

RESEARCH ARTICLE

WILEY

Nonstationarity in correlation matrices for wind turbine SCADA-data

Henrik M. Bette  | Edgar Jungblut | Thomas Guhr

Faculty of Physics, University of Duisburg-Essen, Duisburg, Germany

Correspondence

Henrik M. Bette, Faculty of Physics, University of Duisburg-Essen, Lotharstr. 1, 47057 Duisburg, Germany.
Email: henrik.bette@uni-due.de

Funding information

Bundesministerium für Wirtschaft und Energie, Grant/Award Number: 03EE3016C

Abstract

Modern utility-scale wind turbines are equipped with a Supervisory Control And Data Acquisition (SCADA) system gathering vast amounts of operational data that can be used for analysis to improve operation and maintenance of turbines. We analyze high-frequency SCADA-data from the Thanet offshore wind farm in the United Kingdom and evaluate Pearson correlation matrices for a variety of observables with a moving time window. This renders possible a quantitative assessment of nonstationarity in mutual dependencies of different types of data. We show that a clustering algorithm applied to the correlation matrices reveals distinct correlation structures for different states. Looking first at only one and then at multiple turbines, the main dependence of these states is shown to be on wind speed. This is in accordance with known turbine control systems, which change the behavior of the turbine depending on the available wind speed. We model the boundary wind speeds separating the states based on the clustering solution. Our analysis shows that for high-frequency data, the control mechanisms of a turbine lead to detectable nonstationarity in the correlation matrix. The presented methodology allows accounting for this with an automated preprocessing by sorting new data based on wind speed and comparing it to the respective operational state, thereby taking the nonstationarity into account for an analysis.

KEYWORDS

clustering, correlation matrix, nonstationarity, SCADA-data, wind turbine

1 | INTRODUCTION

Renewable energies are indispensable to respond to the temperature rise caused by climate change. Wind turbines are one key technology paving the way toward a green and emission-less energy production. The Global Wind Energy Council (GWEC) reports a growth of 93 GW in wind installations for 2020, bringing the total up to 650 GW with an increase of the growth rate by 53 % compared to 2019. Offshore installations are on the rise, their comparatively low total increased from 29.1 GW by more than a fifth to 35.3 GW in 2019.¹ The GWEC also forecasts more than 205-GW new offshore wind capacity by 2030 in its 2020 Global Offshore Wind Report.² They point out a growing acceptance of the fact that the price of offshore wind power can out-compete fossil and nuclear fuels. The various reports agree that wind power cannot be the only clean energy used to reach a net zero in emissions by 2050, but onshore and especially offshore wind turbines are crucial to reach this goal.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Wind Energy* published by John Wiley & Sons Ltd.

While offshore locations usually provide steadier and higher wind speeds, this comes at the cost of harsh environmental conditions and increased difficulties for operation and maintenance.³ Maldonado-Correa et al⁴ list various authors claiming that Operation and Maintenance (O&M) costs account for 20–35% of the total expenditure for offshore wind farms. The corresponding numbers are lower but still significant (approx. 10–15 %) for onshore turbines.^{5,6}

Not surprisingly, among the many topics to be studied in the wind turbine field,⁷ O&M is on focus for researchers and industry alike. They undertake increased efforts to effectively optimize O&M procedures for wind turbines and thereby reduce this cost factor.^{8–13} Improved understanding of wind turbine behavior is key to achieving this goal. Data driven methods are developed to control problems such as, for example, yaw misalignment or underperformance.^{14,15} Another prominent topic is the prediction of failures in wind turbines with sufficient lead time to react and carry out preemptive maintenance instead of correctional maintenance. This not only reduces the money lost in turbine downtime but also enables cheaper maintenance. The idea is to optimize assets by replacing components exactly when needed.¹⁶ The wind energy branch follows a general trend in most industries, aiming at moving from scheduled maintenance toward condition-based maintenance to reduce costs and efforts¹⁷

A plethora of data are gathered in modern wind turbines. A Supervisory Control and Data Acquisition (SCADA) system is installed in all major wind farms since commission. This SCADA data contains many variables, usually averaged over 10-min intervals. Some further statistical measures, such as standard deviation in the 10-min interval, are often recorded as well. Many developed methods try to employ them for different types of analysis. The reader is referred to Maldonado-Correa et al⁴ and Tautz-Weinert et al¹⁸ for reviews. Common methodologies include neural networks, physical models and statistical analysis.^{11,19–26} These authors also raise two important points: First, it is often complicated or impossible to reliably label events in the data due to scarcity of available log and maintenance data. Second, Ulmer et al,¹⁰ who apply convolutional neural networks for failure detection, mention that the 10-min averaging process naturally leads to a loss of information. This effect is specifically studied in Beretta et al.²⁷ Some researchers have tried to avoid these problems by using simulated high frequency data^{24,28} while the industry uses additionally installed vibration sensors to increase monitoring quality. Stetco et al²⁹ provide a review on approaches using such Condition Monitoring Systems (CMS). However, the goal is to reduce costs and installing additional sensors has its own inherent costs.

Another challenge for wind turbine analysis and monitoring is presented by the varying external (e.g., wind speed and temperature) and internal (e.g., turbine control and curtailment) operation conditions. Such nonstationarity is important also in many applications aside from wind turbines.^{30–35} We have recently shown that nonstationarity in correlated systems is important for the detection of anomalies.^{36,37} For wind turbines, it has been shown to have an effect on failure detection.^{38,39} Furthermore, different states in frequency data measured by a CMS system have been identified due to operational regimes.⁴⁰ Different behavior of the SCADA data for such regimes is to be expected also due to the turbine control mechanisms.⁸ In general, complex systems containing many different variables, mechanisms, and external influences show nonstationarity in their cross-correlations. The stability of correlations in financial stock market data was analyzed, for example, by Buccheri et al.⁴¹ Münnix et al⁴² show that the correlation matrix of these data inhabits different states over time by using cluster analysis. The stability of these states was further analyzed by Rinn et al⁴³ and Stepanov et al.⁴⁴ Similar studies were also done for traffic data.⁴⁵ Some general correlation analysis for wind turbine data was carried out by Braun et al.⁴⁶

In this paper, we aim to quantify the nonstationarity in correlation matrices for high-frequency SCADA data from real wind turbines during normal operations. To this end, we apply cluster analysis to the correlation matrices of different SCADA signals calculated over 30-min time intervals. Distinct states with significantly different structures in the correlations matrices are found. We show that the prime cause for this is the turbine's own control mechanism. This allows us to develop a criterion based on wind speed to separate the cluster states. Such an automated distinction would, in principle, enable the usage of multiple normal states in applications via preprocessing. It is an important step toward accounting for the nonstationarity due to the operational regimes in an analysis such as failure detection with principal components.

This paper is structured as follows: In Section 2, we will introduce the dataset we work with before moving on to the theoretical background of correlation matrices and clustering in Section 3. Then we will present our clustering results for a single turbine in Section 4. We find proof of nonstationarity and identify the turbine control mechanism as the prime influence. Afterwards, we model the boundary wind speeds between the states in Section 5. Finally, we will show that the established method works for multiple turbines without further problems and can even be improved by the increase in available data in Section 6. We present our conclusions in Section 7.

2 | DATA

Our dataset includes 100 Vestas V90 wind turbines from the Thanet offshore wind farm south-east of Great Britain. It contains observables that are measured at approximately 5-s time intervals. To obtain aligned, synchronized data, we average the time series on 10-s time intervals resulting in a data frequency $\nu = 1/10$ s. This ensures continuity in the data even when the actual measurement frequency fluctuates around 5 s. This does not hinder the calculation of correlations for our purposes. In fact, it is rather similar to any measurement process: Every sensor will in reality average over a short time span to obtain a value. Taking the mean over 10 s seems therefore more natural than, for example, using only the last value from each interval. If at some time the deviation from the 5-s interval becomes stronger or if there are actually missing data points, data points will

be missing in our averaged data as well. This might occur if there was a problem with, for example, sensors or communications. Another reason is simply a decreased measurement frequency when the turbine is switched off. This is underlined by the majority of missing values occurring in the very low wind speed regime beneath the turn-on wind speed. As the turbines are not running during these times, this wind range is not of interest for us. These missing values do not pose a problem for our analysis. For these reasons, we decided to transfer missing values from the original dataset into the 10-s data instead of replacing them by other means.

We are interested in identifying changes in the correlation structure, which emerge while the turbine is operating normally, in contrast to changes caused by failures. Therefore, in our main analysis we look at the following basic observables:

- generated active power (ActivePower)
- generated current (CurrentL1)^{*}
- rotation per minute of the rotor (RotorRPM)
- rotation per minute of the high speed shaft at the generator (GeneratorRPM)
- wind speed (WindSpeed)

These observables provide a good picture of the main turbine systems. Wind speed makes the rotor move. Its rotation is transmitted via gears to the rotation of the high-speed shaft at the generator. This, in turn, generates electrical current and power. Two pairs of strongly correlated variables exist in this set. Deviations between generated active power and current could only occur, if large amounts of reactive power are generated. The low- and high-speed rotation of rotor and generator are directly coupled as well. These expected results are confirmed during our analysis. We include these pairs in our analysis as examples for group structures in the correlations. Knowledge of such structures is indispensable for monitoring a complex system: Are the groups stable? Do they break up? Do correlations across groups exist? In the presence of anomalies, such correlations, which are deemed normal and obvious, might be the structures which break up. Grouping is, in general, an important aspect in the study of complex systems.^{45,47–50}

While measurements of temperatures are very common and useful for failure analysis,^{4,18,19} they are rather slowly changing variables. Hence, in our data they are not measured at high frequency and without decimals, which makes the calculation of short-term correlations impossible. Furthermore, it seems reasonable to assume that their behavior in normal states is strongly coupled to mechanical variables, e.g. higher rotation speeds will lead to increased bearing temperatures. Of course, they would also be influenced by, for example, seasonality or cooling mechanisms. Thus, while excluded from the study at hand, further analysis of temperature correlations is nevertheless desirable for future work.

Two additional important control variables are the pitch angle of the rotor blades (BladePitchAngle[†]) and the ratio between the blade tip speed and the current wind speed (TipSpeedRatio). The first is excluded in our main analysis due to many missing values that hinder the calculation of correlation matrices. The second is not directly present in the data, but results from easy linear calculation

$$\text{TipSpeedRatio} = \frac{2\pi \text{RotorRPM} \cdot \text{RotorRadius}}{\text{WindSpeed}}. \quad (1)$$

It is disregarded in the main analysis, because we study linear correlations and it is also linearly derived from two already present variables. As both omissions are prominent observables when studying wind turbines, we include additional results with consideration of the pitch angle of the blades and the tip speed ratio for the basic cluster analysis in Section 4. To do this, we had to fill the missing values for the pitch angle. A possible explanation for the missing values is a data acquisition system that only writes values whenever a new measurement is different from an old measurement, that is, if the value changed. We could not establish whether this is actually the case in our data, but it seems reasonable. Thus, for the additional results including the pitch angle observable, we treated the data as if this assumption were true. This means we filled any missing values in the 10-s data with the last measured value before. Of course, thereby we also fill in any values that might actually be missing instead of being left out for data storage reasons. Furthermore, with the pitch angle being a rather slowly changing variable compared to the other observables, the filling sometimes leads to stable values over long time periods. This, of course, hinders once again the calculation of short-term correlations. Therefore, one must be cautious when considering these additional results and we did not consider pitch angle or tip speed ratio directly for the analysis following the basic clustering in the present work.

We used approximately 3 weeks of data from March 5, 2017, to March 24, 2017. The data from such a time span are still easy enough to handle while providing enough data points to obtain reasonable clustering results. In view of possible practical applications, 3 weeks is a short enough time span to make it easily usable. There would be no need to collect huge amounts of operational data beforehand. However, it is of course necessary that the data used for identifying different operational states covers all possible states. In practice, it turns out later that this means we need a wide range of wind speeds in our data. The actual time span was chosen, because for at least one wind turbine there are no

^{*}As there are no deviations between the three phases in our data, we simply choose one of them.

[†]Our data do not contain three separate measurements for the blades, but only one. We assume this to be a mean of the three individual blades.

manual or automatic alarms or services during this period (cf. Section 4). Two turbines have no recorded data for this time span, effectively reducing our data set to 98 turbines.

Due to confidentiality agreements, we will never show absolute values of any observable. In fact, only wind speed is shown directly and is then presented in units of the nominal wind speed \tilde{v}_{nom} at which the turbine starts to produce its nominal power output according to the manufacturer. The tilde is introduced to mark it as the rated value provided by the manufacturer as we will later on try to infer this value also from the data.

Each of the measured signals k for each turbine l yields a time series of data points $S_k^{(l)}(t)$, $k = 1, \dots, K$, $l = 1, \dots, L$, $t = 1, \dots, T_{\text{end}}$. Here, we assume that all times are given as unit free steps. In the case of our dataset, we have $K = 5$, $L = 98$ and—assuming complete data— $T_{\text{end}} = 20 \cdot 24 \cdot 60 \cdot 6 = 172800$. The data are arranged into L rectangular $K \times T_{\text{end}}$ data matrices

$$S^{(l)} = \begin{bmatrix} S_1^{(l)}(1) & \dots & S_1^{(l)}(T_{\text{end}}) \\ \vdots & & \vdots \\ S_k^{(l)}(1) & \dots & S_k^{(l)}(T_{\text{end}}) \\ \vdots & & \vdots \\ S_K^{(l)}(1) & \dots & S_K^{(l)}(T_{\text{end}}) \end{bmatrix}, \quad (2)$$

where each row is the time series of signal k .

3 | THEORETICAL BACKGROUND

Our analysis to distinguish different states is based on identifying differences in the correlation matrix of observables listed in Section 2. In Section 3.1 we define the way in which we are calculating the correlation matrices.

To identify nonstationarity in the time series of these matrices we will use a distance measure and a clustering algorithm. These are introduced in Section 3.2.

3.1 | Correlation matrices

To identify changes over time in the correlation structure, the correlation matrices are calculated with a moving time window of 30 min. The time intervals do not overlap. We effectively create a time series of matrices. To this end, the signal time series $S_k^{(l)}(t)$ are divided into disjoint intervals of 30 min; that is, the lengths of the intervals is $T = 180$ in our dimensionless time variable. We refer to these intervals as epochs. Hence, we have $T_{\text{end}}/T = 960$ epochs. To avoid notational confusion, we introduce the new time variable τ labeling the epochs. We reserve the notation $S_k^{(l)}(t)$ for the original time series and write $S^{(l)}(\tau)$ for the $K \times T$ data matrix containing the different time series for turbine l from τ to $\tau + T - 1$. The length of 30 min represents a compromise. Longer time spans would provide more data points per correlation matrix and would thereby decrease noise. However, we want to distinguish different states in time. Considering external conditions, for example, wind, changing on short time scales of several minutes to hours, we have to choose relatively short epochs to ensure resolution of the nonstationarity. Such compromises are common when dealing with correlation matrix time series.^{48,51}

The first step toward calculating correlation matrices is the normalization of $S_k^{(l)}(t)$ to zero mean and standard deviation one in each epoch. With the mean value

$$\mu_k^{(l)}(\tau) = \frac{1}{T} \sum_{t=\tau}^{\tau+T-1} S_k^{(l)}(t) \quad (3)$$

and the standard deviation

$$\sigma_k^{(l)}(\tau) = \sqrt{\frac{1}{T} \sum_{t=\tau}^{\tau+T-1} (S_k^{(l)}(t) - \mu_k^{(l)}(\tau))^2} \quad (4)$$

of the epoch τ , the normalized time series for signal k and turbine l is given by

$$M_k^{(l)}(t) = \frac{S_k^{(l)}(t) - \mu_k^{(l)}(\tau)}{\sigma_k^{(l)}(\tau)}, k = 1, \dots, K, l = 1, \dots, L, \tau \leq t < \tau + T. \quad (5)$$

The normalized $K \times T$ data matrix $M^{(l)}(\tau)$ for each epoch is then defined analogous to $S^{(l)}(\tau)$.

The Pearson correlation matrix for each turbine l during the time $\tau \leq t < \tau + T - 1$ is then easily calculated as

$$C^{(l)}(\tau) = \frac{1}{T} M^{(l)}(\tau) M^{(l)\dagger}(\tau), \quad (6)$$

where $M^{(l)\dagger}(\tau)$ denotes the transpose of $M^{(l)}(\tau)$. In the matrix $C^{(l)}(\tau)$, each element $C_{ij}^{(l)}(\tau)$ is the Pearson correlation coefficient of the signals i and j for turbine l during the epoch from τ to $\tau + T - 1$. The diagonal values are one by definition.

While the dependency of variables in a wind turbine is not always linear, which is already seen in the well-studied power curve, the linear Pearson correlation yields important and good results for the structure of the mutual dependencies. We have repeated our analysis with Spearman's rank correlation, which also measures nonlinear dependencies, but did not find substantial differences. Results for the case with five variables are shown for comparison in Appendix B.

3.2 | Clustering

We will now introduce the clustering, which allows us to sort the correlation matrices into groups (clusters) and check whether different typical states do exist. If we can identify these, we will refer to them as *operational states*. An integer will be assigned to each of them and the algorithm will label each matrix in the time series with one such integer. Instead of a time series of correlation matrices, we then have a new integer time series $n(\tau)$ with the range $n \in \{1, \dots, N\}$ when N is the number of clusters created.

The first outcome will then be as follows: If any decent clustering solution can be found, it is proof that typical states of the correlation structure exist. Then, analyzing the resulting integer time series $n(\tau)$ can much easier reveal dependencies of the state on time or other factors.

Any method separating objects into groups needs a distance measure defined between those objects. For the correlation matrices, we choose the euclidean distance.⁴⁸ The reader can imagine that all matrix entries are written into a vector, effectively arranging the columns of the matrix underneath each other, so that the standard euclidean distance between vectors can be applied. The distance between the correlation matrices for the epochs starting at τ and τ' of turbine l is then

$$d^{(l)}(\tau, \tau') = \sqrt{\sum_{i,j} (C_{ij}^{(l)}(\tau) - C_{ij}^{(l)}(\tau'))^2} = \|C^{(l)}(\tau) - C^{(l)}(\tau')\|. \quad (7)$$

We choose the bisecting k -means algorithm to perform our clustering.^{52,53} This is the algorithm that was also used to determine states in the financial markets by Münnix et al⁴² and Heckens et al.⁴⁸ It can be described as a hybrid of standard k -means clustering^{54,55} and hierarchical clustering.⁵⁶ While the former directly divides the whole set of objects into k groups (see Appendix A), the latter is performed step-wise. In each step, either two groups are merged (agglomerative) or one group is divided into two (divisive). Bisecting k -means is a divisive clustering algorithm, meaning that at the start, all objects belong to one big cluster and during each step, one of the existing clusters is split into two new clusters. This bisection is performed by running a standard k -means on all objects within the group to be split with $k = 2$. Which of the \tilde{N} currently existing clusters is split during a step is decided based on the average internal distance of all objects in a cluster $z_n, n = 1, \dots, \tilde{N}$

$$d_n^{(l)} = \frac{1}{|z_n^{(l)}|} \sum_{\tau \in z_n^{(l)}} \|C^{(l)}(\tau) - \langle C^{(l)} \rangle_n\|. \quad (8)$$

Here, $|z_n^{(l)}|$ denotes the number of objects in cluster $z_n^{(l)}$ and $\langle C^{(l)} \rangle_n$ is the centroid of cluster $z_n^{(l)}$ defined by the element-wise mean:

$$\langle C_{ij}^{(l)} \rangle_n = \frac{1}{|z_n^{(l)}|} \sum_{\tau \in z_n^{(l)}} C_{ij}^{(l)}. \quad (9)$$

Each step the cluster with the largest average internal distance is bisected. The algorithm is terminated when a set number N of clusters is reached.

This is slightly different from the approach used in former works,^{42–44,48} where a threshold is introduced and all clusters are bisected until no single existing cluster has an average internal distance larger than the threshold. However, the threshold is then set based on the number of clusters wanted. It can easily be understood that the resulting clustering will be the same if one either uses our approach to produce N clusters or chooses the threshold in such a way that N clusters are produced.

Applying this algorithm, we split the set $Z^{(l)} = \{C^{(l)}(1), C^{(l)}(1+T), \dots, C^{(l)}(T_{\text{end}} - T)\}$ of all correlation matrices into N subsets $\{z_1^{(l)}, \dots, z_n^{(l)}, \dots, z_N^{(l)}\}$. The centroid of each cluster according to Equation (9) is interpreted as the mean correlation matrix of a cluster representing its typical correlation structure. Thereby, we only need to look at N matrices and a series of T_{end}/T integers $n(\tau)$ instead of as many matrices.

Later on, we will see that the emerging typical correlation matrices correspond to different control settings of the turbines. This explains in a simplified way why the hierarchical k -means works better than a normal k -means in our case. Approximately, we can describe the controller of a wind turbine as a mechanism fixing certain signals to a fixed value. This means the correlation of that signal with other nonfixed signals should vanish. The divisive clustering will first extract a group where signal A might be fixed. Then this group might be further divided into subgroups where signal B is either fixed or not. And in theory, this could go on. Such a problem is very well suited for divisive clustering.

In order to check, if our clustering is sensible, we will do two things. First, we will just look at the cluster centroids and see if we can interpret them and if they are substantially different from each other. Second, we will calculate silhouette coefficients⁵⁷

$$s^{(l)}(\tau) = \begin{cases} \frac{b(\tau) - a(\tau)}{\max(b(\tau), a(\tau))} & , |z_{n(\tau)}| > 1 \\ 0 & , |z_{n(\tau)}| = 1 \end{cases} \quad (10)$$

with the average distance to all other matrices in the same cluster

$$a(\tau) = \frac{1}{|z_{n(\tau)}| - 1} \sum_{\tau' \in z_{n(\tau)}, \tau' \neq \tau} d^{(l)}(\tau, \tau') \quad (11)$$

and the smallest average distance to a single other cluster

$$b(\tau) = \min_{m \neq n(\tau)} \left(\frac{1}{|z_m|} \sum_{\tau' \in z_m} d^{(l)}(\tau, \tau') \right). \quad (12)$$

This coefficient will take values between -1 and $+1$ with larger positive values representing matrices that are well clustered and negative values showing matrices that are closer to another cluster than to their own. To get an indicator for the overall clustering, we will use the average silhouette coefficient

$$S^{(l)} = \frac{T}{T_{\text{end}}} \sum_{\tau} s^{(l)}(\tau). \quad (13)$$

4 | STATE IDENTIFICATION VIA CLUSTERING

In the following, we analyze the correlation matrix time series of one turbine, which will henceforth be referred to as turbine 1 (WT1). We are singling out this wind turbine, because the time frame of the total dataset was selected in such a way that for WT1 no problems were listed in the automatic alarm logs and manual service reports. The idea is that this will make analysis and definition of normal states easier as there was no (reported) unusual behavior.

The correlation matrices are calculated for non-overlapping epochs of 30 min each. This results in 960 matrices per turbine. However, due to several reasons, there might be missing data in the time series. In such a case, any time stamp missing one or more of the measured variables is excluded from the calculation of the correlation matrix. Hereby, no estimation of values, which could influence the actual correlation coefficient, is necessary. More data could be used by calculating the correlation coefficients pair-wise; that is, for any two observables, just remove the time stamps where one of them is missing data. However, this does not result in a well-defined positive semidefinite correlation matrix. For further analysis, only those matrices for which at least half of the expected data points (90 out of 180) exist are considered. Furthermore, epochs in which the standard deviation $\sigma_k(\tau) = 0$ for any signal k have to be disregarded as they cannot be normalized.

We will first provide extensive results for the five considered observables in Section 4.1 and afterwards repeat some analysis including the pitch angle and tip speed ratio in Section 4.2.

4.1 | Main results for five observables

Without pitch angle, no epoch includes a time series k , for which the standard deviation becomes zero. The disregarding of epochs with too many missing values is not a problem when looking at the five main observables as missing measurements usually stem from turbines being operational but switched off during times of very low wind speeds v smaller than turn-on wind speed v_{on} . For WT1, the average wind speed for 746 epochs with enough data is 10.01 ms^{-1} , while the average for 214 epochs where no correlation matrix could be calculated is only 4.34 ms^{-1} . It is obvious that these times where a wind turbine is not operating are unsuited for an identification of operational states. Of course, there might also be other reasons causing the missing data, for example, a problem with the measurement of a signal. However, as for WT1, there are no alarms or services logged, and we would not know what happened in those cases anyway. Any estimation of missing values would therefore need considerable guessing. As our results show that using only the epochs with enough data points is sufficient to reach a good differentiation of operational states, we are confident that just excluding missing values instead of estimating them is a good approach for our purposes.

When applying the hierarchical k -means algorithm described in the previous section to the set of matrices, the first step is to decide how many clusters provide a good solution. To this end, we calculate the silhouette coefficients for solutions with two to five clusters. The silhouette plots can be found in Figure 1. The fifth cluster is almost imperceptible as it consists only of three matrices. Some descriptive statistics for these silhouette coefficients are shown in Table 1. The *mean* corresponds to the average silhouette coefficient from Equation (13). All statistics provided decrease with increasing cluster number, implying that a few different states are sufficient to describe the dynamics of the analyzed correlation matrices. In the plots, we can see some negative coefficients implying elements that would fit better into a different cluster. Such imperfections are to be expected when using heuristics like clustering algorithms. It is however clear that all solutions provide a good grouping with largely positive silhouette coefficients. This is a strong indication that nonstationarities influence the correlation matrix. The influence is strong enough to be detected via simple clustering. Here, we observe that the assumption of a stationary correlation structure, for, for example, principal component analysis, is not justified.




FIGURE 1 Silhouette plots for clustering solutions with two to five clusters. Each clustered element (matrix) is represented by a horizontal line the length of which is the silhouette coefficient for that element. Different clusters are color coded.

TABLE 1 Minimum, first quartile, median, mean, third quartile, and maximum of silhouette coefficients for the clustering solutions with two to five clusters for correlation matrices of WT1.

Clusters	Min	1st Qu.	Median	Mean	3rd Qu.	Max
2	-0.046	0.473	0.571	0.540	0.664	0.734
3	-0.157	0.379	0.537	0.491	0.640	0.718
4	-0.262	0.343	0.508	0.465	0.631	0.716
5	-0.352	0.314	0.479	0.439	0.626	0.707

As mentioned before, we also look at the cluster centroids to see if the matrices show indeed different behavior and if this distinction facilitates clear identification. Figure 2 shows the matrices calculated via Equation (9) in a dendrogram for the hierarchical clustering. The solutions for two and three clusters show distinctly different structures in the matrices, whereas the fourth cluster stems from cluster three but is structurally very similar to cluster one, only differing in the strength of the mean correlation. The introduction of a fifth cluster only produces a very small cluster with only three elements. The algorithm does not identify new groups but rather starts to classify outliers. While the average silhouette coefficient favors two clusters, we continue our analysis with three clusters as we have seen that up to this point structural differences in the matrices occur and we will later see that these can be interpreted very reasonably. Here, we also point out that structural differences in the matrix have a stronger influence on the structure of the eigenvectors, that is, principal components, than differences in average correlation strength. They are therefore more important to distinguish when using methods like principal components and Mahalanobis distance.

The classification of the matrices for three clusters is shown as an integer time series in Figure 3. All three states appear to have a certain stability. Consecutive epochs often belong to the same cluster. However, there is no obvious behavior in dependency of the time. State 3 appears far less often than states 1 and 2. There is no emergence of new or disappearance of old states over time as is sometimes seen in other complex systems.⁴² To get a better idea what each state might represent, we look at the matrices for the cluster centroids calculated according to Equation (9) once more. They are seen in the third row of Figure 2. Generally, as the differences between the matrices are quite clear, we can conclude—in accordance with the silhouette coefficient—that the clustering does indeed separate the matrix time series into meaningful groups. The correlation matrices are nonstationary and automatically separable with a clustering algorithm.

In every cluster, the strongest correlations are clearly visible between the observable pairs ActivePower-CurrentL1 and RotorRPM-GeneratorRPM. This was to be expected as these pairs are very directly linked. Apart from this, we can see that for cluster 1, both of these pairs and the WindSpeed are all correlated with each other. Put differently, higher wind speed leads to faster rotation and thus to higher power. In cluster 2, this changes and the observables RotorRPM and GeneratorRPM, while still being closely correlated with each other, decouple from the other observables. Cross-correlations between these two and any other measurement vanish. The remaining cross-correlation between WindSpeed and the two observables ActivePower and CurrentL1 also vanish in cluster 3. Clearly, the three different states of the correlation structure identified via clustering are meaningful: They show distinctly different correlation behavior between different observables.

To interpret the meaning of the clustering solution, it is helpful to look at turbine control systems. The basic functionality of such a system is, for example, described by Schütt.⁵⁸ The specific functionality varies for individual turbine types, so it is likely that not all turbine types will show the same operational states. The turbine control system of the Vestas V90 turbines analyzed in this paper is one with variable pitch (*Vestas OptiTip™*) and variable speed (*Vestas OptiSpeed™*). We can connect the three clusters to different operational states of this turbine type, which are separated by boundary wind speeds v_{on} , v_1 , v_2 , and v_{nom} . For very low wind speeds just above the turn-on wind speed $v_{on} \leq v < v_1$, the generator rotation is kept constant at the lowest possible value defined by the maximum slip in the generator. This results in a correlation structure as seen in cluster 2. Already for slightly higher wind speeds $v_1 \leq v < v_2$, the system controls the rpm proportional to the wind speed to operate at maximum aerodynamic efficiency of the rotor. This corresponds to cluster 1. With even more wind, but still not enough to reach nominal power output $v_2 \leq v < v_{nom}$, the turbine operates at fixed nominal rpm by controlling the torque. The rotational variables decouple again as for very low wind speeds and the correlation structure corresponds to cluster 2 again. Of course, fluctuations in wind speed will cause the rotation to fluctuate around the nominal value leading to some noise in the correlation structure. If wind speed is high enough to allow full power output $v_{nom} \leq v$, the nominal power output is reached and therefore kept constant alongside the rpm. This results in correlations as seen in cluster 3. All boundary wind speeds v_{on} , v_1 , v_2 , v_{nom} are turbine dependent and usually not public knowledge.

We can see that our reasoning for the turbine at hand is correct by plotting the cluster state over the mean wind speed of the epoch instead of the time stamp in Figure 4. For now disregarding the interval $v_{on} \leq v < v_1$ due to lack of enough data, cluster 1 represents low wind speeds, cluster 2 intermediate wind speeds, and cluster 3 high wind speeds, where nominal power output can be reached. Some exceptions are to be expected due to either wrong sorting during the clustering or wind speeds changing during the 30-min epoch. Another but probably less important factor is the finite response time of a wind turbine controller. It ranges in seconds or minutes and therefore the correlation structure does not respond instantly to fluctuations and changes in wind speed.⁵⁹ One example are the epochs sorted into cluster 3 whose wind speeds seem to lie in the range of cluster 1. They are all moved into the fourth cluster if we take another step in the algorithm. Its centroid is shown in the

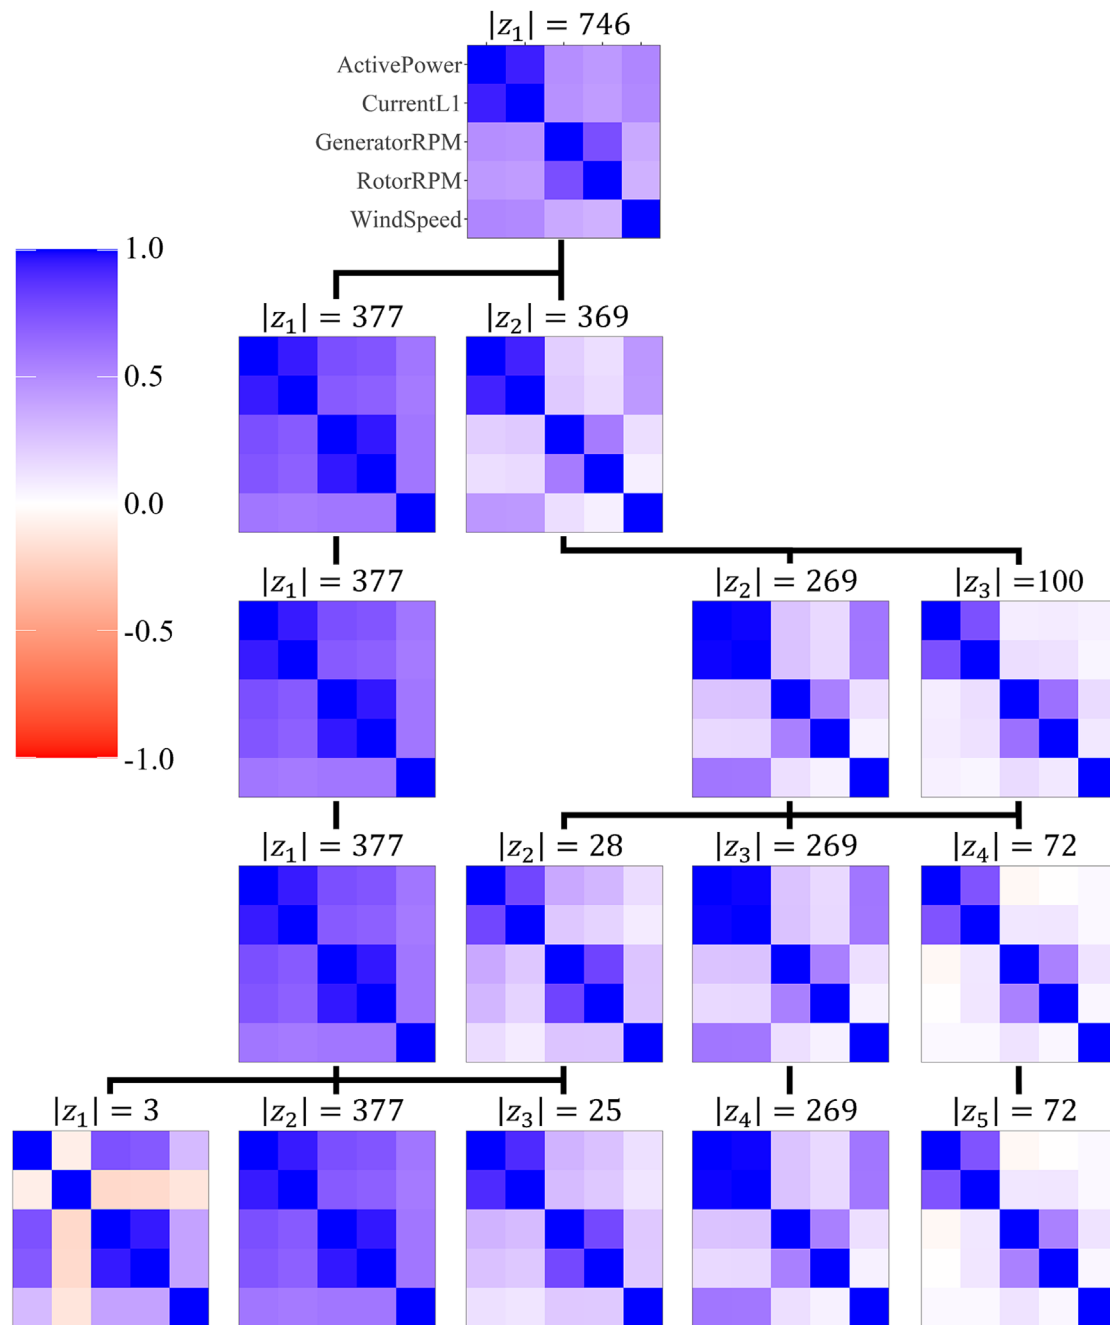


FIGURE 2 Cluster centroids as calculated in Equation (9) for WT1 for different numbers of clusters. The color indicates the value of the correlation coefficient. Black lines connect child and parent clusters of the hierarchical algorithm and the number of cluster elements is given as $|z_n|$. Each cluster solution is ordered from low wind speeds (left) to high wind speeds (right) according to the average wind speed in a cluster.




FIGURE 3 Cluster identifier n over time for WT1. Each dot represents a 30-min epoch.




FIGURE 4 Cluster identifier n over wind speed for WT1. Each dot represents a 30-min epoch.

dendrogram Figure 2 and exhibits a structure very close to cluster 1. Such mismatches occur due to the heuristic nature of the algorithm and noise and fluctuations in the data, which results in matrices lying on the edge between two clusters. We will see in the following section that we can use the silhouette coefficient to identify them. One might also imagine high turbulence intensities, that is, the standard deviation of wind speed divided by the mean wind speed in an epoch, leading to strange behavior in an epoch. We have tested filtering the epochs based on this turbulence intensity and did not find significant changes in the results. The dependency on wind speed is also in accordance with the stability in time as periods with stable wind speeds are common.⁶⁰

We have seen that the correlation matrix is nonstationary in time. The clustering has confirmed a primary influence of the control strategies in dependence of the wind speed. While the existence of different control regimes is not new, our analysis proves that they have strong influence on the structure of the correlation matrix. This automatic separation of states is a vital first step to account for nonstationarities when performing any analysis on high frequency SCADA data.

4.2 | Additional results with pitch angle and tip speed ratio

The inclusion of the tip speed ratio does not affect the number of calculable epochs as it is directly derived from two other observables. Including the pitch angle variable, we only get 623 epochs, for which it is possible to calculate a correlation matrix. This is 123 epochs less than before. As missing values in the time series of pitch angle were filled as described in Section 2, this can only be due to standard deviations in the pitch angle being zero for the pitch angle data. We want to point out that this can be a direct result of the filling mechanism used for missing values. This goes to show that the results with pitch angle while being interesting have to be treated with caution.

The calculation of matrices and the clustering are carried out in exactly the same way as before.

As we cannot assume that the number of relevant clusters stays the same when looking at a different set of observables, we look again at silhouettes in Figure 5 and Table 2 and the cluster dendrogram in Figure 6.

The silhouette coefficients are again largely positive with some expected negative values from imperfect clustering heuristics. On average, the values of the silhouette coefficients are smaller than in the analysis with only five variables. They still indicate a good grouping. The minimum and first quartile even increase in comparison to before, which indicates less poorly sorted matrices. A slight overall decrease in silhouettes is to be expected when clustering larger matrices as more pairs of single correlation coefficients need to be compared and each of them adds fluctuations. Once again, we see decreasing silhouette coefficients for larger numbers of clusters while still indicating that the grouping is reasonable.

Looking at the centroids in Figure 6, we see that with four clusters, the three-cluster solution for with five observables has re-emerged. Only cluster 1 from the previous solution is already split again because of the pitch angle. The numbering of clusters is done based on the average wind speed in each cluster, that is, a low cluster number indicates low wind speeds.

Interestingly, with pitch angle and tip speed ratio considered, we cannot stop at three clusters. The four- and five-cluster solutions still show centroids that are structurally different. The sixth cluster distinguishes stronger and weaker values (mainly in the pitch angle) of the same type of structure and is already quite close to outsider classification with only nine inhabitants. Considering this as well as the sharply falling minimal value of silhouette coefficients from the five-cluster solution to the one with six, we will take a closer look at five clusters. The cluster number over wind speed is shown in Figure 7.

For low wind speeds, the five main observables and the tip speed ratio are always correlated, while the pitch angle is either anticorrelated to all other observables or decouples from them. For the second case, it is likely that it simply stays constant in this regime as the intake of energy from the wind does not need to be reduced here. The anticorrelations might stem from a turbine being turned on underlined by wind speeds being slightly lower for cluster 1 than cluster 2. While it is turned off, large pitch angles are used to minimize strain on a still standing rotor and must then be reduced as wind speed increases above the cut-in point. It is not clear, however, if this effect is strong enough to produce anticorrelations for a period of 30 min. Cluster 3 of the five-cluster solution appears to be an intermediate state, where the rotations are already decoupled from the rest of the system, but the pitch angle is not changed. Tip speed ratio is now anticorrelated to wind speed, active power, and




FIGURE 5 Silhouette plots for clustering solutions with two to six clusters with pitch angle. Each clustered element (matrix) is represented by a horizontal line the length of which is the silhouette coefficient for that element. Different clusters are color coded.

TABLE 2 Minimum, first quartile, median, mean, third quartile, and maximum of silhouette coefficients for the clustering solutions with two to six clusters for correlation matrices of WT1 with pitch angle.

Clusters	Min	1st Qu.	Median	Mean	3rd Qu.	Max
2	0.071	0.521	0.631	0.580	0.672	0.706
3	-0.183	0.380	0.500	0.461	0.609	0.668
4	-0.129	0.239	0.396	0.353	0.490	0.587
5	-0.133	0.270	0.418	0.366	0.495	0.582
6	-0.429	0.243	0.402	0.347	0.479	0.569

current as the rotation—and therefore the tip speed—does not increase any longer. Contrary to the average wind speed sorting, cluster 3 shares a parent cluster with cluster 5 instead of 4. This might be because the decoupling of pitch angle from active power and current is a clearer distinction than the decoupling of the rotations from the rest. As the turbine controller tries to keep rotation constant, it still fluctuates creating some weak correlations, whereas the pitch angle usually stays constant when it is decoupled from everything else. For intermediate wind speeds in cluster 4, the pitch angle is coupled to wind speed, active power, and current. In cluster 5, at high wind speeds, active power and current decouple from wind speed and pitch angle as well. In both states, the pitch angle is used to decrease the intake of power of the turbine. The tip speed ratio behaves as expected from Equation (1).

Figure 7 shows that the clusters are not as clearly separated over wind speed alone as with only five variables. This is mainly true for low wind speeds. The three regimes we identified in our main analysis can again be distinguished. Furthermore, we see in Figure 8 that the small range intermediate cluster 3 can be distinguished from cluster 4 when looking not only at the average wind speed in the epoch but also at the standard deviation of wind speed in the epoch. Cluster 3 exists for small standard deviations. One possible explanation is that the wind changed so little that the controller did not change the pitch angle even though it would already do so in this regime as seen in cluster 4. There could also be a small wind speed regime where the controller already tries to keep rotation constant but does not change the pitch angle to this end. We must also mention that the amount of filled missing values for the pitch angle time series is quite high in cluster 3 as can be seen in Figure 9. This is also true for the overlapping clusters 1 and 2, for which we did not find a clear distinction criterion. The high amount of filled missing values in clusters 2 and 3 could be reasonable as the pitch angle is decoupled from other variables. This would lead to a constant value, which would lead to many missing values in the time series, if the reasoning in Section 2 is correct. However, without knowing for certain that this reasoning is correct, the decoupling of pitch angle from the rest that we find could also be an artifact of the data manipulation.

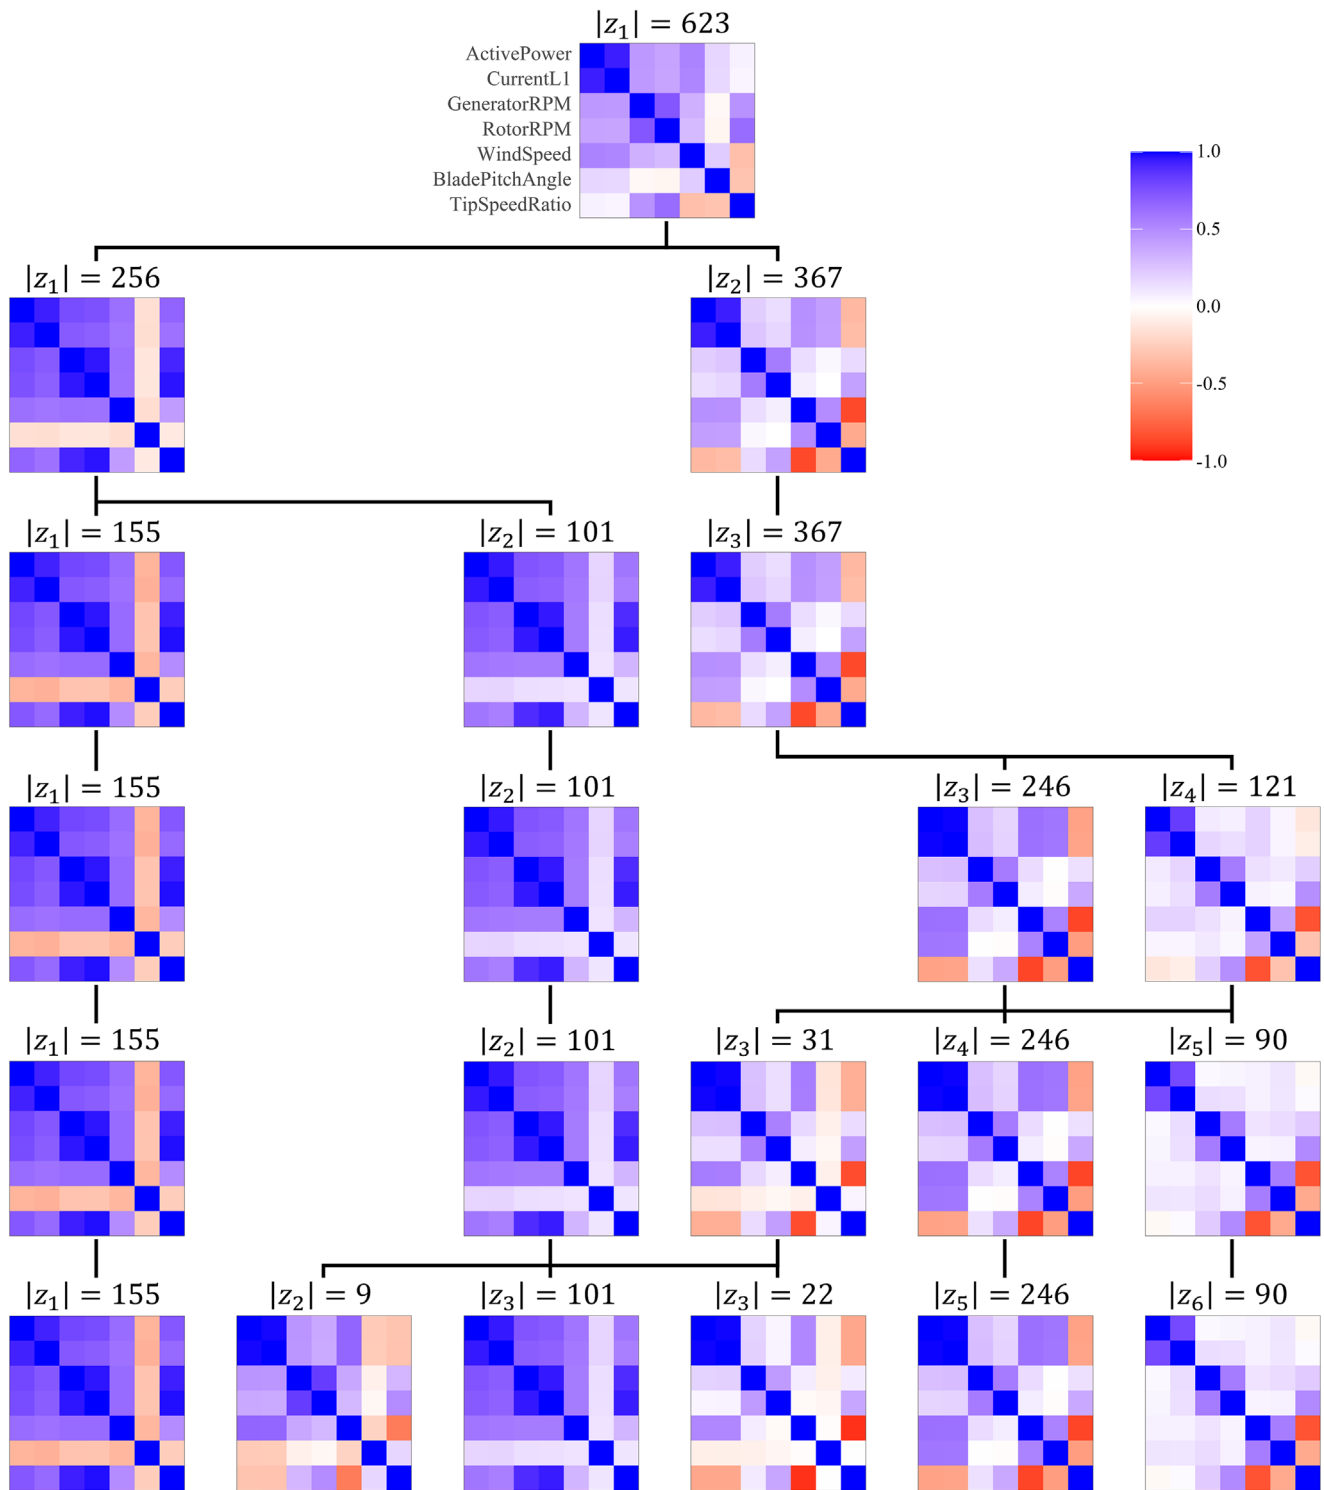


FIGURE 6 Cluster centroids as calculated in Equation (9) for WT1 for different numbers of clusters with pitch angle. The color indicates the value of the correlation coefficient. Black lines connect child and parent clusters of the hierarchical algorithm and the number of cluster elements is given as $|z_n|$. Each cluster solution is ordered from low wind speeds (left) to high wind speeds (right) according to the average wind speed in a cluster.

In conclusion, we have seen that nonstationarity can also be detected with our clustering when including pitch angle and tip speed ratio. The primary influence still stems from the control strategies in dependence of the wind. In general, it can be necessary to have more than one variable, on which the distinction between clusters is based. It could also happen that some overlap cannot be distinguished and more than one normal state would need to be considered for analysis of the system under those operating conditions.

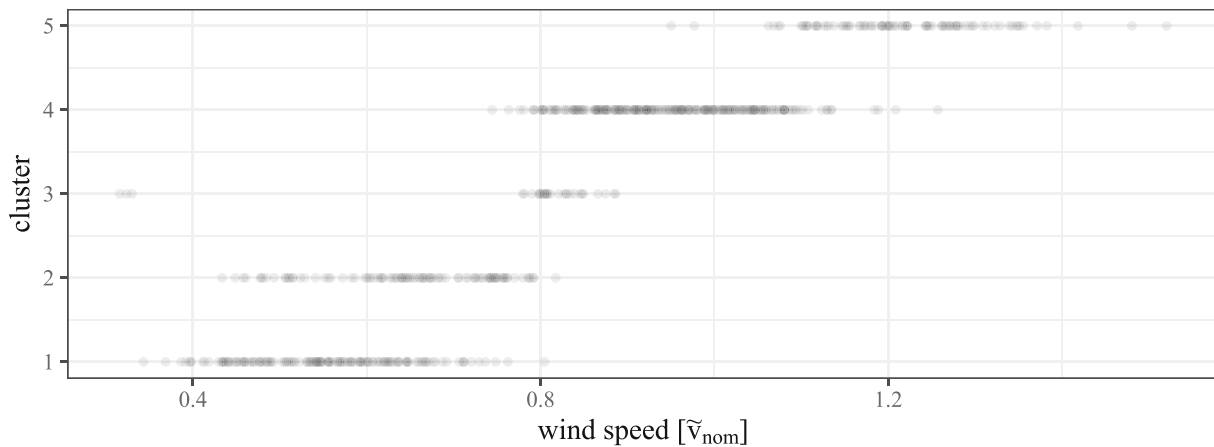


FIGURE 7 Cluster identifier n over wind speed for WT1 and included pitch angle. Each dot represents a 30-min epoch.




FIGURE 8 Cluster allocation as colored points over the average wind speed in an epoch on the x-axis and the standard deviation of wind speed during an epoch on the y-axis.

Further analysis of matrices with more and different variables as well as an attempt to distinguish other influencing factors is interesting for future work. Regulatory impacts on the correlation structure as, for example, curtailment should also be considered. We will continue the current work with an analysis of the possibility to predict the state solely based on wind speed for our five main observables.

5 | CLUSTER PREDICTION BY WIND SPEED

Having established a strong influence of the control system on the structure of the analyzed correlation matrices, we will now try to predict which correlation matrix state, that is, operational state, the turbine should be in based on the wind speed. We have seen in the analysis with additional variables that overlaps between states can happen when differentiating solely based on wind speed. In such cases, more variables or external parameters might be necessary for state prediction. For the current work, we will stick with the simpler example of five observables. Here, we are confident when separating the clusters solely based on wind speed and show that such a distinction can, in general, work. This makes the example easy to follow. Also, this way we do not need to worry about the filled in pitch values. We now present a method that allows separation of the




FIGURE 9 Cluster allocation as colored points over the average wind speed in an epoch on the x-axis. The y-axis shows the proportion of filled missing values in the pitch angle time series relative to all values in an epoch.




FIGURE 10 Probability density functions for the 30-min epoch mean wind speed per cluster. The width of bins is $0.05\tilde{v}_{nom}$.

three states found based on wind speed. To this end, we look at the distribution of wind speeds in the different states, analyze them and then predict the boundary wind speeds that separate the distinct groups. This will show that it is possible to account for found nonstationarity, even though small adaptations are likely to be necessary when considering different sets of observables or turbines. In Figure 10, one can see the empirical probability density functions (pdfs) for wind speeds per cluster state. As already expected from Figure 4, we can clearly distinguish the different regimes. However, we identify much more clearly what we are calling mismatches: Epochs that are sorted into cluster 3 but have mean wind speeds in the range associated with cluster 1. They make up the left peak of the distribution for cluster 3. Furthermore, we can see a small peak in the pdf for cluster 2 lying at very low wind speeds beneath the distribution for cluster 1. These could be reasonable as the rotation of the generator shaft is kept at a minimum rotational speed needed for operation of the turbine for very low wind speeds as discussed in the previous section. However, the data in the very low wind regime is sparse and not as reliable as the turbines often move in and out of operation during these times due to shutting off below a certain minimal wind speed, therefore we will disregard this first boundary v_1 for now.

Before modeling the boundaries between the distributions, we try to compensate for mismatches due to matrices lying at the edge of two clusters, matrices being wrongly sorted, or singular outliers. This can easily be done by using the silhouette coefficient we have introduced before. It gives an indication of how good a member of a cluster fits into this cluster compared to the other clusters. This means that any 30-min epoch

that has been sorted into cluster i but should rather be in cluster j will have a very small or negative silhouette coefficient. We can use this fact and remove from the calculation of the pdf all epochs with a silhouette coefficient below the first quartile of all silhouette coefficients, which can be seen in Table 1. The resulting pdfs can be seen in Figure 11. The second peak at low wind speeds for cluster 3 disappears. This indicates that our reasoning of a mismatch was correct. The persistence of the small peak at very low wind speeds for cluster 2 on the other hand shows that it indeed points toward a control of the rotational speeds in this regime.

The empirical distributions are noisy due to the finite amount of data points. This is especially true for cluster 3, which contains the least epochs. However, it is very clear that every cluster is representing a wind speed regime. There are now two basic approaches to defining the boundaries between these regimes. One can simply look at the empirical data and define for each value of wind speed the maximum likelihood state based on the empirical pdf. Second, one can fit a distribution to the data and calculate the intersections of these, which represent the boundaries. For now, we choose to fit distributions as it turns out that a normal distribution is a good choice for each wind speed regime (see Figure 11) and the other method could be heavily influenced by noisiness in the empirical data. Of course, this will dismiss the smaller peak of cluster 2. It should be taken into account if enough data exists in this regime (cf. Section 6). For now, cluster 1 has a mean of $0.603\tilde{v}_{\text{nom}}$ and a standard deviation of $0.101\tilde{v}_{\text{nom}}$. Clusters 2 and 3 are centered at $0.943\tilde{v}_{\text{nom}}$ and $1.255\tilde{v}_{\text{nom}}$ with standard deviations of $0.101\tilde{v}_{\text{nom}}$ and $0.071\tilde{v}_{\text{nom}}$, respectively. These values lead to boundaries at $v_2 = 0.774\tilde{v}_{\text{nom}}$ and $v_{\text{nom}} = 1.118\tilde{v}_{\text{nom}}$. Interestingly, the last value shows that v_{nom} as calculated in our analysis is larger than \tilde{v}_{nom} . The reason for this discrepancy lies in realistic operational conditions. The power curve (active power in dependency of wind speed) as given by the manufacturer is one line. Accordingly, there is exactly one value \tilde{v}_{nom} that marks the starting point for nominal power production. In reality, especially when looking at high-frequency data, there will always be an area around this line which is realized. The value \tilde{v}_{nom} lies in the middle of this smeared out power curve. At this wind speed nominal power output can be reached but is not yet constant. With even higher wind speeds, it becomes less and less likely that the actual power produced lies beneath the nominal value. Only when this probability nearly vanishes, a change in correlation structure is detectable by our method. It is therefore reasonable that our value v_{nom} lies higher than \tilde{v}_{nom} . While our value is therefore well suited to distinguish correlation states, it cannot be directly compared with the nominal wind speed given by the manufacturer of a turbine. We have confirmed this by looking at scatter plots of our data but cannot show them in this paper due to data confidentiality.

We want to point out two things. First, when using our method as a pre-processing for an analysis it needs to be run on the same observables that are to be considered in the analysis (cf. Section 4.2). Some adaptations for additional influences from the external conditions, regulatory influences (e.g., curtailment and deration) or overlaps of clusters might be necessary. Second, it is not proven that a normal distribution will always provide the best fit. For example, if much larger wind speeds exist in the data set, the distribution for cluster 3 would have much heavier tails in the large wind speed regime unless cut off by the cut-out wind speed of the turbine.

In the current case for WT1, the model fitted works very well. Compared to the clustering solution, we have 9.9% of changes in group assignment. If one only looks at epochs with silhouettes above the first quartile, this number reduces to 3.7%. This is clear as the epochs previously characterized as mismatches obviously change their cluster allocation when applying the model to all epochs. The mean correlation matrices of the states as sorted by the model are shown in Figure 12. They clearly exhibit the different structures discussed above produced by the control system of the turbine showing that our state prediction works. This state prediction is an essential first step for using the different operational regimes as a preprocessing for data analysis. Using the fitted criterion, one can predict what state the turbine should be in and run the analysis for the corresponding operational regime. This is necessary if the analysis itself (e.g., principal components) directly involves the correlation matrix.




FIGURE 11 Probability density functions for the 30-min epoch mean wind speed per cluster using only epochs with a silhouette coefficient above the first quartile of all silhouette coefficients. For each cluster, a normal distribution was fitted. Then black vertical lines indicate the intersections of these distributions and thereby the boundary wind speeds. The width of bins is $0.05\tilde{v}_{\text{nom}}$.




FIGURE 12 Matrices corresponding to the group centroids after sorting with the epochs according to the calculated boundaries for WT1. The mean matrices were calculated for all epochs, not only those used to determine the boundaries. The color indicates the value of the correlation coefficient.

Otherwise, it is also possible to make direct use of the clustering and simply identify the state by comparing the current correlation matrix to the previously identified cluster centers. A direct application of both that we want to test in the future is monitoring the correlation structure. Using the fitted criterion we predict a correlation matrix and compare it to the current one. With sufficiently labeled data we plan to analyze if deviations might signify operational problems or failures.

As we have seen in our analysis with additional variables in Section 4.2 for low wind speeds, large overlaps between multiple states can occur when differentiating by wind speed. If this is the case in an analysis at hand, one could look for other distinguishing factors. However, it is not a given that these exist. Another possibility is to accept that more than one control state is normal for the given conditions and compare new data to all possible states. If for example the goal is to minimize false alarms in a failure detection procedure, one could run failure analysis in all likely states and then choose the one that gives the least indication for failure. An alternative could be weighing the failure indicators with the likeliness calculated for each state under the given conditions.

6 | APPLICATION TO MULTIPLE TURBINES

For a single wind turbine, we identified different operational states in the correlation structure and presented a model to distinguish these states based on wind speed. To be useful for applications, our findings need to be general characteristics and not be specific for one turbine. We proceed and test our methodology for all turbines in the data set. We want to emphasize that in a first step this does not mean assuming one model with fixed wind speed boundaries and applying it to all turbines. Rather, we test if the procedure described in previous chapters can be automatically transferred to other turbines without supervision. Hence, we perform cluster analysis and fit the boundaries for each turbine separately.

An easily comparable indicator for the quality of the proposed methodology is the relative numbers of cluster allocation changes from the model compared to the clustering itself. We already discussed that for WT1 at the end of the previous section. This number will drastically increase if either of the two steps in the calculation does not work well on a turbine: If the clustering algorithm returns a solution that is not grouped by wind speed, sorting on the basis of wind speed will change the allocation of many epochs. If they are clustered according to wind speed, but the boundaries are less sharp than for WT1, it will again result in many changed allocations.

A histogram over all calculated allocation changes for the 98 turbines is shown in Figure 13. We can see that the changes for WT1 lie in the lower end of changes, as expected due to it not showing any alarms or failures in the chosen time span. This does not hold true for the other turbines, most of which exhibit a few more changes. However, there are multiple turbines which show no more changes than WT1. Also, for those that do the fitted boundaries still work remarkably well. Some allocation changes are always expected for the reasons discussed above and especially near the boundaries the distinction between two states is not always perfect. Concluding that our proposed method works well for all turbines in our dataset, we proceed with an optimization.

It stands to reason that the pdfs for wind speeds per cluster should be much smoother, if we look at all turbines at the same time. Combining the data from all turbines before fitting the model, we have approximately 98 times more data points than for a single turbine. The resulting distributions are shown in Figure 14. They are indeed much smoother. Furthermore, we can see that the previously assumed Gaussian fit would not work well anymore. Especially the distribution shown for cluster 1 is skewed. Also, the peak in the distribution for cluster 2 beneath wind speeds of $0.4\tilde{v}_{nom}$ becomes more explicit. This second point is explained in Sections 4 and 5 and actually underlines our reasoning: For very low wind speeds, rotation is kept constant and with more data from all turbines we can distinguish this regime more clearly than before.

At a first glance, the skewness of the pdfs does not fit our theory so well. If the controller would work perfectly and instantly and the wind speeds were constant during each epoch, we would expect the distinctions between the operational states to be represented by rectangular




FIGURE 13 Histogram counts showing the frequency of relative changes in state allocation when comparing clustering and individual models per turbine.




FIGURE 14 Probability density functions for the 30-min epoch mean wind speed per cluster considering all turbines. Only those epochs with a silhouette coefficient above the first quartile of all silhouette coefficients for each turbine were used (calculated separately per turbine). The underlying wind speed distribution without cluster separation is shown as inlay. The width of bins is $0.02\tilde{v}_{nom}$.

functions as probability densities. In nonperfect conditions, these would overlap and smooth out to be something similar to a Gaussian distribution as assumed before, but they should not become skewed. However, the reason can be found in the underlying distribution of wind speeds in the inlay in Figure 14. It follows roughly the expected Weibull distribution. The deviations could be explained by combinations of influences in the environment of the wind farm, overlaying of different Weibull distributions for different wind directions or maybe even measurement effects due to the sensor being placed behind the rotor. Some differences might also be introduced by the removal of low silhouette coefficients as these will appear often in the regimes of the boundaries between states where two correlation structures might be mixed during an epoch. This is of interest for future studies. For now, we can take away that the skewness of this underlying distributions might lead to the skewness in the cluster pdfs. To check this, we replace the histogram count of epochs $h_i(v_w)$ for wind speed v_w and cluster i with the rescaled count

$$\tilde{h}_i(v_w) = \frac{h_i(v_w)}{h_{total}(v_w)}$$

by dividing with the total histogram count of epochs $h_{total}(v_w)$ for that wind speed and all clusters. This basically removes the effect of the underlying wind speed distribution by transforming it into an equal distribution. The resulting pdfs for each cluster can be seen in Figure 15. Indeed, we can now see symmetric areas, showing behavior very similar to rectangular functions for cluster 1 and 3 with cluster 2 being smoothed out to a more Gaussian curve, because its wind regime is quite narrowly bounded by overlaps with the other states. This strongly underlines the existence of three regimes corresponding to operational states of the turbines.

As the functions for all turbines are much smoother and the bin size can be reduced, we can apply the direct maximum likelihood method instead of fitting a continuous curve to decide cluster allocation based on the epoch wind speed. This leads to three instead of the previous two boundary wind speeds to account for the appearance of cluster 2 in the very low wind regime. The values of the boundaries are $v_1 = 0.38\tilde{v}_{nom}$, $v_2 = 0.80\tilde{v}_{nom}$, and $v_{nom} = 1.10\tilde{v}_{nom}$. The resulting histogram of changes due to model allocation compared to the clustering is presented in




FIGURE 15 Probability density functions for the 30-min epoch mean wind speed per cluster considering all turbines after dividing by the total number of counts per wind speeds to rescale the underlying distribution to an equal one. Only those epochs with a silhouette coefficient above the first quartile of all silhouette coefficients for each turbine were used (calculated separately per turbine). The width of bins is $0.02\tilde{v}_{nom}$.




FIGURE 16 Histogram counts showing the frequency of relative changes in state allocation per turbine when comparing the clustering solution and the maximum likelihood model based on data from all turbines.

Figure 16 and shows less changes compared to Figure 13. One reason for this is the taking into account of the very low wind speed regime in cluster 2. Such a method without need for fitting could be easily transferred to other data and turbine types providing high usability as a preprocessing step for data analysis. It will need to be tested in future work, how best to deal with wind speeds where the clusters overlap. The simplest method proposed above is an all-or-nothing approach choosing the likeliest cluster. Contrary to that, if one wants to minimize false alarms in a failure detection for example, it could be useful to compare current data to all clusters, which are possible for the current wind speed and choose the one indicating the least likelihood for a failure. Additional filters alongside the current wind speed can also be necessary as seen in Section 4.2, where the standard deviation of wind speed helped separating clusters. Also, operating measures such as curtailment might lead to temporary changes in the boundaries between clusters and thereby create the need for an additional filter. In general, when applying this method, one should always check the clustering results beforehand.

Overall, we conclude that the results formerly shown for WT1 are easily transferred to multiple wind turbines. Furthermore, the model to decide state allocation based on wind speed can be optimized by taking into account more turbines and thereby more data.

7 | CONCLUSION

Using a matrix distance measure and clustering algorithm formerly applied to other complex systems, we were able to identify different operational states in 30-min correlation matrices of high frequency wind turbine data without prior knowledge of the control system. This demonstrates the nonstationarity of the correlation matrix for wind turbines and its automated detectability. While the states quite often exhibit stability over a certain time period, the real dependency lies with wind speed. This is expected for wind turbine control regimes. In the analysis with additional variables, the standard deviation of wind speed during a 30-min epoch was also shown to have an influence. Furthermore, it was possible to model the boundary wind speeds between the different states for the main analysis of five observables—again without knowledge about the

actual parameters used in the control system. This allowed us to recreate the cluster allocation solely based on the 30-min average wind speeds. Being developed on one turbine, the method is transferred easily to multiple turbines. Results were improved by this increased amount of data. Our study shows clearly that the control system causes detectable nonstationarity in the correlation structure of high frequency wind turbine SCADA data. The automatic separation of states is important to account for this nonstationarity when analyzing such data, for example to monitor a turbine during operation.

While it is of course known that the control system of the turbine changes its operational behavior based on the external influences, our analysis proves that the influence on the correlation structure of the SCADA data is significant and an automatic distinction based on the correlation matrix is possible. Therefore, assuming a stationary correlation matrix, e.g. when applying principal component analysis for dimensionality reduction on a dataset, is unjustified.

Furthermore, it could potentially be important for monitoring with high frequency SCADA-data, e.g. when applying failure detection. Especially methods directly dependent on the correlation matrices such as principal components⁶¹ and Mahalanobis distance⁶² might benefit from the definitions of multiple, distinct normal states in the correlation behavior as they usually assume stationarity. They are commonly applied to wind turbines,^{11,19,24-26,63} see also the reviews mentioned in the introduction. We have recently shown that for generic correlated systems with distinct normal states, the knowledge of these states increases the sensitivity of change detection based on principal components.^{36,37} This is possible via pre-processing: Using a criterion based on historical data—wind speed in the presented case—new or live data could be compared to the respective operational states. Charmingly, the proposed ansatz does not require changes in established techniques, it just requires their application to multiple subgroups and is therefore easily implemented.

As nonstationarity has been shown to influence failure analysis for wind turbines,^{38,39} we aim to test the benefit of our method for this use-case in future studies using SCADA-data with labeled failure events and encourage others to do so. Handling of cluster allocation at the boundaries between clusters, where overlaps in the probability distribution exist, and abnormal operating conditions such as curtailment should then be studied as well. Additionally to using the proposed pre-processing for failure analysis methods, we intend to test also a more direct application: Predicting an expected correlation matrix based on a fitted criterion and comparing it to the current one. Subsequently, deviations might point toward operational problems. Labeled data will also allow training of the model based on all intervals without abnormalities instead of a continuous time span thereby increasing the amount of available data per turbine. Furthermore, it will then be possible to study long-term nonstationarity indicated by Jin et al.¹⁹

As another obvious extension of the current work, we hope to analyze the nonstationarity for other wind turbine types and different observables.

AUTHOR CONTRIBUTIONS

Thomas Guhr initiated the research. Henrik M. Bette and Thomas Guhr defined the objectives. Henrik M. Bette developed the tools and conducted the data analysis with support from Edgar Jungblut. Henrik M. Bette wrote the first draft of the manuscript. Edgar Jungblut and Thomas Guhr contributed to the writing of the final version.

ACKNOWLEDGEMENTS

We are most grateful to *Vattenfall AB* for providing the data. We acknowledge fruitful conversations with Joachim Peinke, Matthias Wächter, Christian Phillipp, Timo Lichtenstein, and Anton Heckens. This study was carried out in the project *Wind farm virtual Site Assistant for O&M decision support - advanced methods for big data analysis (WiSAbigdata)* funded by the Federal Ministry of Economics Affairs and Energy, Germany (BMWi). One of us (H.M.B.) thanks for financial support in this project. Open Access funding enabled and organized by Projekt DEAL.

CONFLICT OF INTEREST STATEMENT

The authors declare no potential conflict of interests.

PEER REVIEW

The peer review history for this article is available at <https://www.webofscience.com/api/gateway/wos/peer-review/10.1002/we.2843>.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from *Vattenfall AB*. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors with the permission of *Vattenfall AB*.

ORCID

Henrik M. Bette  <https://orcid.org/0000-0002-8644-5398>

REFERENCES

1. Lee J, Zhao F. Global Wind Report 2021. tech. rep., Global Wind Energy Council; 2021. <https://gwec.net/global-wind-report-2021/>
2. Lee J, Zhao F. Global Offshore Wind: Annual Market Report 2020. tech. rep., Global Wind Energy Council; 2020. <https://gwec.net/global-offshore-wind-report-2020/>
3. Feuchtwang J, Infield D. Offshore wind turbine maintenance access: a closed-form probabilistic method for calculating delays caused by sea-state. *Wind Energy*. 2013;16(7):1049-1066.
4. Maldonado-Correa J, Martín-Martínez S, Artigao E, Gómez-Lázaro E. Using SCADA data for wind turbine condition monitoring: a systematic literature review. *Energies*. 2020;13(12):3132.
5. Zhao Y, Li D, Dong A, Kang D, Lv Q, Shang L. Fault prediction and diagnosis of wind turbine generators using SCADA data. *Energies*. 2017;10(8):1210.
6. Tautz P. *Offshore Wind Turbines: Reliability, Availability and Maintenance*: Institution of Engineering and Technology; 2012.
7. Veers P, Dykes K, Lantz E, et al. Grand challenges in the science of wind energy. *Science*. 2019;366(6464):eaau2027.
8. Menezes EJM, Araújo AM, da Silva NSB. A review on wind turbine control and its associated methods. *J Cleaner Prod*. 2018;174:945-953.
9. Dao CD, Kazemtabrizi B, Crabtree CJ, Tavner PJ. Integrated condition-based maintenance modelling and optimisation for offshore wind turbines. *Wind Energy*. 2021;24(11):1180-1198.
10. Ulmer M, Jarlskog E, Pizzi G, Manninen J, Huber LG. Early fault detection based on wind turbine SCADA data using convolutional neural networks. *PHM Soc Europ Conf*. 2020;5(1):9.
11. Yan H, Mu H, Yi X, Yang Y, Chen G. Fault diagnosis of wind turbine based on PCA and GSA-SVM. In: 2019 Prognostics and System Health Management Conference (PHM-Paris); 2019:13-17.
12. Martín-del Campo S, Al-Kahwati K. Unsupervised ranking of outliers in wind turbines via isolation forest with dictionary learning. *PHM Soc Europ Conf*. 2020;5(1):9.
13. Chen W, Qiu Y, Feng Y, Li Y, Kusiak A. Diagnosis of wind turbine faults with transfer learning algorithms. *Renew Energy*. 2021;163:2053-2067.
14. Pandit RK, Infield D. SCADA-based wind turbine anomaly detection using Gaussian process models for wind turbine condition monitoring purposes. *IET Renew Power Gen*. 2018;12(11):1249-1255.
15. Meyer A, Brodbeck B. Data-driven performance fault detection in commercial wind turbines. *PHM Soc Europ Conf*. 2020;5(1):7.
16. Colone L, Reeder M, Dimitrov N, Straub D. Assessing the utility of early warning systems for detecting failures in major wind turbine components. *J Phys: Conf Ser*. 2018;1037(3):32005.
17. Li Y, Peng S, Li Y, Jiang W. A review of condition-based maintenance: its prognostic and operational aspects. *Front Eng Manag*. 2020;7(3):323-334.
18. Tautz-Weinert J, Watson SJ. Using SCADA data for wind turbine condition monitoring—a review. *IET Renew Power Gen*. 2017;11(4):382-394.
19. Jin X, Xu Z, Qiao W. Condition monitoring of wind turbine generators using SCADA data analysis. *IEEE Trans Sustain Energy*. 2021;12(1):202-210.
20. Kolumbán S, Kapodistria S, Nooraee N. Short and long-term wind turbine power output prediction. <https://arxiv.org/abs/1707.06497>; 2017.
21. Gottschall J, Peinck J. How to improve the estimation of power curves for wind turbines. *Environ Res Lett*. 2008;3(1):15005.
22. Aziz U, Charbonnier S, Bérenguer C, Lebranchu A, Prevost F. Critical comparison of power-based wind turbine fault-detection methods using a realistic framework for SCADA data simulation. *Renew Sustain Energy Rev*. 2021;144:110961.
23. Helbing G, Ritter M. Deep learning for fault detection in wind turbines. *Renew Sustain Energy Rev*. 2018;98:189-198.
24. Pozo F, Vidal Y, Salgado O. Wind turbine condition monitoring strategy through multiway PCA and multivariate inference. *Energies*. 2018;11(4):749.
25. Vidal Y, Pozo F, Tutivén C. Wind turbine multi-fault detection and classification based on SCADA data. *Energies*. 2018;11(11):3018.
26. Pliego Marugán A, García Márquez FP. Advanced analytics for detection and diagnosis of false alarms and faults: a real case study. *Wind Energy*. 2019;22(11):1622-1635.
27. Beretta M, Pelka K, Cusidó J, Lichtenstein T. Quantification of the information loss resulting from temporal aggregation of wind turbine operating data. *Appl Sci*. 2021;11(17):8065.
28. Odgaard PF, Johnson KE. Wind turbine fault detection and fault tolerant control—an enhanced benchmark challenge. In: 2013 American Control Conference. IEEE; 2013:4447-4452.
29. Stetco A, Dinmohammadi F, Zhao X, et al. Machine learning methods for wind turbine condition monitoring: a review. *Renew Energy*. 2019;133:620-635.
30. Kazemi P, Giralt J, Bengoa C, Masoumian A, Steyer J-P. Fault detection and diagnosis in water resource recovery facilities using incremental PCA. *Water Sci Technol*. 2020;82:2711-2724. <http://iwaponline.com/wst/article-pdf/82/12/2711/803450/wst082122711.pdfbyguest>
31. Shang L, Liu J, Turksoy K, Shao QM, Cinar A. Stable recursive canonical variate state space modeling for time-varying processes. *Control Eng Pract*. 2015;36:113-119.
32. Canonaco G, Restelli M, Roveri M. Model-free non-stationarity detection and adaptation in reinforcement learning. *Front Artif Intell Appl*. 2020;325:1047-1054.
33. de Mesquita CPB, White CT, Farrer EC, Hallett LM, Suding KN. Taking climate change into account: non-stationarity in climate drivers of ecological response. *J Ecology*. 2021;109:1491-1500. <https://onlinelibrary.wiley.com/doi/full/10.1111/1365-2745.13572>
34. Chen G, Lu G, Liu J, Yan P. An integrated framework for statistical change detection in running status of industrial machinery under transient conditions. *ISA Trans*. 2019;94:294-306. <https://doi.org/10.1016/j.isatra.2019.03.026>
35. Huang J, Sun X, Yang X, Peng K. Fault detection for chemical processes based on non-stationarity sensitive cointegration analysis. *ISA Trans*. 2022;129:321-333.
36. Bette HM, Guhr T. Sensitivity of principal components to changes in the presence of non-stationarity. <https://arxiv.org/abs/2208.04651>; 2022.
37. Tveten M. Which principal components are most sensitive in the change detection problem? *Stat*. 2019;8(1):e252. <https://doi.org/10.1002/sta4.252>
38. Zimroz R, Bartelmus W, Barszcz T, Urbanek J. Diagnostics of bearings in presence of strong operating conditions non-stationarity—a procedure of load-dependent features processing with application to wind turbine bearings. *Mech Syst Signal Process*. 2014;46:16-27.
39. Bull T, Ulriksen MD, Tcherniak D. The effect of environmental and operational variabilities on damage detection in wind turbine blades. In: 9th European Workshop on Structural Health Monitoring; 2018:84. <https://www.ndt.net/article/ewshm2018/papers/0084-Bull.pdf>
40. Avendano-Valencia LD, Barahona B, Chatzi E, Hoelzl C. Operational regime clustering for the construction of PCE-based surrogates of operational wind turbines. In: Proceedings of the 7th International Conference on Advances in Experimental Structural Engineering; 2017:961-964. <https://www.researchgate.net/publication/319880225>

41. Buccheri G, Marmi S, Mantegna RN. Evolution of correlation structure of industrial indices of US equity markets. *Phys Rev E*. 2013;88(1):12806.
42. Münnich MC, Shimada T, Schäfer R, et al. Identifying states of a financial market. *Scientif Rep*. 2012;2:644.
43. Rinn P, Stepanov Y, Peinke J, Guhr T, Schäfer R. Dynamics of quasi-stationary systems: finance as an example. *Europhys Lett*. 2015;110(6):68003.
44. Stepanov Y, Rinn P, Guhr T, Peinke J, Schäfer R. Stability and hierarchy of quasi-stationary states: financial markets as an example. *J Stat Mech: Theory Experiment*. 2015;2015(8):P08011.
45. Wang S, Gartzke S, Schreckenber M, Guhr T. Quasi-stationary states in temporal correlations for traffic systems: Cologne orbital motorway as an example. *J Stat Mech: Theory Experiment*. 2020;2020(10):103404.
46. Braun T, Waechter M, Peinke J, Guhr T. Correlated power time series of individual wind turbines: a data driven model approach. *J Renew Sustain Energy*. 2020;12(2):23301.
47. Newman MEJ. Communities, modules and large-scale structure in networks. *Nature Phys*. 2012;8:25-31. <https://www.nature.com/naturephysics>
48. Heckens AJ, Krause SM, Guhr T. Uncovering the dynamics of correlation structures relative to the collective market motion. *J Stat Mech: Theory Experiment*. 2020;2020(10):103402.
49. Patankar SP, Kim JZ, Pasqualetti F, Bassett DS. Path-dependent connectivity, not modularity, consistently predicts controllability of structural brain networks. *Netw Neurosci*. 2020;4(4):1091-1121.
50. Toju H, Yamamoto S, Tanabe AS, Hayakawa T, Ishii HS. Network modules and hubs in plant-root fungal biomes. *J Royal Soc Interf*. 2016;13(116):20151097.
51. Marti G, Nielsen F, Bińkowski M, Donnat P. A review of two decades of correlations, hierarchies, networks and clustering in financial markets. In: Nielsen F, ed. *Progress in Information Geometry. Signals and Communication Technology*: Springer Cham; 2021:245-274.
52. Steinbach M, Karypis G, Kumar V. A comparison of document clustering techniques. In: *KDD Workshop on Text Mining*; 2000.
53. Tan P-N, Steinbach M, Karpatne A, Kumar V. *Introduction to Data Mining*. 2nd ed.: Pearson Education; 2019.
54. Lloyd SP. Least Squares Quantization in PCM. *IEEE Trans Inform Theory*. 1982;28(2):129-137.
55. Jain AK. Data clustering: 50 years beyond K-means. *Pattern Recogn Lett*. 2010;31(8):651-666.
56. Kaufman L, Rousseeuw PJ. *Finding Groups in Data, Wiley Series in Probability and Statistics*. Hoboken, NJ, USA: John Wiley & Sons, Inc.; 1990.
57. Rousseeuw PJ. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math*. 1987;20(C):53-65.
58. Schütt RJ. Control of wind energy systems. *Understanding Wind Power Technology: Theory, Deployment and Optimisation*, Vol. 9781118647: Wiley Blackwell; 2014:340-368.
59. Behnken C, Wächter M, Peinke J. Multipoint reconstruction of wind speeds. *Wind Energy Sci*. 2020;5(3):1211-1223.
60. Weber J, Reyers M, Beck C, et al. Wind power persistence characterized by superstatistics. *Scientif Rep*. 2019;9(1):19971.
61. Denis DJ. Principal components analysis. *Applied Univariate, Bivariate, and Multivariate Statistics*. 2nd ed. Hoboken: John Wiley & Sons, Ltd.; 2021: 423-448.
62. Maesschalck RD, Jouan-Rimbaud D, Massart DL. The Mahalanobis distance. *Chemometr Intell Lab Syst*. 2000;50:1-18.
63. de la Hermosa González-Carrato RR. Wind farm monitoring using Mahalanobis distance and fuzzy clustering. *Renew Energy*. 2018;123:526-540.
64. Gartzke S, Wang S, Guhr T, Schreckenber M. Spatial correlation analysis of traffic flow on parallel motorways in Germany. *Phys A: Stat Mech Appl*. 2022;599:127367.

How to cite this article: Bette HM, Jungblut E, Guhr T. Nonstationarity in correlation matrices for wind turbine SCADA-data. *Wind Energy*. 2023;26(8):826-849. doi:[10.1002/we.2843](https://doi.org/10.1002/we.2843)

APPENDIX A: STANDARD K-MEANS CLUSTERING

The standard k -means algorithm sorts every object $O_i, i = 1, \dots, N$ from a set of N objects into k subsets $\{z_1, \dots, z_1, \dots, z_k\}$. Every subset is called a cluster. Generally, the optimal number k has to be determined by a separate method.^{55,56} The input for the algorithm are the objects $O_i, i = 1, \dots, N$ and a distance measure $d(O_i, O_j) \geq 0, i, j = 1, \dots, N$ as well as a method to compute the centroid of any cluster z_l . Note that the distance measure must also be defined for the centroids. Then the algorithm works as follows:

1. Select k objects as starting cluster centroids.
2. Assign every object to the nearest cluster based on the distance from the object to the cluster centroid.
3. Calculate the new cluster centroids.
4. Repeat steps 2 and 3 until no allocation changes occur.

In this paper, the objects are the correlation matrices of the subset we wish to split into two in the hierarchical approach, therefore $k = 2$ always. The cluster centroids are calculated according to Equation (9) and the distance is calculated by Equation (7).

APPENDIX B: CLUSTERING RESULTS WITH SPEARMAN'S RANK CORRELATIONS

The results presented in this study are based on the Pearson correlation coefficient that has been proven useful to establish structural features in complex systems.^{42,45,48,64} However, some of the variables used in our analysis, such as wind speed and active power have a well-known nonlinear dependency. We have therefore tested the robustness of our results, when applying the nonlinear measure of Spearman's rank correlation coefficient.

To calculate the time series of Spearman correlation matrices, we rank the individual time series $S_k^{(l)}(t)$, $k = 1, \dots, K$, $l = 1, \dots, L$, $t = \tau, \dots, \tau + T - 1$ for signal k of turbine l for one epoch. Then we proceed by calculating the Pearson correlation matrices for the ranked time series. The following clustering procedure and analysis is carried out in exactly the same way as for the results in Section 4. We present here results for the case of five variables, which can be directly compared to the results with Pearson correlation coefficients in Section 4.1.

The silhouette coefficients shown in Figure B1 and Table B1 indicate good grouping. They show on average marginally larger values than for the Pearson correlation. Comparing the resulting cluster centers for Spearman correlations in Figure B2 with those for simple Pearson correlations in Figure 2, it is obvious that the differences for the structural features revealed in this analysis are minimal. The number of elements changes slightly for some clusters, but the overall result and interpretation are the same for both correlation measures.

As expected from the similarity of the results, also the plots of the cluster allocation over time in Figure B3 and over wind speed in Figure B4 are very similar to their counterparts in Section 4.1.

We conclude that for the analysis carried out in this study, the simple Pearson correlation measure is sufficient. The structural differences in the correlation matrices and thereby also the structural differences in the eigenvectors, that is, principal components, are well captured in the linear correlations.

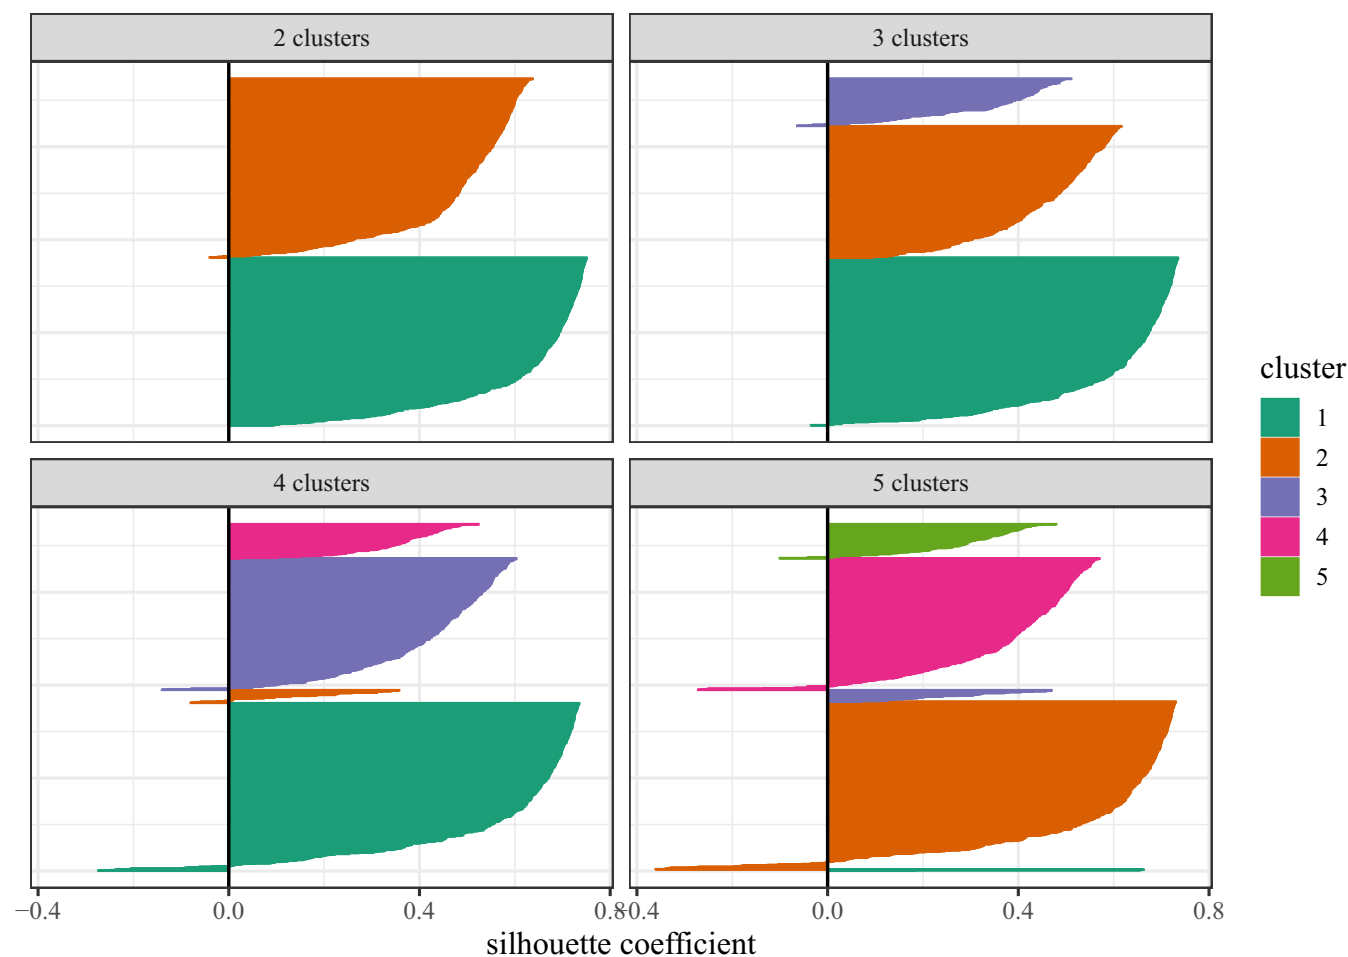


FIGURE B1 Silhouette plots for clustering solutions with two to five clusters for Spearman's rank correlation matrices. Each clustered element (matrix) is represented by a horizontal line the length of which is the silhouette coefficient for that element. Different clusters are color coded.

TABLE B1 Minimum, first quartile, median, mean, third quartile and maximum of silhouette coefficients for the clustering solutions with two to five clusters for Spearman ranked correlation matrices of WT1.

Clusters	Min	1st Qu.	Median	Mean	3rd Qu.	Max
2	-0.040	0.479	0.579	0.550	0.677	0.751
3	-0.064	0.406	0.533	0.507	0.656	0.735
4	-0.274	0.365	0.510	0.476	0.649	0.735
5	-0.361	0.325	0.483	0.450	0.641	0.730




FIGURE B2 Spearman's rank correlation matrix cluster centroids as calculated in Equation (9) for WT1 for different numbers of clusters. The color indicates the value of the correlation coefficient. Black lines connect child and parent clusters of the hierarchical algorithm and the number of cluster elements is given as $|z_n|$. Each cluster solution is ordered from low wind speeds (left) to high wind speeds (right) according to the average wind speed in a cluster.




FIGURE B3 Cluster identifier n over time for WT1 and Spearman's rank correlations. Each dot represents a 30-min epoch.




FIGURE B4 Cluster identifier n over wind speed for WT1 and Spearman's rank correlations. Each dot represents a 30-min epoch.

DuEPublico

Duisburg-Essen Publications online

UNIVERSITÄT
DUISBURG
ESSEN

Offen im Denken

ub

universitäts
bibliothek

This text is made available via DuEPublico, the institutional repository of the University of Duisburg-Essen. This version may eventually differ from another version distributed by a commercial publisher.

DOI: 10.1002/we.2843

URN: urn:nbn:de:hbz:465-20240222-164803-4



This work may be used under a Creative Commons Attribution 4.0 License (CC BY 4.0).