Understanding Anomalies: Visualizing Sensor Data for Condition Monitoring of Manufacturing Machines

Daniel Olivotti[1], and Dennis Eilers[1]

[1] Leibniz Universität Hannover, Institute for Information Systems, Hannover, Germany
{olivotti,eilers}@iwi.uni-hannover.de

Abstract. In a more and more service oriented manufacturing industry new data driven challenges like predictive maintenance arise. For example, machine learning models can use sensor data to predict anomalies during machine operation. Such models usually base on experiences from past data to train these algorithms. However, since components are often used in different machines with different and partly unique or new domains, experiences about mutual interferences are missing. In this study we try to address this issue by introducing a visualization technique for an intuitive anomaly detection which allows domain experts and engineers to monitor the condition of a machine over time. The heat map based visualization highlights unusual operation measurements dependent on different sensor data combinations. With domain and engineering knowledge, the insights can be used to identify case based reasons for a changing behavior. The application is tested with a demonstration machine.

Keywords: anomaly detection, condition monitoring, sensor data, manufacturing machines, visualization

1 Introduction and Motivation

In manufacturing industries services play an important role. Machine builders as well as component suppliers are offering physical products along with services, so called product-service systems, more and more [1], [2], [3]. Services lead to further revenue channels and long-term customer satisfaction [1], [4].

Examples for these services are condition monitoring and predictive maintenance. Potential downtimes of machines in manufacturing industries usually result in high breakdown costs. To counteract this, maintenance helps to keep machine availability at a high level. With the help of predictive maintenance, anomalies can be detected early, and the necessary measures can be taken. This results in high cost savings as well as high machine output. The combination of actual sensor values with prediction models in order to forecast a degradation of components is a challenging task [5]. Machine builders integrate components from different component suppliers in their machines. Individual requirements and preferences of machine owners are considered and increases the variety of machines. The interplay between components in machines is essential when anomaly detection is performed, and breakdown possibilities of machines are calculated. Component suppliers have detailed knowledge of their own

Multikonferenz Wirtschaftsinformatik 2018,
March 06-09, 2018, Lüneburg, Germany

products but combining this company specific knowledge and different sensor data for a whole machine can lead to a better understanding of observed phenomena.

In this study, we introduce a heat map based visualization technique to better identify the reasons for anomalies, depending on different sensor data combinations. The idea is to use machine learning models trained with time series operation data (e.g. motor speed, or torque) in order to learn how the normal operation of a unique and individual machine looks like. In a subsequent step the models are applied to constantly predict future operations states such that actual and predicted values can be compared. The resulting residuals are then visualized on a heat map color scale depending on different sensor values of different machine components. The visualization allows to intuitively identify clusters of unusual operation states in the context of sensor data like the temperature or the vibration of certain components. The technique also allows continuous updates to assess mutual interferences between components over time. By using this visual analytics approach in combination with machine learning models, application domain experts and engineers can discuss possible reasons or improvements regarding unusual observations more easily. We illustrate the functionality by using a unique demonstration machine which is built to test the interplay of different components during operation.

The remainder of this paper is structured as follows. The research background with focus on predictive maintenance and condition-based monitoring as well as visualizations for data analytics applications is stated out in Section 2. The constructed demonstration machine along with a description of relevant sensor data are presented in Section 3. The anomaly detection approach based on Artificial Neural Networks (ANN) is presented in Section 4. In Section 5 the heat map based visualization applied to the sensor data of the demonstration machine is focused. The results and limitations are discussed in Section 6. Conclusions and an outlook are presented in Section 7.

2 Research Background

Maintenance of machines is a widely discussed topic in the literature. Predictive maintenance is not based on fixed time schedules but aims to predict failures before they occur which reduces machine downtimes and costs of maintenance. Condition monitoring and predictive maintenance received increasing attention recently due to the availability of massive amounts of machine data. The goal is to identify possible anomalies in the operation of machines in an early stage of the impending problem to prevent major failures and corresponding downtime cost or more expensive repair measures [6]. [7] define predictive maintenance as a condition-based preventive maintenance approach. The actual condition derives from sensors or models or combination of both. [8] divides prognostic models in physical models, knowledgebase models and data-driven models. There exists for example applications of assessing the condition of wind turbines by analyzing sensor data [6], [9]. In the literature, several methodologies like statistical analysis or artificial intelligence are discussed. Various models and algorithms are used for condition-based maintenance [8], [10], [11], [12]. Also for machinery diagnostics and prognostics several research projects are already conducted [10], [13]. Predictive maintenance is seen as IT-based service in the context of the Internet of Things [14]. A service platform is required to offer such services in value networks [15]. Visualization techniques which goes further than simple line or bar charts or the display of key performance indicators are rarely discussed in literature.

Today, more and more individualized machines offer a great opportunity for the production of smaller lot sizes and thus also more individual products but also increases the complexity of tasks like predictive maintenance since no reference values and experiences with certain combinations of machine components are available. Therefore, individual application domain and engineering knowledge is required to be still able to identify possible anomalies or unusual behavior. Visualization is a suitable vehicle to trigger human inference from massive amounts of data [16]. Using visualization to support data analyst can be summarized under the term visual analytics which is first defined by [17]. For comprehensive surveys we refer to [18], [19]. Studies like [20], [21] even show that with help of visualizations human can outperform pure machine learning models in certain analytics task. This shows that human reasoning and domain knowledge is especially important in a sparse data space and with noise and uncertainty in the data. In addition, visualization can improve the communications which is essential for the overall team performance [22]. This is necessary to discuss possible reasons and solutions for identified problems with machines. Promoting a visual literacy can establish a common language on the same level of complexity [23] between different expert groups with different proficiencies which also helps to facilitate efficient problem identification. Therefore, visualization in general can help to tackle the challenges which arise in today’s service based manufacturing.

Since nowadays the condition of a machine can be measured by a huge amount of different sensor data, for example to capture the temperature or the vibration, a visualization technique must be able to handle large and diverse datasets. Heat maps are often applied to large amounts of data to provide an intuitive and easy understandable overview in a two-dimensional data space. They are mainly used to show particularly distinctive patterns in certain regions and are therefore suitable for detecting unusual patterns during the operation of a machine. Researchers have shown the applicability of heat maps for several different problem domains. For example [24] use heat maps to illustrate results of a risk management model, [25] visualizes the density of forecasting ensembles over certain time steps and [26] use 2D and 3D heat maps to visualizes performance measures of regression models.

Standard approaches for processing sensor data in condition monitoring and predictive maintenance applications use simple threshold values as problem indicators or more advanced machine learning algorithms to classify machine-related anomalies [27]. Both require experience and, in particular, previous data about a certain machine type. This is often not the case with custom designs, so there is no experience about the interaction of the components. The visualization approach in this study should addresses the question of how to better understand and assess normal or unusual behavior of machines where no or only little experience and past data are available. Therefore, we use the afore mentioned ideas to propose a comprehensible and intuitive heat map visualization which supports application domain experts and engineers in assessing machine conditions to identify the reasons for a problem based on large amount of sensor data.

3 Demonstration Machine and Sensor Data

A demonstration machine within a project with an automation and engineering company is used to provide data in a flexible way. The demonstration machine represents a practical example of a manufacturing machine. The machine is equipped with hoist and conveyor applications to move goods. These exemplary applications are typically for the involved company. Each of these applications has different technical requirements and specifications based on positioning, movement and dynamic. The interplay of sensor data from different applications and motion axes is not known in advance and therefore the anomaly detection, presented in the following sections, is used. Especially when goods with different sizes, weights and position need to be processed, the physical strain cannot be estimated simple by the mean or intervals. In total seven axes provide a circular material flow of goods in the demonstration machine. Two hoist axes are equipped each with a conveyor and are on the left and right side of the machine. In the center of the machine three conveyors are arranged one above each other. All axes are powered by electrically-geared motors. These geared motors are powered by inverters.

A programmable logic controller (PLC) send commands to the inverters and from there they are passed on to the geared motors. Actual status information and sensor data is transferred in reverse order to the PLC. In the PLC a logic controls the process based on defined states and actions. The inverters provide data within cycle times of less than 10 milliseconds (ms). These is identified as a stable cycle time for recording data as the necessary fieldbus communication is not interfered. From the PLC the data is provided to an industrial PC where the data is stored. There exist several protocols to provide data over the network e.g. MQTT and OPC Unified Architecture. Voltage, current, torque and rotational speed of the engines are recorded and used for anomaly detection. These are also the main process variables for the movement control of the seven axes. Other data from inverters and geared-motors are not considered at the moment. The built-in components are able to retrieve this process parameters without the usage of additional sensors. A visualization terminal displays relevant parameters to the machine operator.

4 Time Series Anomaly Detection with Artificial Neural Networks

In this section we describe our analysis approach to detect anomalies of the demonstration machine. In order to evaluate the current operation status, we investigate a time series of torque measured in newton metre (Nm) which reflects the different positions of a weight on a conveyor belt. Figure 1 shows a plot of the torque values in

Nm for six cycles. This time series captures the occurring movement patterns during the ongoing operation and thus serves as a measure for the condition of the machine.

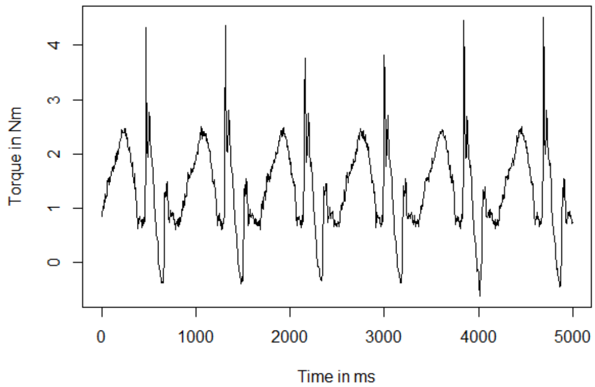

Figure 1. Time series of torque in Nm

To identify an unusual behavior during the operation we use the following anomaly detection approach: Firstly, we sample 60 cycles of sensor data which defines the expected behavior of the demonstration machine in normal conditions. Secondly, ANNs are used to learn and afterwards predict the sensor values for the next time step. Thirdly, the measured value at the next time step is compared to the predicted value and the resulting residuals are evaluated based on a confidence level of how likely an occurrence of this magnitude actually is under the assumption of normal operation conditions. In our example we use simple feed-forward ANNs trained with lags of the 20 past steps of the time series. 30 ANNs with different random weight initializations are trained to build an ensemble based on averaged predictions. To prevent the problem of overfitting we use an early stopping approach. We therefore split the available training data into 80% actual training data and 20% validation data. The 80% portion is used for the iterative learning process, while the 20% portion simulates the out-of-sample results. The actual out-of-sample performance is then observed on a completely independent holdout test dataset. The training process is terminated as soon as the error on the validation data no longer decreases. The ANNs consist of two hidden layers with 50 and 25 hidden neurons using the hyperbolic tangent as activation function. ANNs in general serve as a machine learning method for classification and regression problems and are applied to a variety of forecasting problems [28]. They are able to learn complex, non-linear patterns from data and are reasonably robust against noisy data [29]. To implement the described approach, we use the H2O framework for deep learning (version 3.10.5.3) and perform all necessary data pre-processing and result analysis in R (version 3.4.1) and the H2O R interface. Since ANNs and their learning process is not the focus of this work we refer to [30]. But there is much leeway to improving the models further for example by an extensive hyper parameter optimization (number of hidden layers and neurons), different ensemble methods and more sophisticated regularization techniques like dropout. In this study however, we explicitly focus on a simple approach which performs reasonably well on our data set to demonstrate the applicability of the approach and the subsequent anomaly visualization. Figure 2 compares the actual measurement with the predicted values on an out-of-sample test set, while Figure 3 shows the residual distribution (root-mean-square error: 0.05156, mean absolute error: 0.03864).

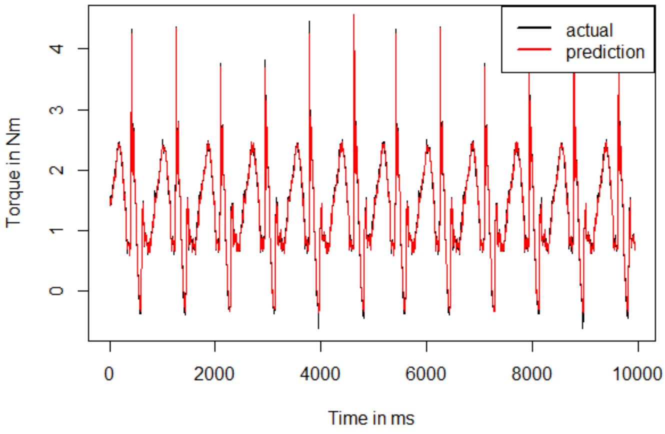

Figure 2. Actual time series vs. prediction

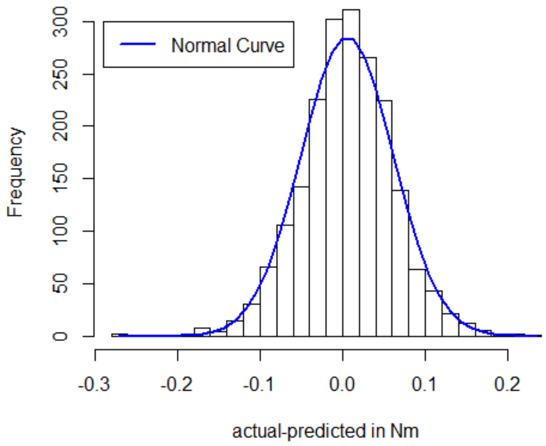

Figure 3. Residual distribution and normal distribution fit

The presented procedure is a standard approach in time series anomaly detection [31] and serves as an early warning system to trigger a subsequent investigation of the reasons for the deviations by means of the analysis method presented in the following section.

5 Heat Map based Visualization

In this section we use the residuals of the previously described time series forecast approach as an anomaly indicator during the operation of the demonstration machine to find the reasons for possible problems. In real world industry scenarios, the interplay of different components in different areas of application of machines can lead to unexpected wear or malfunctions, which should be detected as early as possible, corrected before major failures occur and avoided in the future. This is a challenging task especially with individually-engineered machines due to the lack of experience and missing data. Therefore, the goal is to better understand possible reasons for an unusual behavior of machines based on engineering and application domain knowledge. For this purpose, we propose to use a heat map based residual visualization technique, first introduced by [32] which was previously applied as an evaluation method for machine learning regression models. The idea is to visualize model errors in the context of two different possible features (independent variables) which may have an influence on the distribution of the errors and therefore contain valuable information for explaining or mitigating the problem. To illustrate the functionality, we first artificially simulate two sensor measure time series from two components C-A and C-B of the demonstration machine in addition to the already available torque time series. The first sensor measures the vibration of C-A and the second sensor the vibration of C-B over time. Figure 4 shows the time series during an out of sample period which was not part of the ANN training process.

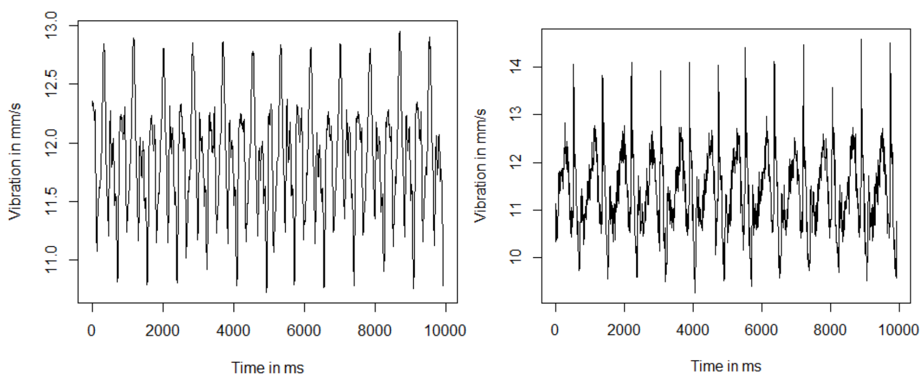

Figure 4. Time series of simulated vibration data for two different components

The idea is to intuitively link the residuals of the time series forecast with measured data of any sensor from any component of the machine to highlight possible dependencies and reasons for observed unusual behavior. In this example we use the two simulated sensors to construct a two-dimensional data space which results in a simple scatter plot as shown in the left plot of Figure 5. We incorporate the density information of the distribution as suggested by [33] to ensure a meaningful representation of the data even for large sample sizes. Darker regions represent a higher density. To incorporate the information about possible anomalies, each data point is now assigned its corresponding residual from the forecasting application. The values of the residuals should be the basis for the color scale of a heat map which incorporates the forecasting error as a further dimension into the graphic. It is assumed that each residual represents the local model results depending on the values of the sensor data. To show the dependencies between two different sensor data sets and the residuals by means of a smoothed color scale in a heat map approach, a regular grid with the respective color codes needs to be calculated. Since the data points are usually not on a regular grid, a method for weighting the influence of each residual on each point on the regular grind in the data space is necessary. Therefore, we use a two-dimensional Gaussian kernel which ensures a strong influence of a residual near the position of the corresponding data point and a decreasing influence with increasing distance. A detailed mathematical description of the procedure can be found in [32]. The results of this simulation can be seen in the right plot of Figure 5.

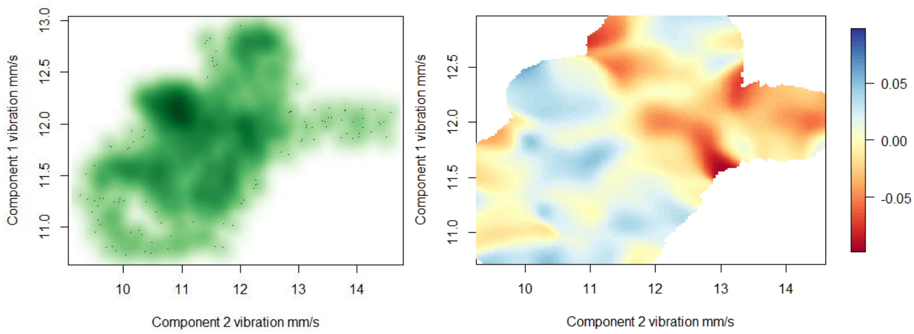

Figure 5. Scatter plot (left) and heat map (right) of residuals depending on sensor data

The color represents the local model errors when predicting the measured Nm of the machine for the next time step which therefore highlights possible unusual behavior of the machine depending on the vibration of component C-A and C-B. In this example, we chose a diverging color palette with 64 divisions. Blue regions represent high positive deviations from the prediction model (actual value in this state is higher than expected by the model) and red regions represent high negative deviations from the prediction model (actual value in this state is lower than expected by the model). In this example, one can observe an overestimation of torque in regions of the data space with higher vibration of both sensors. Underestimations can be observed in conjunction with less vibrations. If the sensor data and therefore the respective component of the machine would have no influence on the torque the expected heat map would have no visible patterns which represents a normal distribution of the residuals. Here, however, various reasons can be responsible for the observed patterns. For example, wear on the motor, incorrectly placed loads or frictions of various components should be taken into account. In applications with many more and different kinds of available sensor data, the visualization can be extended by using a scatter plot matrix to get a quick and comprehensible overview of sensor values/pairwise sensor combinations and the corresponding possible anomalies during the machine operation. The presented technique also allows the user to monitor unusual behavior over time by comparing the heat maps/scatter plot matrix from different time periods which can reveal changes in the condition of special components in the context of the overall machine behavior. Therefore, the approach links unusual machine behavior with sensor data in an intuitive graphical view which can be better interpreted by human experts.

6 Discussion and Limitations

In this first application example, we have shown how heat map based visualizations can be used to monitor the condition of individual machines in the context of measured sensor data from different machine components. The idea follows a two-step approach: Firstly, a machine learning model is used to identify deviations between an expected operation state and the actual observed state of a machine. The resulting residuals of the model are used to construct the heat map which visualizes these deviations on a defined color scale, depending on two different sensor datasets which construct the two-dimensional data space. This enables the user to identify possible clusters of unusual machine behavior in specific regions of sensor measurements. The connection between the anomalies (color scale) and sensor data (axes) can help to identify possible reasons and solutions for identified problems. Such a combination of a machine learning and visual analytics approach can be beneficial especially in situation where only little experience and past data about the machines are available. In addition, the visualization is easy to interpret which can facilitate discussion about the results between people with different background like engineers, data scientist or application domain experts. The technique also enables a monitoring over time by comparing the visualizations from different periods. This allows an assessment about structural changes in the condition of certain components or the combinations of certain components which can be used to perform predictive maintenance measures.

However, both steps of the proposed approach face some difficulties and drawbacks. First the machine learning model currently uses one simple time series of torque for the training process which represents the operation status of the machine. It is assumed that the data for the training represent a normal operation of the machine. The anomaly detection is therefore only able to identify deviations from an assumed ideal situation but there is no possibility to optimize the initial operation behavior. A further limitation is the method to learn the normal operation behavior. Currently a simple feed-forward ANN is used, with some basic optimization and regularization procedures. In further research this will be benchmarked or replaced by more sophisticated algorithms like recurrent neural networks, especially Long Short-Term Memory (LSTM), which are more suitable for time series prediction. The modeling process in general needs therefore further optimization and benchmarks to find the best solution for learning normal operation behaviors of machines.

Regarding the visualization approach, there also exist several limitations and drawbacks. Visualizations in general always require human judgements which includes errors and wrong conclusions. Heat maps in particular are restricted to two dimensions and an additional color scale. The increasing number of sensors, however, make the visualization of all possible sensor combinations more difficult. A scatter plot matrix can help to mitigate the problem. Another problem of heat maps is that they are hardly comparable to each other since the color scale is relative to the measured values and residuals. Quantitative comparisons between different points in time for example can lead to wrong conclusions. Currently also no information about the density is available in the representation, so the heat maps must be shown in conjunction with scatter plots which contain density information.

Another important part of this study is the demonstration machine which provides sensor data from key process parameters. This is advantageous because this data is already used for the control of the motion axes and no additional sensors needed. Cycle times of less than 10 ms are currently used to process sensor data from the PLC to the industrial PC. A tradeoff between low cycle times and high field bus occupancy rate need to be considered. The cycle time must be defined for each parameter to provide the required data quality. In general, additional sensors provides more insights in the condition of machines. The costs of additional sensors must be considered. Common practice is to equip key components for the functions of machines with appropriate sensors. Frequency and volume of sensor data can lead to high amount of data. Policies how long which data is store need to be defined. Interdependencies of components in machines are often complex to model. In some cases, they can solely be found out experimentally. The presented heat map based visualization techniques enables first insights on dependencies of different components. Domain experts or engineers can analysis the affected components in detail. This knowledge can be stored in a knowledge management system to get insights for future applications in advance.

Overall it can be argued that the presented approach can help to trigger human creativity and judgement despite the mentioned difficulties and drawbacks. Especially a visualization approach can only be a support for human analysts who must draw their conclusions based on the provided systems. But to measure the actual benefit of the proposed technique is a challenging task, which is why the evaluation is currently a major limitation of our study.

7 Conclusions and Outlook

Condition monitoring and predictive maintenance aim to reduce maintenance efforts while keeping machine availability high. This is achieved by the use of sensor data from various components and machines. To detect anomalies, the interplay of components in machines is essential. In this article, a heat map based visualization technique is applied to sensor data from manufacturing machines. By combining information of different sensors in a visualization, domain experts or engineers are facilitated to find possible reasons for anomalies during machine operation.

In further research, more components from different component suppliers should be included in the demonstration machine. Long-time test runs enable an evaluation of the presented application of heat map based visualizations in detail. In this article we limited our analysis on different components in machines. An analysis of machines in interchained production lines is also planned to predict machine park availability.

We will also test different machine learning models to better understand the current state of a machine. Especially time series models like LSTM networks might be suitable for this purpose. For live updates of the heat map visualization, the performance of the calculations need to be improved. The current approach is implemented in the programming language R. To increase the performance, other lower level languages like C or languages for high performance computing like Julia might be better choices. Different cloud solutions are also taken in consideration since service-based solutions should be provided.

References

1. Oliva, R., Kallenberg, R.: Managing the transition from products to services. Int. J. Serv. Ind. Manag. 14 (2), 160–172 (2003)
2. Schrödl, H., Bensch, S.: E-Procurement of Cloud-based Information Systems – a Product-Service System Approach. In: Proceedings of the Thirty Fourth International Conference on Information Systems (ICIS), Milan, Italy, December 15–18. (2013)
3. Neff, A.A.: Developing a Maturity Model for Product-Service Systems in Manufacturing Enterprises. In: Proceedings of the Pacific Asia Conference on Information Systems (PACIS), Jeju Island, Korea, June 18–22. (2013)
4. Barrett, M., Davidson, E., Prabhu, J., Vargo, S.L.: Service Innovation in the Digital Age: Key Contributions and Future Directions. MISQ 39 (1), 135–154 (2015)
5. Kaiser, K.A., Gebraeel, N.Z.: Predictive Maintenance Management Using Sensor-Based Degradation Models. IEEE Trans. Syst., Man, and Cybern. - Part A: Syst. and Hum. 39 (4), 840–849 (2009)
6. Garcia, M.C., Sanz-Bobi, M.A., Del Pico, J.: SIMAP: Intelligent System for Predictive Maintenance Application to the health condition monitoring of a windturbine gearbox. Comput. Ind. 57, 552–568 (2006)
7. Mobley, R. D.: An Introduction to Predictive Maintenance, Second Edition, Butterworth-Heinemann, Oxford, UK (2002)
8. Peng, Y., Dong, M., Zuo, M.J.: Current status of machine prognostics in condition-based maintenance: a review. Int. J. Adv. Manuf. Technol. 50, 297–313 (2010)
9. Tchakoua, P., Wamkeue, R., Ouhrouche, M., Slaoui-Hasnaoui, F., Tameghe, T.A., Ekemb, G.: Wind Turbine Condition Monitoring: State-of-the-Art Review, New Trends, and Future Challenges. Energies 7, 2595–2630 (2014)
10. Jardine, A.K.S., Lin, D., Banjevic, D.: A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mech. Syst. Signal Process. 20, 1483-1510 (2006)
11. Ahmad, R., Kamaruddin, S.: An overview of time-based and condition-based maintenance in industrial application. Comput. Ind. Eng. 63, 135–149 (2012)
12. Tsang, A.H.C.: Condition-based maintenance: tools and decision making. J. Qual. Maint. Eng. 1 (3), 3–17 (1995)
13. Kothamasu, R., Huang, S.H., VerDuin, W.H.: System health monitoring and prognostics - a review of current paradigms and practices. Int. J. Adv. Manuf. Technol. 28, 1012–1024 (2006)
14. Wortmann, F., Flüchter, K.: Internet of Things - Technology and Value Added. Bus. Inf. Syst. Eng. 57 (3), 221–224 (2015)

15. Gerloff, C., Cleophas, C.: Excavating the Treasure of IoT Data: An Architecture to Empower Rapid Data Analytics for Predictive Maintenance of Connected Vehicles. In: Proceedings of the Thirty Eighth International Conference on Information Systems (ICIS), Seoul, South Korea, December 10-13. (2017)
16. Keim, D., Andrienko, G., Fekete, J.-D., Görg, C., Kohlhammer, J., Melançon, G.: Visual Analytics: Definition, Process, and Challenges. In: Kerren, A., Stasko, J.T., Fekete, J.-D., North, C. (eds.) Information Visualization, Lecture Notes in Computer Science, pp. 154–175. Springer, Berlin, Heidelberg (2008)
17. Thomas, J.J., and Cook, K.A.: A Visual Analytics Agenda. IEEE Comput. Graph. Appl. 26 (1), 10–13 (2006)
18. Sun, G.-D., Wu, Y.-C., Liang, R.-H., Liu, S.-X.: A Survey of Visual Analytics Techniques and Applications: State-of-the-Art Research and Future Challenges. J. Comput. Sci. Technol. 28 (5), 852–867 (2013)
19. Endert, A., Hossain, M.S., Ramakrishnan, N., North, C., Fiaux, P., Andrews, C.: The Human Is the Loop: New Directions for Visual Analytics. J. Intell. Inf. Syst. 43 (3), 411–435 (2014)
20. Tam, G.K.L., Fang, H., Aubrey, A.J., Grant, P.W., Rosin, P.L., Marshall, D., Chen, M.: Visualization of Time-Series Data in Parameter Space for Understanding Facial Dynamics. Comput. Graph. Forum 30 (3), 901–910 (2011)
21. Tam, G.K.L., Kothari, V., Chen, M.: An Analysis of Machine- and Human-Analytics in Classification. IEEE Trans. Vis. Comput. Graph. 23 (1), 71–80 (2017)
22. Bresciani, S., Eppler, M.J.: The Benefits of Synchronous Collaborative Information Visualization: Evidence from an Experimental Evaluation. IEEE Trans. Vis. and Comput. Graph. 15 (6), 1073–1080 (2009)
23. Fransecky, R.B., Debes, J.L. Visual Literacy: A Way to Learn–A Way to Teach: Association for Educational Communications and Technology, Washington D.C. (1972)
24. Buehler, K.S., Pritsch, G.: Running with Risk. McKinsey Q. 4, 40–49 (2003)
25. Köpp, C., von Mettenheim, H.-J., Breitner, M.H.: Decision Analytics with Heatmap Visualization for Multi-Step Ensemble Data. Bus. Inf. Syst. Eng. 6 (3), 131–140 (2014)
26. Klemm, P., Lawonn, K., Glaßer, S., Niemann, U., Hegenscheid, K., Völzke, H., Preim, B.: 3D Regression Heat Map Analysis of Population Study Data. IEEE Trans. Vis. Comput. Graph. 22 (1), 81–90 (2016)
27. Lee, J., Kao, H.-A., Yang, S.: Service Innovation and Smart Analytics for Industry 4.0 and Big Data Environment. Procedia CIRP 16 (Supplement C), 3–8 (2014)
28. Adya, M., Collopy, F.: How Effective Are Neural Networks at Forecasting and Prediction? A Review and Evaluation. J. Forecast. 17 (5–6), 481–495 (1998)
29. Schocken, S., Ariav, G.: Neural Networks for Decision Support: Problems and Opportunities. Decis. Support Syst.11 (5), 393–414 (1994)
30. Bishop, C.M.: Neural Networks for Pattern Recognition, Oxford University Press. Oxford (1995)
31. Malhotra, P., Vig, L., Shroff, G., Agarwal, P.: Long Short Term Memory Networks for Anomaly Detection in Time Series. In: 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, Bruges, Belgium, April 22-24 April. (2015)
32. Eilers, D., Köpp, C., Gleue, C., Breitner, M.: It’s Not a Bug, It’s a Feature: How Visual Model Evaluation Can Help to Incorporate Human Domain Knowledge in Data Science. In: Proceedings of the Thirty Eighth International Conference on Information Systems (ICIS), Seoul, South Korea, December 10–13. (2017)
33. Carr, D.B., Littlefield, R.J., Nicholson, W.L., Littlefield, J.S.: Scatterplot Matrix Techniques for Large N. J. Am. Stat. Assoc. 82 (398), 424–436 (1987)