Context-Aware Performance Benchmarking of a Fleet of Industrial Assets

Alessandro Murgia, Elena Tsiporkova, Mathias Verbeke and Tom Tourwé

Abstract Industrial assets are instrumented with sensors, connected and continuously monitored. The collected data, generally in form of time-series, is used for corrective and preventive maintenance. More advanced exploitation of this data for very diverse purposes, e.g. identifying underperformance, operational optimization or predictive maintenance, is currently an active area of research. The general methods used to analyze the time-series lead to models that are either too simple to be used in complex operational contexts or too difficult to be generalized to the whole fleet due to their asset-specific nature. Therefore, we have conceived an alternative methodology allowing to better characterize the operational context of an asset and quantify the impact on its performance. The proposed methodology allows to benchmark and profile fleet assets in a context-aware fashion, is applicable in multiple domains (even without ground truth). The methodology is evaluated on real-world data coming from a fleet of wind turbines and compared to the standard approach used in the domain. We also illustrate how the asset performance (in terms of energy

Alessandro Murgia · Elena Tsiporkova · Mathias Verbeke · Tom Tourwé
Sirris, Brussel, Belgium
✉ Alessandro.Murgia@sirris.be
✉ Elena.Tsiporkova@sirris.be
✉ Mathias.Verbeke@sirris.be
✉ Tom.Tourwe@sirris.be

ARCHIVES OF DATA SCIENCE, SERIES A
(ONLINE FIRST)
KIT SCIENTIFIC PUBLISHING
Vol. 5, No. 1, 2018

DOI: 10.5445/KSP/1000087327/17
ISSN 2363-9881

production) is influenced by the operational context (in terms of environmental conditions). Moreover, we investigate how the same operational context impacts the performance of the different assets in the fleet and how groups of similarly behaving assets can be determined.

1 Introduction

Increasing global competition and compulsory cost reductions force industries to develop models for managing their fleet of assets in an optimal manner. These models rely on the stream of data generated by sensorized assets to identify problems such as underperformance and degradation. This data, generally recorded as time-series, is rich of useful information for fleet optimization. However, its full exploitation is hampered due to several shortcomings. The first shortcoming is due to the fact that generally used methods struggle to properly describe how asset performance changes along with the operational context. On the one hand, methods based on simple bivariate analysis of asset performance (e.g. power curve (Lydia et al, 2014)) cannot handle operational contexts where the performance depends on multiple factors. On the other hand, advanced data mining and machine learning methods for exploiting time-series extract features that are time-dependent and consequently cannot abstract the real operational context. As a consequence, the resulting models are asset-specific and difficult to generalize to the whole fleet. A second shortcoming comes from the lack of ground truth. Indeed, in many real-world cases, there is no explicit indication of the exact periods of abnormal behavior of the asset, which makes the adoption of traditional supervised learning methods not possible. A third shortcoming depends on the continuous evolution of the fleet where assets are constantly added, removed and/or replaced. In this context, methods for optimization of fleet performance are fragile since they might require continuous updates to cope with the new fleet configuration.

In this paper, we propose a methodology to address the previous shortcomings. The methodology characterizes the operational context of an asset and quantifies its impact on asset performance. The operational contexts of the asset are extracted by exploring the data (e.g. associated with exogenous factors) in order to derive relatively homogeneous windows of operating conditions. We define these multi-dimensional windows as “hypercubes”. Subsequently, diverse performance indicators of the assets are extracted for each hypercube. These indicators can then be used for benchmarking and profiling. In both cases, assets are compared based on their performance across the different operational contexts. Finally, thanks to the identification of similarly behaving assets, our methodology provides the building blocks for modeling a fleet in a scalable manner that is also more resilient to fleet changes. Our methodology supports explorative analysis and, being unsupervised, can be applied in multiple real-world domains where there is no labeled data (no ground truth).

The performance of the methodology is evaluated on real-world data from the wind energy domain. We describe the current state of the art in this domain and we show that our methodology is better for describing the connection between asset performance (energy production) and operational context (environmental conditions). We highlight how to find clusters of wind turbines that behave similarly through the different operational contexts and how the same operational context impacts differently the performance of the assets.

2 State of the Art

This section describes how time-series are exploited in the literature, which models for fleet optimization are more resilient to changes and which are design concepts for a data warehouse. Time-series provided by sensorized assets can be exploited for diagnostics, and prediction of failures and performance analysis (Bell and Foslien, 2005; Gorinevsky et al, 2002; Kim and Mylaraswamy, 2006). The data stream generated by the assets can be analyzed in the frequency or time domain. In the frequency domain, generally the analysis focuses on the power spectrum of the assets, or the vibration frequencies or acoustic emission of its sub-system (Gong and Qiao, 2015; Wang and McFadden, 1996; Gong and Qiao, 2013; Iyer et al, 2013). In the time domain, the analysis generally is based on tracking the trends of sensor values (mean, range, variance, correlation, etc.) in order to detect abnormal behavior (Sanchez et al, 2015; Blesa et al, 2015; Freire et al, 2013). In this domain, the raw time-series are exploited with different levels of re-engineering. On the one hand, the data is exploited to build models based on linear regression, Support Vector Regression (SVR), kNN regression, etc. In this case, despite the existence of advanced data mining and machine learning methods for exploiting time-series, the features used in these models cannot abstract the real operational context that influences asset performance. Indeed, these features are entangled with time-dependent information such as the hour, day, week, or month in which the data was gathered. On the other hand, the raw data is used as-is to draw simple operative curves. In this curve, the dependent variable describes the performance of the asset (e.g. yield); whereas the independent variable changes according to the domain (e.g wind speed for the wind domain, irradiation for the photovoltaic domain, etc.). In the wind turbine domain, the commonly used operative curve is the power curve. The power curve is a two-dimensional representation that correlates the wind speed with the energy production of the turbine. The power curve, defined following the specific protocols (IEC, 2008), is used to characterize the performance of a wind turbine without specific knowledge of the assets and its sub-parts (Gill et al, 2012). Generally, manufacturers use power curves as part of the technical specification of the wind turbine and as a guarantee of its performance whereas researchers use power curves to reveal performance degradation and failures in the wind turbine (Jia et al, 2016; Uluyol et al, 2011; Gill et al, 2012; Lapira et al, 2012). Being a simple two-dimensional plot (performance versus wind speed), the power curve cannot be used to properly characterize (or compare) the asset performance in cases of strong dependencies on an operational context described by multiple factors. For instance, having two wind turbines with different power curves may not imply that one of them is not operating correctly or underperforming. Indeed, beyond wind speed there are other factors, like temperature and wind direction, that need to be considered for performance comparison.

For fleet optimization it is necessary to adopt methods that are perfectly tuned to handle the different configurations of assets. This is complex to achieve in practice since assets are continuously added, updated or removed which consequently leads to a continuous update of the fleet model. To make this model more resilient to asset changes it is possible to use an integrative modeling. This approach leverages on scalable models (e.g. prototypical model of assets with similar behavior) and encompasses relations among these models. The rationale behind integrative modeling is to decompose iteratively a complex system in (simpler) sub-systems. Sub-systems that present similarities can be then modeled by the same model. Two machine learning approaches can be used to express and learn integrative models: Statistical relational learning (Koller et al, 2007) and probabilistic programming (De Raedt and Kimmig, 2015). The former takes relations between co-dependent models into account explicitly, by combining probabilistic graphical models with relational representations. The latter allows to model complex and probabilistic situations.

A fleet of sensorized assets produces a huge amount of time-variant data, rich of actionable information that can be exploited for query and analysis. For this reason, the collected data needs to be properly organized and stored since it represents an asset for the management's decision. This scenario is common for many companies and is handled by building a data warehouse (Inmon and Hackathorn, 1994). Ponniah provides a general overview of the leading principles for defining the data warehouse (Ponniah, 2001). He introduces the concepts of dimension analysis and information package. The former describes, usually in business terms, the type of information that is relevant to be stored (e.g. product, factory, etc.). The latter is a diagram for collecting the relevant metrics associated with each dimension. The metric describes the measured facts associated to the dimension (e.g. sales per product). The metric is organized in a hierarchic manner, namely with different level of details, to support the drilling down of the analysis (e.g. sales in city, region, county).

3 General Methodology

Benchmarking and profiling of assets are activities that depend on asset performance comparison. Assets are compared based on how they behave when they operate in the same context. This relationship input-output is often not known beforehand (e.g. there are no physical models to rely on) and needs to be inferred. For this purpose, it is possible to bin all factors that define the operation contexts in homogeneous multi-dimensional windows and then characterize asset performance within each window. In this manner, within the same window, assets are compared in the same conditions. Finally, extending the comparison to all windows, assets are compared taking into account all the possible contexts in which they can operate. We call these multi-dimensional windows "hypercubes" and this section describes how to construct and exploit them.

Characterization of the Operational Context

The operational context of an asset is characterized via hypercubes. The design of the hypercube reflects Ponniah's principles for building a data warehouse (Ponniah, 2001). Here, the business dimensions are the factors that influence asset behavior, the hierarchical dimension is based on the required resolution for drilling down the analysis, the facts are business metrics used to describe asset performance. In our context, the hypercube is defined in three steps:

1. Decide which *performance indicator* of the asset needs to be monitored along with the statistic that is used for describing the performance. The performance indicator can be energy production, number of failures, efficiency etc. whereas the statistic can be the median, average, etc.

2. Decide the N factors that are relevant to describe the operating context of the asset. Without the support of a domain expert, this may require data exploration techniques, e.g. multi-variate analysis, clustering, pattern mining, to identify potentially relevant factors that can influence the performance indicator, e.g. exogenous factors such as temperature, location, light, etc.

3. Bin the values of each one of the N factors. This can be done using equal-width binning, equal-frequency binning, domain-specific binning, adaptive intelligent binning (De Meyer et al, 2008), etc. Once defined, the bin partitions remain fixed. Fixed boundaries allow the analysis of performance evolution within the same operational context. The resolution used for the bin is a trade-off for having a good resolution along with a significant number of points per bins. High resolution is needed to characterize properly the asset behaviour. On the other hand, a too high resolution generates too many bins with few or no data points. The latter case is an issue since only a limited amount of points are used to characterize the typical effect of the context on asset performance.

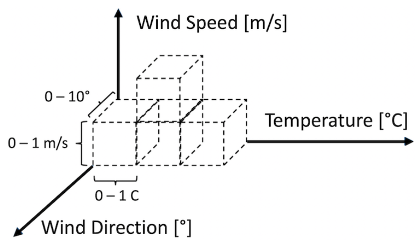

Figure 1: Example of 4 hypercubes of 3 dimensions. In this example, the resolution of the hypercubes is for wind speed 1 m/s, for temperature 1 °C and for wind direction 10 °.

We define "hypercube" as a cube of N dimensions (one dimension per factor). The hypercubes are a partition of a bounded box of values with all mathematical properties of such a partition. Figure 1 reports an example of 4 hypercubes. Each hypercube stores all data points that are included within its boundaries. These points are then used to compute for each hypercube the statistic of the performance indicator. Storing this data is useful when the user wants to change the hypercube resolution or wants to compute other performance indicators.

Context-aware Performance Benchmarking

Hypercubes can be used for context-aware performance benchmarking of the single asset with respect to

a) its own historical records or

b) the other assets of the fleet.

In the first case, we use the historical records of the single asset to compute for each hypercube the statistic of the performance indicator (e.g. median energy production). Then, we label any point within the same hypercube as "underperforming" if its value is below a certain user-defined threshold (e.g. 80% below the median yield of that hypercube). In the second case, we collect the data generated by all fleet assets. The gathered data is used to compute the statistic for each hypercube. Like the previous case, we label any point within the same hypercube as underperforming if its value is below a certain user-defined threshold. In both cases, the statistics median and average can be replaced by other statistics.

It is important to observe that our method, being unsupervised, can be used to identify abnormal behaviors in *unlabeled* datasets. Consequently, it can be used in any domain where sensorized assets keep track of historical data (without the need of ground truth).

Context-aware Performance Profiling

Hypercubes can be used to profile assets, namely identify assets with similar performance behavior across different operational contexts. Hypercubes can be used as features which have as a value the chosen statistic for the performance indicator (e.g. median energy production). These features can then be exploited for clustering (e.g. via Hierarchical Clustering, K-Means, DBSCAN) in order to identify similarly behaving assets. In this sense, hypercubes represent a unique contextual fingerprinting of asset performance.

Optimization of fleet performance generally requires continuous updates of the fleet model that cannot be easily automatized. Indeed, assets are continuously added, removed and/or replaced from the fleet. Consequently, a human intervention is needed to develop the new models for new asset types, their interactions or to update the existing models (e.g. due to component replacement or new setups). Considering that even identical assets can behave differently (since they may operate in different conditions), it becomes clear why creating, optimizing and maintaining a model for each asset is not scalable or cost-efficient. Finally, in some cases this activity may not even be feasible since historical data is not yet available (cold start problem).

Our methodology provides the basis for handling these issues. Thanks to the context-aware asset profiling, assets with similar behavior can be described by a single prototypical model. Then the prototypical models can be combined to build the fleet model. On the one hand, a prototypical model helps to handle the cold start problem of a new asset since it provides a bootstrapping model. On the other hand, addition (update, or removal) of an asset would concern only one prototypical model, making the fleet model more resilient to changes.

4 Implementation and Dataset

We apply our methodology in the wind turbine domain. Energy production is the most important characteristic for a wind turbine. For this reason, we select it as a performance indicator and its median as a statistic. To characterize the operating context we use wind speed, wind direction and temperature. The values of these environmental factors are uniformly binned and define a hypercube of three dimensions. The resolution used for the hypercube is 1 m/s for wind speed, 10 ° for wind direction and 1 °C for temperature. Table 1 describes the distribution of data points for a sample of wind turbines in a sample of hypercubes. As can be observed, the distribution of data points over the hypercubes differs across turbines, illustrating the difference in operational behavior. This is mainly due to the position that these wind turbines have in the wind farm.

Table 1: Distribution of data points for a sample of wind turbines in a sample of hypercubes.

	Wind Turbine				
Hypercube	1	2	3	4	5
3.5–4.5 m/s ; 5.5–6.5 °C ; 90–100 °	14	12	14	30	35
4.5–5.5 m/s ; 5.5–6.5 °C ; 90–100 °	11	13	8	15	18
5.5–6.5 m/s ; 5.5–6.5 °C ; 90–100 °	21	9	15	21	23

As a dataset, we use real-world data coming from a fleet of more than 20 wind turbines (for privacy reasons, we cannot disclose further information on the wind farm). For each wind turbine, the data contains the evolution of wind speed, wind direction, temperatures and active power (energy production) averaged over a 10 minute interval across two years. From the original dataset, we inspect whether turbines have outliers for temperature, wind speed, wind direction or active power. Only for temperature, we removed from two wind turbines in the fleet respectively 42 and 22 data points. These points were removed since the temperature detected was 10 degrees higher than the average temperature of the fleet.

5 Results and Discussion

This section presents concrete examples on how the methodology can be used for benchmarking and profiling of asset(s) in the wind turbine domain.

5.1 Context-aware Performance Benchmarking of a Single Asset

In the wind turbine domain the operational context of the asset is evaluated via power curves. An example of power curve is reported in Figure 2. Beyond the wind speed, the power curve does not account for other exogenous factors. Consequently, it limits the possibility to understand in which operational context the turbine is underperforming. To address this shortcoming, we show how hypercubes can be used for a context-aware performance benchmarking:

1. From the turbine historical records, we extract the hypercubes and retain only the ones that host at least 10 points.
2. For each hypercube the median energy production is computed.
3. Each point is labeled as underperforming if its energy production is more than 80% below the median energy production for that hypercube.

Figure 3 provides a benchmark of turbine performance. Here, we can see that the turbine underperformance is located in specific ranges of the wind speed and wind direction. Within these ranges, the underformance is more often observed when wind direction is between 200–250 ° and wind speed is between 9–17 m/s. This scatterplot can be used to easily identify whether there are patterns where the turbine underperforms and also to prioritize inspections in the operational contexts where underperformance happens more frequently. None of these insights could have been extracted using a standard operative curve. Finally, it is worth to point out that this type of plot is just one possible example on how the operator can use the methodology for gaining insights on asset behavior via visual inspection. Further insights would be easily extracted just changing the statistic, the performance indicator or using the fleet of assets as a baseline for comparison.

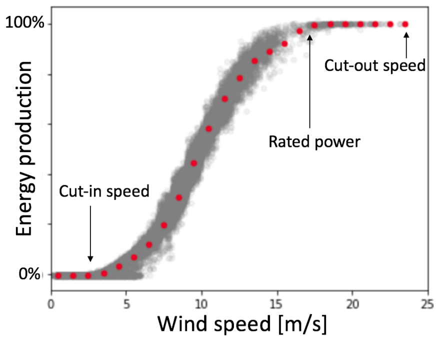

Figure 2: The theoretical power curve is represented by the red circles. The grey circles describe the raw data. The curve is characterized by a rated power (maximum generated energy) and a cut-in speed cut-out speed (the speed at which the turbine starts and stops generating energy, respectively).

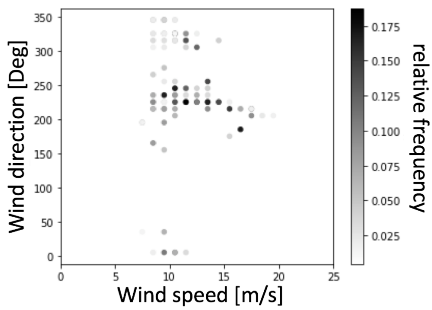

Figure 3: The circles represent hypercubes where at least one point is labeled as underperforming. The grey-scale on the right side describes the relative frequency of underforming points within that hypercube.

5.2 Context-aware Performance Profiling and Benchmarking of a Fleet of Assets

The use of power curves for identifying wind turbines with similar behaviors is problematic. Indeed, also two identical wind turbines may have different power curves due to other exogenous factors (e.g. temperature and wind direction) that beyond the wind speed can influence the energy production. Our methodology can be used to address this issue since the performance of the asset is compared across all operational contexts.

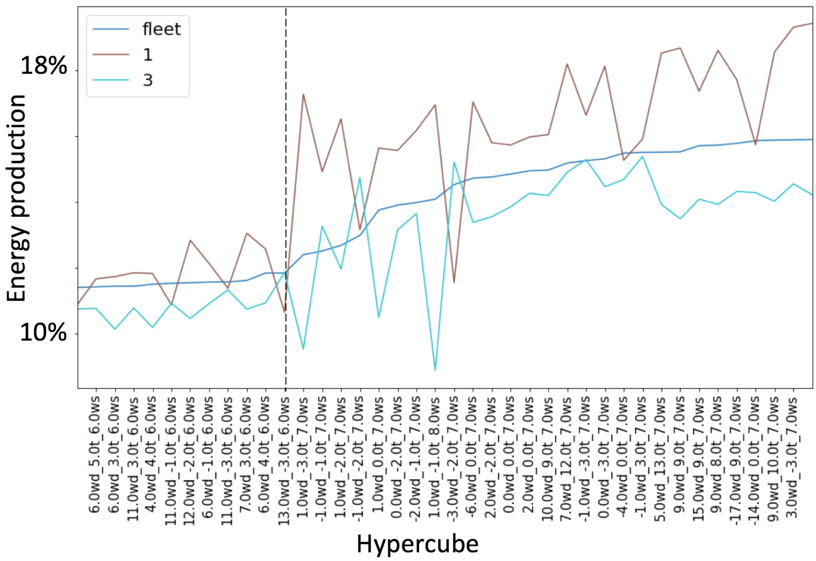

Figure 4: Context-sensitive performance benchmarking of clusters over a sample of hypercubes. The y-axis reports the energy production as a percentage with respect to the maximum rated power. The x-axis reports the hypercubes as a tuple: wind direction (wd), temperature (t) and wind speed (ws). The hypercubes are sorted according to the fleet average active power.

From the operational point of view, we proceed in the following way. We extract hypercubes and retain only the ones that have median energy production above 0 for all wind turbines. Wind turbines are then clustered using as features the median energy production of each hypercube. As a clustering algorithm, we adopt Ward's method of hierarchical clustering (Friedman et al, 2001). At the end of the clustering, we obtain 5 clusters with 5–9 assets plus one cluster containing only one asset. We refer to these clusters as 1, 2, 3, 4 and 5. We can benchmark these clusters by computing the difference of their performances with respect to the fleet performance (the same analysis can be done also at asset level).

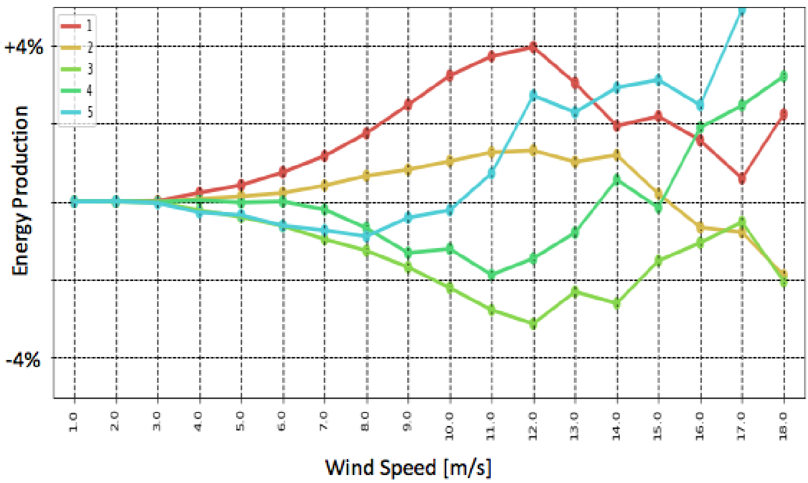

Figure 5: Difference of energy production between cluster and fleet (x-axis) for different wind speeds.

Figure 4 shows the performance of the fleet and two clusters (1 and 3) over a subset of hypercubes sorted according to the fleet average active power. As we can see, the hypercubes allow to compare visually the performance of group of asset clusters for different operational contexts. Generally, cluster 1 and 3 are respectively above and below the fleet average. However, when there is a change of wind speed, the benchmarking of the two clusters is more complex. To highlight this aspect, we can focus on Figure 4. Here, the dashed line separates hypercubes with wind speed above 6 m/s with respect to hypercubes with higher wind speed. Just after this step increase of the wind speed, the performance of the two clusters is more variable and dependent on secondary factors, such as the wind direction and temperature.

We can further explore the behavior of the wind turbines by considering only one dimension of the hypercubes each time. Figure 5 shows the evolution of the difference between the median active power of several clusters with respect to the fleet (the x-axis). As we can see, the performance of one cluster can be better or worse than another one depending on the wind range analyzed. For instance, the energy production of cluster 5 remains below the fleet average as long as the wind speed is lower than 10 m/s. On the other hand, when the wind speed is above 14 m/s, the energy production is the highest of the fleet which is probably due to faulty sensor measurement.

6 Conclusions

We developed a new methodology to characterize the behavior of assets in terms of performance and operational context. This methodology is domain-agnostic, does not require ground truth and can handle operational contexts described by multiple factors. The methodology facilitates explorative analysis and can be used to benchmark and profile assets.

We evaluated our methodology using real-world data coming from the wind energy domain. Here, we described how this methodology overcomes the limitation of the standard power curve approach. Our methodology was able to detect autonomously when the asset was underperforming based on asset historical data and user defined threshold. Finally, we also showed how to apply our methodology for partitioning a fleet of wind turbines into groups that have similar performance in a given operational context. For these clusters, we also illustrated how to investigate the influence of the environmental factors on the performance of the assets. As a future step, we will leverage on context-aware profiling to develop scalable models via integrative modeling and context-aware performance characterization for monitoring asset performance decay.

Acknowledgements This work was supported by FOD Economie through the project BitWind.

References

Bell MB, Foslien W (2005) Early Event Detection – Results From A Prototype Implementation. In: 17th Annual Ethylene Producers Conference, Session TA006-Ethylene Plant Process Control, Vol. 1, pp. 727–741.

Blesa J, Jiménez P, Rotondo D, Nejjari F, Puig V (2015) An Interval NLPV Parity Equations Approach for Fault Detection and Isolation of a Wind Farm. IEEE Transactions on Industrial Electronics 62(6):3794–3805. DOI: 10.1109/TIE.2014.2386293.

De Meyer T, Sinnaeve D, Van Gasse B, Tsiporkova E, Rietzschel ER, De Buyzere ML, Gillebert TC, Bekaert S, Martins JC, Van Criekinge W (2008) NMR-based Characterization of Metabolic Alterations in Hypertension Using an Adaptive, Intelligent Binning Algorithm. Analytical Chemistry 80(10):3783–3790. DOI: 10.1021/ac7025964.

De Raedt L, Kimmig A (2015) Probabilistic (Logic) Programming Concepts. Machine Learning 100(1):5–47. DOI: 10.1007/s10994-015-5494-z.

Freire NMA, Estima JO, Cardoso AJM, et al (2013) Open-circuit Fault Diagnosis in PMSG Drives for Wind Turbine Applications. IEEE Transactions on Industrial Electronics 60(9):3957–3967. DOI: 10.1109/TIE.2012.2207655.

Friedman J, Hastie T, Tibshirani R (2001) The Elements of Statistical Learning, 1st edn. Springer Series in Statistics, Springer Science & Business Media, New York (USA). ISBN: 978-0-387952-84-0.

Gill S, Stephen B, Galloway S (2012) Wind Turbine Condition Assessment Through Power Curve Copula Modeling. IEEE Transactions on Sustainable Energy 3(1):94–101. DOI: 10.1109/TSTE.2011.2167164.

Gong X, Qiao W (2013) Bearing Fault Diagnosis for Direct-drive Wind Turbines via Current-demodulated Signals. IEEE Transactions on Industrial Electronics 60(8):3419–3428. DOI: 10.1109/TIE.2013.2238871.

Gong X, Qiao W (2015) Current-based Mechanical Fault Detection for Direct-drive Wind Turbines via Synchronous Sampling and Impulse Detection. IEEE Transactions on Industrial Electronics 62(3):1693–1702. DOI: 10.1109/TIE.2014.2363440.

Gorinevsky D, Dittmar K, Mylaraswamy D, Nwadiogbu E (2002) Model-based Diagnostics for an Aircraft Auxiliary Power Unit. In: Proceedings of the International Conference on Control Applications, Institute of Electrical and Electronics Engineers (IEEE), New York (USA), Vol. 1, pp. 215–220. DOI: 10.1109/CCA.2002.1040188.

IEC (2008) Wind Turbines Part 12-2: Power Performance of Electricity Producing Wind Turbines based on Nacelle Anemometry. Tech. Rep., International Electrotechnical Commission (IEC). URL: `https://webstore.iec.ch/publication/5430`.

Inmon WH, Hackathorn RD (1994) Using the Data Warehouse, 1st edn. John Wiley & Sons, Somerset (USA). ISBN: 978-0-471059-66-0.

Iyer KLV, Lu X, Usama Y, Ramakrishnan V, Kar NC (2013) A Twofold Daubechies-wavelet-based Module for Fault Detection and Voltage Regulation in SEIGs for Distributed Wind Power Generation. IEEE Transactions on Industrial Electronics 60(4):1638–1651. DOI: 10.1109/TIE.2012.2188258.

Jia X, Jin C, Buzza M, Wang W, Lee J (2016) Wind Turbine Performance Degradation Assessment based on a Novel Similarity Metric for Machine Performance Curves. Renewable Energy 99:1191–1201. DOI: 10.1016/j.renene.2016.08.018.

Kim K, Mylaraswamy D (2006) Fault Diagnosis and Prognosis of Gas Turbine Engines based on Qualitative Modeling. In: ASME Turbo Expo: Power for Land, Sea, and Air, American Society of Mechanical Engineers (ASME), New York (USA), pp. 881–889. DOI: 10.1115/GT2006-91210.

Koller D, Friedman N, Džeroski S, Sutton C, McCallum A, Pfeffer A, Abbeel P, Wong MF, Heckerman D, Meek C, et al (2007) Introduction to Statistical Relational Learning. Adaptive Computation and Machine Learning, Getoor L, Taskar B (eds), MIT Press, Cambridge (USA). ISBN: 978-0-262072-88-5.

Lapira E, Brisset D, Ardakani HD, Siegel D, Lee J (2012) Wind Turbine Performance Assessment Using Multi-regime Modeling Approach. Renewable Energy 45:86–95. DOI: 10.1016/j.renene.2012.02.018.

Lydia M, Kumar SS, Selvakumar AI, Kumar GEP (2014) A Comprehensive Review on Wind Turbine Power Curve Modeling Techniques. Renewable and Sustainable Energy Reviews 30:452–460. DOI: 10.1016/j.rser.2013.10.030.

Ponniah P (2001) Data Warehousing Fundamentals. John Wiley & Sons, New York (USA). DOI: 10.1002/0471221627.

Sanchez H, Escobet T, Puig V, Odgaard PF (2015) Fault Diagnosis of an Advanced Wind Turbine Benchmark Using Interval-based ARRs and Observers. IEEE Transactions on Industrial Electronics 62(6):3783–3793. DOI: 10.1109/TIE.2015.2399401.

Uluyol O, Parthasarathy G, Foslien W, Kim K (2011) Power Curve Analytic for Wind Turbine Performance Monitoring and Prognostics. In: Annual Conference of the Prognostics and Health Management Society, Vol. 2, pp. 1–8.

Wang W, McFadden P (1996) Application of Wavelets to Gearbox Vibration Signals for Fault Detection. Journal of Sound and Vibration 192(5):927–939. DOI: 10.1006/jsvi.1996.0226.