Predictive Maintenance Service Powered by Machine Learning and Big Data

Dr. Tapio Torikka and Dr. Steffen Haack

Bosch Rexroth AG, Zum Eisengießer 1, Lohr am Main, Germany

E-Mail: tapio.torikka@boschrexroth.de , steffen.haack@boschrexroth.de

We present a service for Predictive Maintenance in which existing machine data from control units or data from retrofitted sensors can be acquired from industrial machines by various gateway solutions. These gateways preprocess the data onsite and transmit it securely to a cloud-based Big Data system without impacting the production process of the industrial machine. Additional servers run Machine Learning algorithms to analyze the incoming data and generate data-based models representing the machine behavior. Results from existing applications show that significant benefits can be created for our customers and that Machine Learning algorithms demonstrate superhuman performance in detecting anomalous machine behavior.

Keywords: big data, digitalization, connectivity, machine learning, predictive maintenance
Target audience: mobile hydraulics, industrial hydraulics, industrial internet

1 Predictive Maintenance

The goal of Predictive Maintenance is to increase the availability of machines by recognizing failures or problems before significant damage is cause to the equipment. With advance information of developing failures, the machine operator is able to prepare required maintenance activities well in advance and avoid unplanned downtime. This is particularly beneficial on mission critical machines in 24/7 operation where downtime is very costly. Predictive Maintenance builds up on Condition Monitoring practises, traditionally carried out using threshold monitoring based on signal analysis: Thresholds for alarms levels are set on individual sensor signals. Also, predefined rules can be used for more complex systems. The interpretation of the data rests on human experts. A Predictive Maintenance system will add an analytics component to assist a human expert in detecting possible failure patterns and maintenance recommendations. This increases the level of automation and accuracy of the predictions. The Predictive Maintenance service presented in this paper uses Machine Learning to reach the required prediction accuracy and to provide an easily observable metric for decision making by human experts. Advantages of this approach, a high prediction accuracy and the reduction of manual setup work, have been demonstrated in academia /3/. The requirement of such an approach is a large volume of data in order for Machine Learning algorithms to learn typical failure patterns. Modern Big Data IT systems are capable of handling huge data volumes and although adoption in traditional industry is not yet common, this is commonplace in IT industry especially among popular cloud services like Facebook and Twitter. At the same time, advances in Machine Learning methods and computing, especially GPU computing, has made the required computing resources available. The platform developed by Rexroth, ODiN, uses these technologies to collect and analyze data from a growing worldwide network of connected machines.

2 Machine Learning

Machine Learning is a field of Artificial Intelligence in which an algorithm constructs a model based on input data in order to make predictions. No explicitly programmed instructions are used to create the model, which enables problems to be solved where limited or no domain knowledge is available. The flow chart in figure 1 shows the process of training and applying a Machine Learning algorithm. After acquisition of data from a selected data source it has to be preprocessed, e.g. by scaling the data and imputation of missing values. After preprocessing a Feature Extraction step extracts significant information from the data in order to improve the following training step. During the training phase a learning algorithm is used to iteratively adjust the internal parameters of the Machine Learning model to improve itself. Once a desired accuracy or a preset training time has been reached, the model can be saved and later used to make predictions with unknown data. /1/

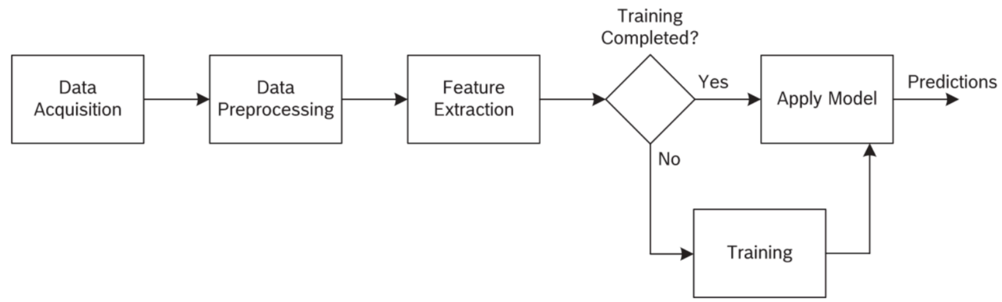

Figure 1: Data flow in a Machine Learning process

In Predictive Maintenance, the goal of a trained Machine Learning model is to predict failures, i.e. assign a known failure category to a presented data set. Two distinct types of learning are of particular interest in this application: Unsupervised and Supervised learning.

2.1 Unsupervised Learning

When no example data of failures is available, unsupervised learning can be used for anomaly detection, i.e. to detect unusual patterns in the data by comparing newly acquired data to previously observed reference data. These patterns can later be labelled and used as input data for a classifier trained via supervised learning.

2.2 Supervised Learning

Supervised learning is preferred when labelled data from known failure patterns is available either through experiments or collection of field data. Also, unsupervised learning can be used for data exploration on unknown data sets to assist the labelling process. The data is labelled with the target category for prediction and used in the training process. After training, a classifier will attempt to classify unknown data to these predefined categories.

2.3 Advantages of Machine Learning

Many machines are purpose built for a specific application and therefore exhibit a unique behavior. Additionally, environmental effects like temperature variations, noise, vibrations, etc. influence the sensor signals. To make matters worse, the data volumes produced by industrial machines are often very large. One reason for this is the need for high frequency measurements due to highly dynamic production cycles of the machine. These characteristics make the assessment of the health state of the machine based on manual inspection of individual sensor signals or simple rules next to impossible. For a Machine Learning algorithm such characteristics do not pose a problem. Learning patterns in the data is possible as long as there is enough data for the algorithms to work with. An example illustrates the capability of this approach: A Machine Learning algorithm reaches a prediction accuracy of over 95% predicting failures on an axial piston pump /3/. The data was collected on a series of test bench experiments by collecting vibration data from accelerometers mounted on a pump. Measurements were carried out on a Good pump as well as after introduction of built in failures. Additional measurements were conducted by running the pump out of specifications. The high variation in the vibration signal which is caused by a different setup of the test bench and maintenance work was taken into account by repeating the measurements after adjustments in the piping of the test bench and after performing a series of simulated maintenance steps on the pump. When the data is presented to a human expert for manual classification, the prediction accuracy drops to 43%.

3 Big Data

Regardless of the definition of the term Big Data based on the data itself (e.g. 3 V:s, 4 V:s), efficiently handling a data set exceeding the storage capacity of a single computer requires the use of a distributed computing framework like Apache Hadoop /2/. In Hadoop each node in a cluster of servers contains storage and computation resources and jobs on the complete data set can be performed using the combined computation and storage capacity of the cluster. Access to the data at the lowest level is provided by a file system, HDFS (Hadoop Distributed File System). Minimal structuring of the data is needed when stored in HDFS, which allows various types of data to be stored as collected from data sources (text files, images, measurement data, logs etc.). Various additional modules in the Hadoop framework allow for data to be structured and processed as required by the use case, e.g. Spark (computation), Hive (data warehouse) and HBase (data structuring into tables). Hadoop contains a management system, YARN, to manage resources and jobs on the cluster. Data on the cluster is typically replicated on several servers/racks. In case of a problem, e.g. hard drive failure, the copies of the data on other hard drives will ensure that no data is lost. This is all managed automatically and is an essential feature for clusters containing thousands of servers. The following chapter shows the architecture of such a system.

4 ODiN Platform

4.1 Architecture

Rexroth has developed a cloud-based Big Data platform, ODiN, in which existing machine data from control units or data from retrofitted sensors can be acquired from industrial machines by various gateway solutions. These gateways collect the data parallel to the control unit of the machine, preprocess the data onsite, prepare it for transmission and forward it to the centralized ODiN Big Data system. The data collection does not impact the production process of the industrial machine since no control functions are performed by the gateway. A secure https data transmission over the Internet is realized in parallel to the existing machine network typically via a cell phone network in order to avoid interference with the onsite network infrastructure.

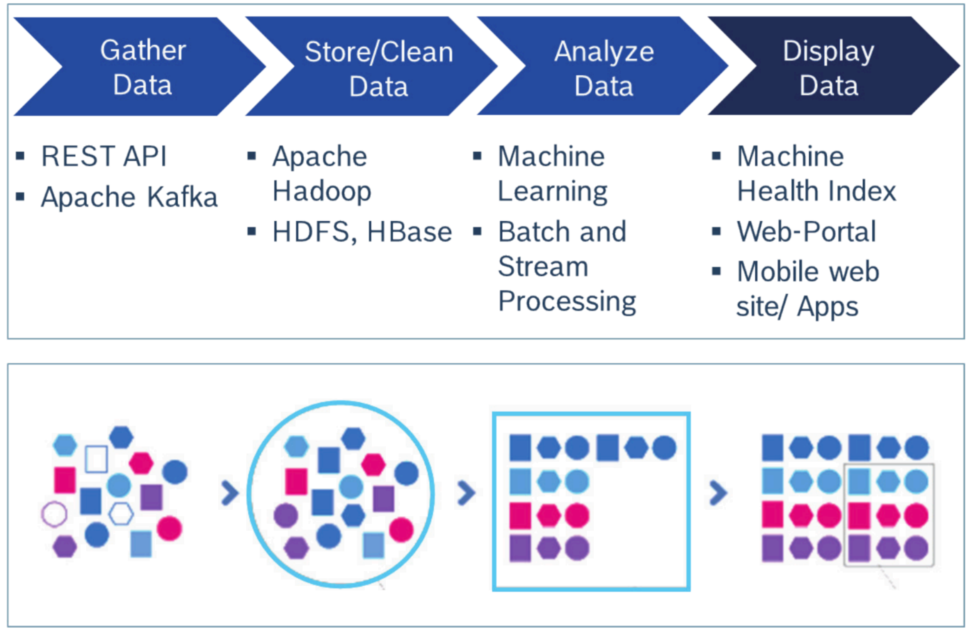

Figure 2: Data processing in ODiN platform

Figure 2 illustrates the data processing within the ODiN system. The cloud-based Big Data system consists of a cluster of servers for hosting data from all connected machines worldwide. The data ingestion is carried out by servers running Apache Kafka receiving data from the gateways on machines via a REST API. These servers feed the data to a Hadoop system where data is stored in HDFS. HBase stores the data in a more structured form for easier access by visualization and analytics processes.

The system implements the lambda architecture consisting of a stream and batch processing part. Stream processing uses Apache Spark and carries out simple statistical calculations (e.g. mean) directly on the input data stream. The results are pushed towards a web portal for fast data visualization purposes. Tasks requiring more computing power, e.g. Machine Learning, are carried out by the batch processing part. Dedicated analytics servers access data in HBase or HDFS, perform predictions on the data and output the results back into these data stores.

The web portal visualizes the results of both stream and batch processing. Quick status information is available on a dashboard and a charting tool gives the users a more detailed access to data and a look at long term trends. The web portal also includes tools for maintenance report generation, accessing equipment information including sensors, a message feed and gateway connection status. Administrators have the possibility to manage users and organizations. A mobile website for handheld devices contains basic information (dashboard and messages) for a quick status check.

4.2 Machine Health Index

The status of machines is displayed using a dedicated metric, the Machine Health Index. Where traditional signal analysis with threshold monitoring involving dozens or hundreds of sensor signals from each connected machine would be laborious, the Machine Learning algorithms condense all input signals from a component or a larger system to a single value, thus simplifying the monitoring of applications. The Machine Health Index quantifies the behavior of the machine or component ranging from 0 to 100 (0 = poor health, high failure probability, 100 = good health, low failure probability). Figure 3 shows a large number of input sensor signals available for monitoring in a traditional approach. The condensed information out of these signals in a Machine Health Index is shown in figure 4.

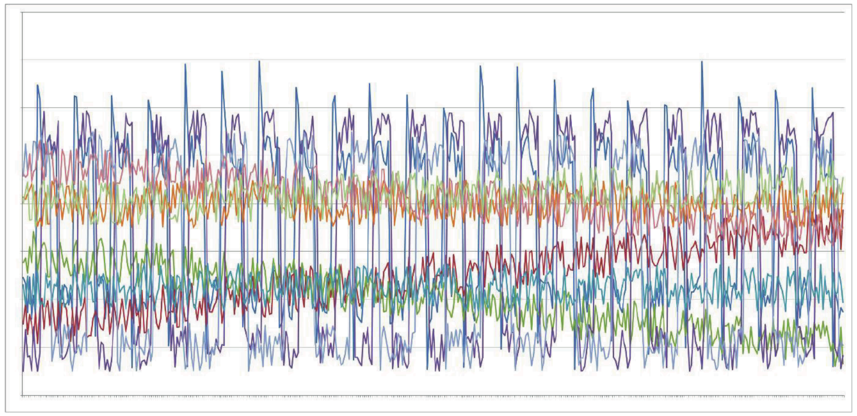

Figure 3: Example sensor signals from industrial equipment

Figure 4: Example of a Machine Health Index with a failure developing over time

In the case of unavailable failure data, the Machine Health Index is based on known failure free reference data from the machine in question. This data is typically collected after installation of the system. In a retrofitted machine a fitness check ensures that the equipment is running nominally. An unsupervised Machine Learning algorithm will then compare incoming data to the reference state, the Machine Health Index will represent the deviation from the reference state. On large machines with a high number of sensor signals the Machine Health Index can be calculated for each component in the machine enabling a more detailed overview of the health of the machine. If known failure data is available, a more precise calculation of the Machine Health Index including the probabilities for the known failures with supervised learning is possible.

4.3 Predictive Maintenance Service

The presented ODiN platform is the basis for a Predictive Maintenance service. This service is offered to customers as a complete package containing data collection, analysis, status reporting and implementation of service actions when necessary. The following chapter shows examples from implemented real life applications.

5 Examples

A break disc manufacturer is using the Predictive Maintenance service presented in the last chapter to increase the availability of a mission critical hydraulic power pack in the production process. The Predictive Maintenance solution would complement an existing scheduled maintenance to further reduce the risk of unplanned downtime. Since no redundancy is built into the power unit, a breakdown of the pump will lead to downtime of the power unit and affect the production process. The power unit was retrofitted with sensors and a gateway collecting data and sending it into the cloud for analysis. After a fitness check the system went live and a Machine Health Index is calculated for each component in the power unit (pump, oil, filters, cooling system). Figure 5 shows an excerpt or the data showing the Machine Health Index for the pump developing over time. After running 70% longer than a scheduled maintenance would have allowed, the behaviour of the pump suddenly changed as indicated by the significant drop in the Machine Health Index. The customer was contacted after a detailed diagnosis by human experts and a replacement of the pump was scheduled a week later. Production continued normally until the replacement. A bearing failure was detected at the inspection of the replaced pump. Due to the advance warning the customer was able to improve organization of the maintenance stop and subsequently cut the maintenance time by 50%.

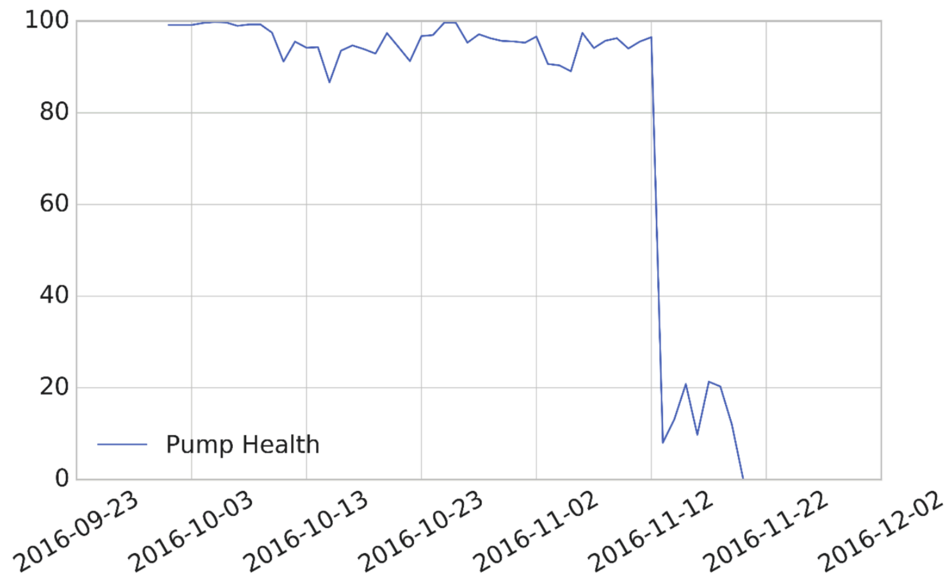

Figure 5: A sudden decrease of the Machine Health Index in early November indicated a significant deterioration of the pump. The customer was immediately notified and maintenance work could be carried out a week later.

A sudden change in machine behaviour like in the example above requires maintenance actions to be taken fast. An example of a slowly developing failure can be shown in figure 6. There is a slowly developing trend in the Machine Health Index several weeks before a serious problem arises. This enables machine operators to be notified well in advance. In the example below, the pump is deteriorating slowly and as soon as a more significant drop to 60 in the Machine Health Index occurred, the customer was notified to schedule maintenance. This time, maintenance could be scheduled three weeks later. During these three weeks the Machine Health Index kept deteriorating but the pump was still operating nominally. This time a less dramatic failure was detected at inspection after pump replacement. However, a failure process was ongoing and the pump would have failed in the near future.

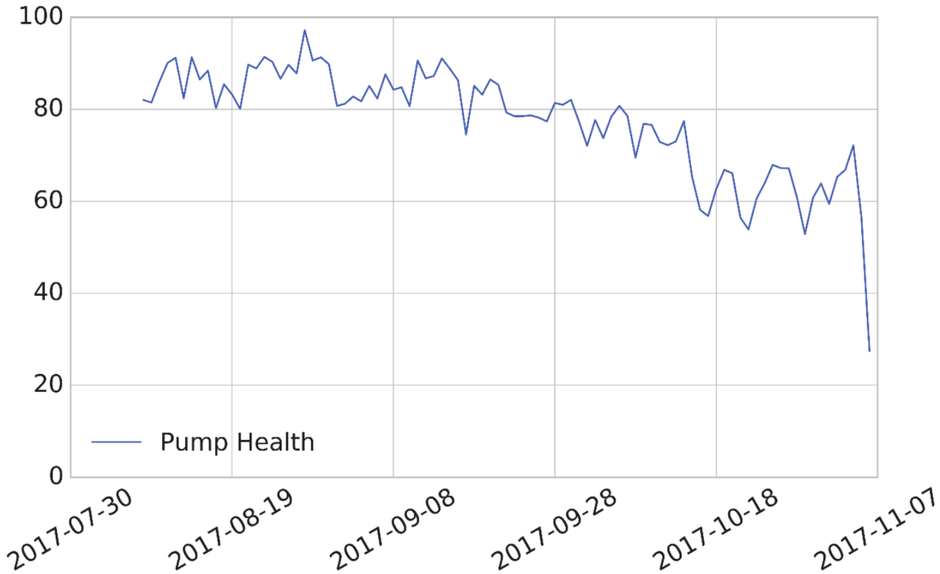

Figure 6: Health of the pump over time. A clear downward trend can be observed. The customer was notified in the middle of October to schedule maintenance. Maintenance was carried out three weeks later.

In the examples above, the Machine Health Index serves as the first indicator of an anomaly in the application and removes the need for humans to continuously analyse the incoming data. Human experts are still needed when an anomaly occurs to decide the correct action to be taken.

With an increasing number of machines connected to the ODiN system the data store is growing rapidly and generating the basis for more detailed failure diagnosis further automating the analytics process.

6 Conclusions

Maximizing availability on industrial machines with Predictive Maintenance requires a higher degree of automation for data analysis when compared to traditional condition monitoring systems. The examples presented in this paper show that Machine Learning algorithms provide this capability by accurately predicting patterns that represent sudden as well as slowly developing failure processes in fluid power equipment without significant human involvement. The required Big Data system containing data from all Predictive Maintenance applications worldwide can be implemented using state of the art software solutions.

When included in a maintenance contract, this solution can significantly reduce downtime on customer machines as demonstrated on real life examples. Due to advance planning, savings of up to 50% in maintenance costs could be achieved. Additionally, lifetime utilization of equipment can be significantly increased in applications where scheduled maintenance is currently used (by up to 70%). Once in place, the data in such a platform can be used by data scientists for the development of additional data based services for end customers as well as component and system manufacturers further increasing productivity.

References

/1/ Duda, H., Hart, P., Stork, D., *Pattern Classification 2nd Edition*, Wiley, 2000.

/2/ N, N., *Apache Hadoop*, http://hadoop.apache.org/, visited on December 12, 2017.

/3/ Torikka, T., *Bewertung von Analyseverfahren zur Zustandsüberwachung einer Axialkolbenpumpe*, Shaker Verlag, Reihe Fluidtechnik, Aachen, Germany, ISBN 978-3844003420, 2011.