect in the environment as it has the most direct impact on the structures and the strongest correlation with power. Therefore, the environmental mapping method described here focusses on modelling the spatial variations of wind conditions in a farm.

2.2 Wake Pattern in Wind Farms

Wind conditions vary across an offshore wind farm mainly due to wake effects. The fact that turbines are designed to extract kinetic energy from the wind explains why the wind becomes slower after passing through a turbine rotor. The moving rotor also brings extra turbulence into the wind, resulting in a region behind the rotor with reduced wind speed and intensified turbulence. Further downstream, the wake tends to spread and eventually return to free stream condition, yielding a cone-shaped wake region [16]. In Figure 3a, the different shades of red indicate the variation in characteristics of the wake, which means that the wake effect on any downstream turbine is highly dependent on the relative position of the turbine in shadow with respect to its upstream neighbour. When multiple wakes superimpose on one another, with an example given in Figure 3b, the combined effect needs to be considered.

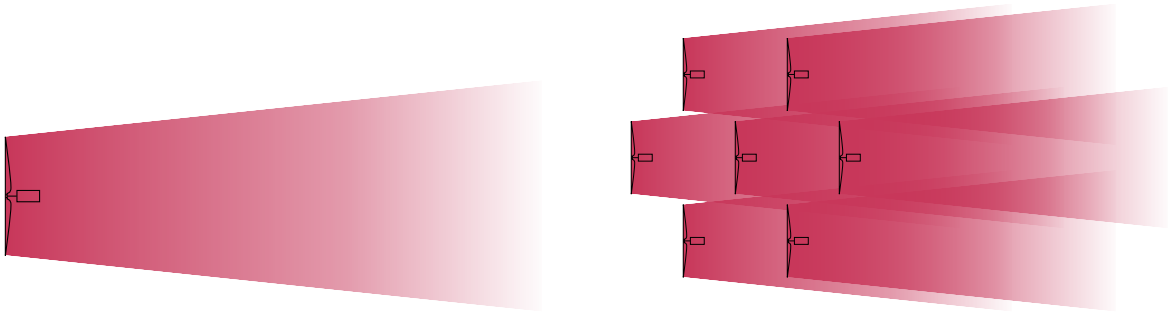

(a) Single turbine wake (top view).

(b) Wake superposition (top view).

Fig. 3: Schematic illustration of turbine wakes, adapted from [16].

It is self-explanatory that the turbine layout and wind direction play important roles in the overall wake pattern, as they affect the position and direction of individual turbine wakes. Given a fixed wind farm layout in this analysis, Figure 4 shows some examples of the wind speed patterns based on different wind directions, visualising the mean

wind speed deficit. The turbulence variations will be accounted for later in the modelling process. It is to be noted that the wind direction at each turbine location is estimated by the nacelle position, given that a good match is found between the wind direction measured on the available weather mast and the nacelle position at the nearest turbine location [15].

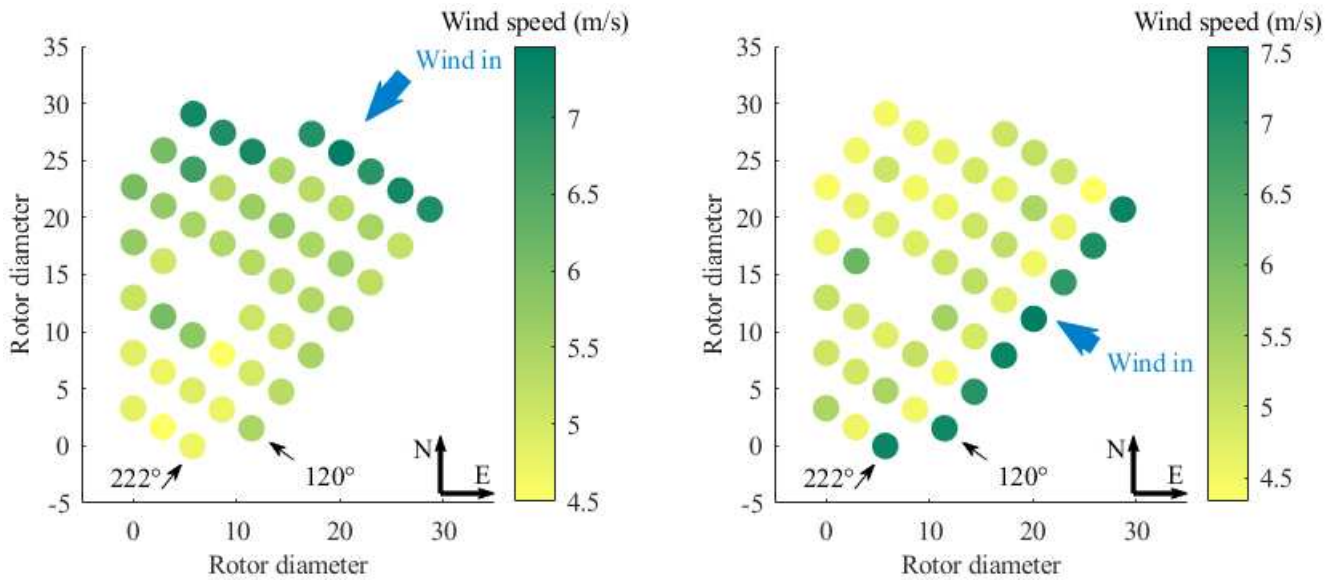

(a) 22°to 62°.

(b) 100°to 140°.

Fig. 4: Mean wind speed distribution across the Lillgrund farm for two incoming wind directions.

2.3 Data-Based Model

To the best of the authors' knowledge, this paper demonstrates the first attempt to build a model that predicts the temporally and spatially-changing EOV fields for SHM purposes. In line with the recent trend in SHM [1], a data-driven approach is preferred for environmental mapping, since it allows for easier integration with other steps in the process of damage identification. Another reason is that, as the physics-based approach has been developed for decades, the current advancement in this field favours the inclusion of more data for better model calibration, namely data-enhanced physics-based models. Given the abundant availability of data in this analysis, an opportunity is given to explore the fully data-based approach in an attempt to achieve relatively high accuracy at a much lower computational cost compared to the physics-based alternatives. As the first step in constructing a reliable model, the model described in this paper is trained to predict only one spatial pattern in wind speed, with the wind coming from 100°to 140°(Figure 4b) specifically.

The relation between wind speeds at various turbine locations is considered nonlinear due to the complex wake effects, especially with regard to turbulence. Gaussian process (GP) regression is used here as it provides an efficient way to learn the nonlinear dependence between continuous variables. GP regression is also a stochastic method that gives a prediction distribution based on which a mean prediction and its associated confidence interval can be computed. The simplest version of the GP is applied here, with a zero mean function and a squared-exponential covariance function. Detailed explanations for GP training and inference are provided by [17].

3 Wind Speed Predictions

The GP model seeks to predict a map of wind conditions by interpolating across the space. That is, it aims to predict the wind speed at all turbine locations, given the wind speed from a fixed subset of the locations as inputs. In the following, the training and testing data sets correspond to different time windows, ranging from 13 to 14 hours' worth of data, that are subjected to similar wind directions (100°to 140°) and wind speeds (4 m/s to 8 m/s). This choice is to make sure that the training and testing sets indicate similar spatial patterns on separate occasions. The training

set is also selected such that the model prediction represents the wind speed variations under normal operational conditions.

It is a standard practice in SHM to use the error of predictive models as an indicator of structural performance. One of the most commonly-used error metrics is the normalised mean-square error (MSE) [6, 18], which is defined as,

$$MSE = \frac{100}{N\sigma_y^2}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{1}$$

where y and $\hat{y}$ represent target data and model prediction respectively, each with a size N. The error is normalised by the variance of the target data, σ_y^2, in order to marginalise the effect due to discrepancy between training and testing data. The scaling factor of 100 ensures that the mean of the data will score 100% if used as the model. This serves as a threshold below which correlation is indicated between prediction and target. On a similar note, error values higher than 100% potentially indicate performance anomalies.

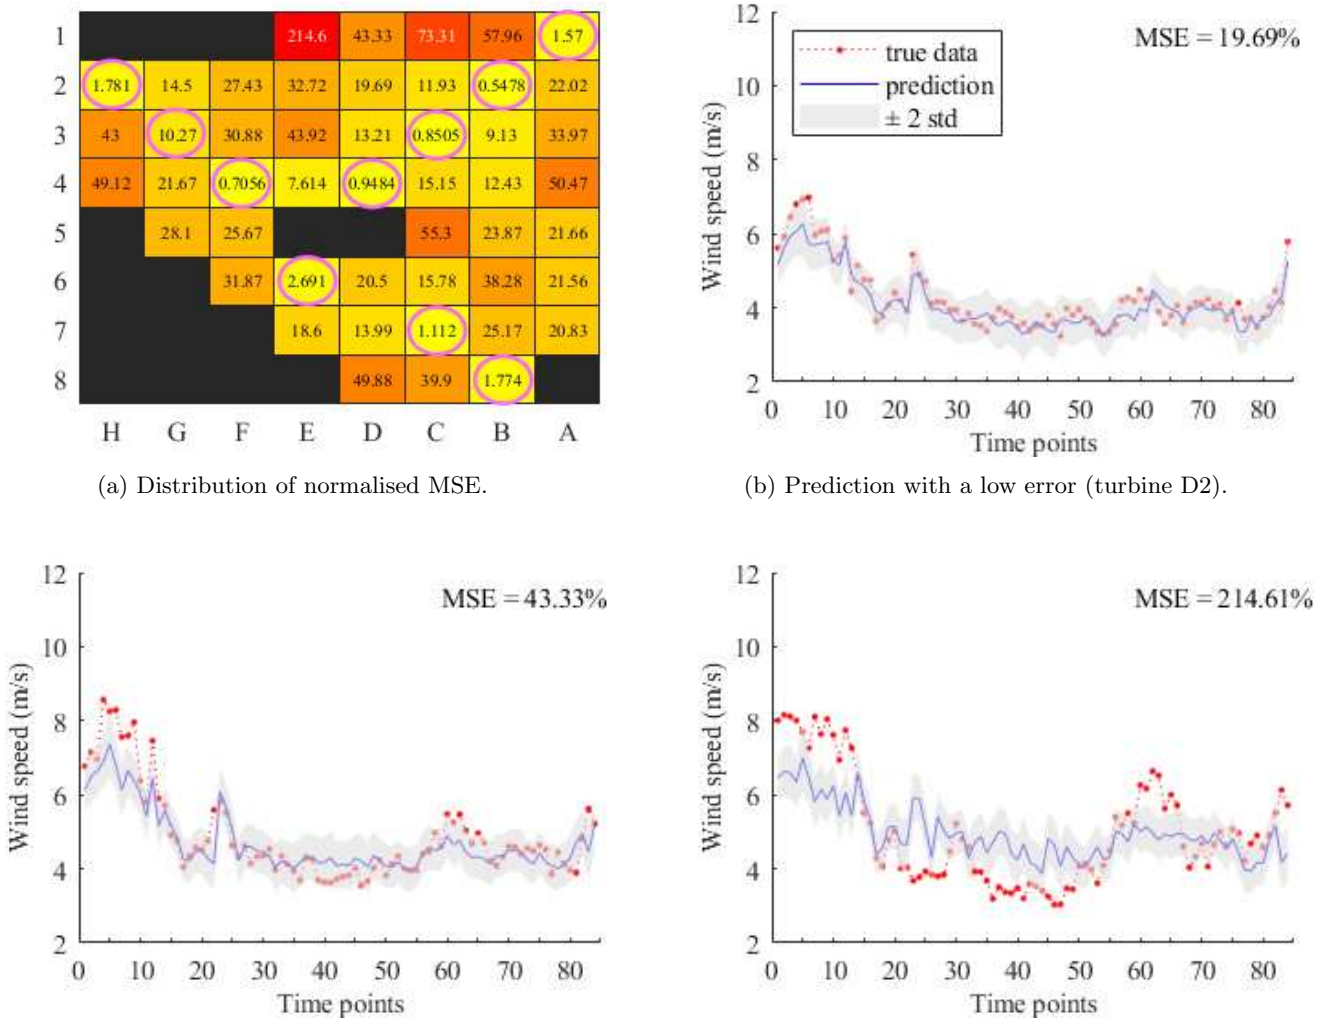

(a) Distribution of normalised MSE.

(b) Prediction with a low error (turbine D2).

(c) Prediction with a medium error (turbine D1).

(d) Prediction with a high error (turbine E1).

Fig. 5: Prediction error distribution, (a), and examples of time series predictions, (b)-(d), for a testing data set.

The error distribution for a testing data set can be seen in Figure 5a. It is worth noting that the heat map is created to reflect on the relative positions between turbines rather than the exact wind farm layout, therefore the map appears to be a rotated and distorted version of the farm. Based on this view, the incoming wind of 100°to 140°appears to proceed from right to left in the heat map. The first thing to be observed was that the normalised MSEs at the ten

reference turbine locations, as indicated by pink circles in Figure 5a, were generally small. This gives an idea of how accurate the GP predictions can possibly be – a mean MSE of about 2.5% was obtained when the model attempted to predict what were given as inputs. For the region where the GP had to interpolate, the majority, except for the highlighted turbine with an error greater than 100%, demonstrated reasonably good predictability, with a mean error of 24%.

In respect of the testing set shown in Figure 5a, examples of time-series predictions with low, medium and high errors are demonstrated in Figures 5b to 5d. In the low-error case (Figure 5b), the predicted line followed the trend of the target data well, with most of the data fluctuations captured by the confidence interval of $\pm$ twice the standard deviation. A short period of slight under-prediction could be seen from time point 0 to 10. The fact that most data points in this period are bounded by the confidence limits means that they are considered by the GP as fluctuations within the acceptable range.

When these fluctuations grew out of the confidence bounds, as seen in Figure 5c, a higher error was assigned by the GP, showing an example of a medium-error prediction. In addition to the under-prediction at the beginning of the time window, the higher error can also be attributed to the over-prediction during time instances 35-50 and the under-prediction during 55-60. Since turbine D2 (Figure 5b) and D1 (Figure 5c) are the same distance away from the first row of turbines encountered by free-stream wind (i.e. D1 and D2 are both in Row D), it is reasonable that predictions with similar trends are made at these two turbine locations. However, such predictions provide a worse fit at D1 than at D2 because they do not take into account how the turbines on the edge of a wind farm may be influenced by the environment from the outside. In this instance, the influence from the outside manifests itself as larger-scale fluctuations in wind speed. This explains the relatively-higher errors at turbine locations along Row 1 (Figure 5a), demonstrating an edge effect.

The most prominent edge effect was found at turbine E1, where the largest-scale fluctuations occurred (Figure 5d). Once again the GP does not expect the fluctuations due to the outside environment, and tends to make a prediction that is similar to that in the neighbourhood (e.g. at turbine location E2). However, in most (if not all) testing data sets that correspond to the same spatial pattern, the highest normalised MSE along Row 1 is found at turbine E1, which suggests a change (in structure or control strategy) at this location. Although the reasons behind it require further investigation, given that very few reference turbines are placed on the edge of the farm and/or that only the information about wind speed is used to train the GP, it is understandable that the model gives a high error to address the unexpected change at turbine E1. To take a step back, putting aside the explanations, it is evident in Figure 5d that the model has the ability to highlight a location, with a prediction error higher than 100%, where an unexpected trend in wind speed occurs.

4 Conclusions

This paper has introduced environmental mapping as a data-based method to understanding the temporally and spatially-changing EOVs for the implementation of PBSHM. The work is based on data from an offshore wind farm, Lillgrund, which exemplifies the concept of homogeneous population. At this stage, an environmental map can be produced by a data-driven model, based on GP regression, which is trained on data from one specific spatial pattern. The GP model has demonstrated its ability to interpolate across the space, obtaining a full map of wind speed on the basis of inputs from a few reference points. Furthermore, if the training data represents the normal wind/operating conditions, the model can also be used to indicate unexpected environmental conditions that might result from different control sequences not present in the training set.

The data-based environmental mapping introduced here can be applied in a manifold manner. Firstly, for the purpose of SHM, environmental mapping can be used as a cost-efficient wind field model for the specific wind farm(s) that the model is trained on. The wind speed predictions can be served as environmental inputs to individual turbine models, allowing them to predict normal structural responses subject to environmental variations. A second opportunity that the authors will be working on is to use a similar mapping approach on power production data, creating a map that aims to establish the relations between power and a number of environmental (e.g. wind speed and direction) and operational (e.g. nacelle position, pitch angle and rotor speed) factors. This new map is then applicable to the optimisation of wind farm power production through wake steering and/or induction control. Similarly, the methodology developed here can be adapted to damage sensitive features (e.g. derived from acceleration data), with the intention of building a more reliable indicator of structural conditions.

Taking a step further, it is also planned to modify the current model to account for a range of various spatial patterns in the environment, aiming for a more generalised model.

Acknowledgements

The authors would like to acknowledge the support of the EPSRC, particularly through grant reference numbers EP/R004900/1, EP/S001565/1 and EP/R003645/1.

References

[1] C.R. Farrar and K. Worden. *Structural Health Monitoring: A Machine Learning Perspective*. John Wiley and Sons Ltd, 2013.

[2] K. Worden, E.J. Cross, N. Dervilis, E. Papatheou, and I. Antoniadou. Structural health monitoring: From structures to systems-of-systems. *IFAC*, 48:1–17, 2015.

[3] L.A. Bull, P.A. Gardner, J. Gosliga, A.E. Maguire, C. Campos, T.J. Rogers, M. Haywood-Alexander, N. Dervilis, E.J. Cross, and K. Worden. Towards population-based structural health monitoring, Part I: Homogeneous populations and forms. In *Proceedings of IMAC XXXVIII – the 38th International Modal Analysis Conference, Houston, TX*, 2020.

[4] H. Sohn. Effects of environmental and operational variability on structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365:539–560, 2007.

[5] K. Worden, C.R. Farrar, G. Manson, and G. Park. The fundamental axioms of structural health monitoring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463:1639–1664, 2007.

[6] E.J. Cross. *On Structural Health Monitoring in Changing Environmental and Operational Conditions*. PhD thesis, Department of Mechanical Engineering, University of Sheffield, 2012.

[7] W. Hu, S. Thöns, R. Rohrmann, S. Said, and W. Rücker. Vibration-based structural health monitoring of a wind turbine system part II: Environmental/operational effects on dynamic properties. *Engineering Structures*, 89:273–290, 2015.

[8] P.A. Gardner and K. Worden. Towards population-based structural health monitoring, Part IV: Heterogeneous populations, matching and transfer. In *Proceedings of IMAC XXXVIII – the 38th International Modal Analysis Conference, Houston, TX*, 2020.

[9] J. Gosliga, P.A. Gardner, L.A. Bull, N. Dervilis, and K. Worden. Towards population-based structural health monitoring, Part II: Heterogeneous populations and structures as graphs. In *Proceedings of IMAC XXXVIII – the 38th International Modal Analysis Conference, Houston, TX*, 2020.

[10] J. Gosliga, P.A. Gardner, L.A. Bull, N. Dervilis, and K. Worden. Towards population-based structural health monitoring, Part III: Graphs, networks and communities. In *Proceedings of IMAC XXXVIII – the 38th International Modal Analysis Conference, Houston, TX*, 2020.

[11] L.J. Vermeer, J.N. Sørensen, and A. Crespo. Wind turbine wake aerodynamics. *Progress in Aerospace Sciences*, 39:467–510, 2003.

[12] A.J. Brand, J. Peinke, and J. Mann. Turbulence and wind turbines. *Journal of Physics: Conference Series*, 318(7), 2011. `https://iopscience.iop.org/issue/1742-6596/318/7`.

[13] F.D. Bianchi, H. de Battista, and R.J. Mantz. *Wind Turbine Control Systems: Principles, Modelling and Gain Scheduling Design*. Springer Verlag, London, 2007.

[14] P.S. Veers. Three-dimensional wind simulation. Technical report, Sandia National Laboratories, Albuquerque ABQ, 1988.

[15] J.-A. Dahlberg. Assessment of the Lillgrund windfarm: Power performance and wake effects. Technical report, Vattenfall Vindkraft AB, 2009.

[16] F. González-Longatt, P. Wall, and V. Terzija. Wake effect in wind farm performance: Steady-state and dynamic behavior. *Renewable Energy*, 39:329–338, 2012.

[17] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge MA, 2006.

[18] K. Worden. *Parametric and Nonparametric Identification of Nonlinearity in Structural Dynamics*. PhD thesis, Heriot-Watt University, Edinburgh, 1989.