Proceedings of 8th Transport Research Arena TRA 2020, April 27-30, 2020, Helsinki, Finland

Detecting anomalous behavior of railway switches under real operation conditions: workflow and implementation

Daniela Narezo Guzman[a]*, Edin Hadzic[b], Henk Samson[b], Serge van den Broek[b], Jörn C. Groos[a]

[a]*German Aerospace Center (DLR), Institute of Transportation Systems, Berlin/Braunschweig, Germany*
[b]*Strukton Rail, 3600BA Maarssen (Utrecht), Netherlands*

Abstract

Railway switches are crucial assets since they enable trains to change tracks without stopping. Larger parts of the infrastructure are compromised when certain switches fail. Regular inspections, maintenance and repairs are required to increase switch reliability, making them costly assets. Monitoring systems help determining the condition of assets. Nowadays nearly thousand switches in the Netherlands are remotely monitored by Strukton Rail. The current version of this monitoring system has helped to identify degrading and failing switches, but it also generates false alarms. There is room for improvement in how the monitoring system supports asset managers in making decisions regarding the asset. Here, we present a workflow that exploits switch monitoring data under real operation conditions. The running workflow implements a machine-learning model for automatic detection of anomalous switch functioning. Models for predicting switch degradation and failure evolution, and for identifying failure types are under development and remain to be implemented.

Keywords: intelligent systems; infrastructure; railway engineering; algorithms; asset management

* Corresponding author. Tel.: +49-30-67055-9650;
E-mail address: Daniela.NarezoGuzman@dlr.de

1. Nomenclature

SR	Strukton Rail
CC	Current curve(s)
SCC	Strukton Control Centre
T^2	Hotelling's parameter
SPE	Square Prediction Error
FMECA	Failure Mode, Effect and Criticality Analysis
ARIMA	Autoregressive Integrated Moving Average
TRL	Technology Readiness Level

2. Introduction

Condition monitoring is the continuous update of parameters that characterize the health state of a system or component, and is made available through sensors. Condition monitoring may enable substituting preventive (i.e. scheduled) maintenance by condition-based maintenance. The latter is a more efficient approach than the former (see Vinberg et al. (2018)) for repairing and maintaining, since it leads to improved reliability, availability, safety and reduced costs. Condition monitoring is of great interest for the railway system and is an active field of research; it has been applied for the detection of railway track irregularities and faults by Tsunashima (2019), wheel and wheelset defects by Alemi et al. (2016), vehicle dynamics faults in Ngigi et al. (2012) and switch faults in Silmon and Roberts (2010).

SR recently developed an application for collecting, handling and managing warnings and alarms coming from railway assets monitored by different systems. This includes POSS®, the system that monitors the condition of switches via sensors that acquire point machine (or switch engine) current during each turnover. This signal (referred here to as CC) provides a useful representation of the energy required by the switch engine to relocate the switch-blades from one end-position to the other according to INNOTRACK (2009). The blades end-position determines the tracks that a transiting vehicle takes. In its current version POSS® provides warnings and alarms when CC exceed manually-selected reference thresholds (see Narezo Guzman et al. (2018b) for further details). Multiple warnings and alarms may be triggered when anomalous behaviour is detected, reflecting that there is a problem with the asset. However, warnings and alarms might also be triggered due to ill-defined reference thresholds. Moreover warnings and alarms of an asset are clustered according to their time of occurrence. An expert at the SCC assesses these clusters and creates notifications for work i.e. a list of follow-up maintenance and repair actions. By consulting multiple information sources, such as video inspections, these experts also provide long-term maintenance recommendations. In the past, many problems were detected and solved using this method, often even preventing upcoming failures. However, many of the clusters are not worth of following up since they are a collection of low priority and/or false alarms. The analysis of clusters and the creation of notifications is a time consuming job. Clearly, there is a need for a more accurate warning system which minimizes unnecessary work, reduces costs and assures asset availability.

In a real operational environment there are thousands of assets being monitored and maintained. Each of them is unique and has its own degradation curve, which is partially determined by local parameters. In the Netherlands the execution of repair and maintenance actions needed for preventing further degradation is subject to very tight schedules and the strong incentive to reduce maintenance slots. As an organization responsible for the maintenance and availability of railway assets, SR would greatly benefit from models exploiting the monitored data, which can trigger accurate alarms, identify the failure cause, detect asset degradation and predict upcoming problems at an early stage. The goal of this research is to successfully develop such models, as well as to implement the necessary IT infrastructure for supporting the workflow under real operation conditions. These efforts can be considered a significant step towards condition-based and predictive maintenance. The outcome of the models should enable the operators and analysts at the SCC (the end users) to visualise asset current and future conditions in an intuitive way, allow them to easily add or remove switches from the workflow, as well as to train the models efficiently and reliably.

This paper reports on a fully automated workflow for monitoring and assessing the condition of switches in a real operation environment. The workflow constituent parts, their dependencies and the output are described in section 3.1. Overviews of currently running machine-learning models embedded in the workflow, as well as of the models under development to be implemented in the workflow, are provided in sections 3.2 and 3.3, respectively. In section 4 closing remarks are provided.

3. Workflow under real operation conditions

3.1. Description of workflow

Fig. 1 shows a schematic representation of the workflow including all input information accessing it (represented by three boxes on the left side). These data come from different sources and include:

- POSS® monitored CC acquired as raw data (in hexadecimal system).
- FMECA system and domain knowledge.
- Additional data: executed and planned maintenance actions, repair of asset failures, asset information (e.g. location, switch characteristics, etc.), switch usage and load (e.g. total weight of passing trains).
- Weather data: weather variables with high temporal and spatial resolution that have an impact on switch performance such as temperature, snow, rain and humidity.

The server is configured to run the workflow every hour. In the workflow, raw monitoring data is pre-processed in the production environment and stored in a SQL database on a SR server. The data in the database is then further processed in the demonstrator, computing features derived from the acquired CC (such as CC duration, maximum value, etc.) measured for every switch-blades repositioning. The machine-learning model for anomaly detection exploits the derived features as well as weather information gathered at the time the repositioning takes place. Once the model is trained with historical records of the aforementioned information/data, it assesses every CC acquired in the previous hour and triggers an alarm in case anomalous behaviour is detected. Several consecutive anomalies for one switch can indicate an evolving/sustained failure. If the failure is detected at an early stage, the problem can be addressed and the switch repaired before its functionality is completely compromised. Moreover, some failures (e.g. a rusting gear-box) follow a clear degradation trend that can be captured by a regressive model, making the evolution of the failure, in principle, predictable. Such failures can be prevented provided they are detected early enough. Besides failure detection and the prediction of future switch condition, it is of great interest to identify the failure type and provide an automated diagnosis. Given the lack of adequate training data required for purely supervised machine-learning approaches, the diagnostic model under development takes a hybrid approach. It consists of an extensive modular Bayesian network (see Koller and Pfeffer (1997)) whose structure is based on expert knowledge (for instance collected via FMECA or similar methods) and additional information from literature.

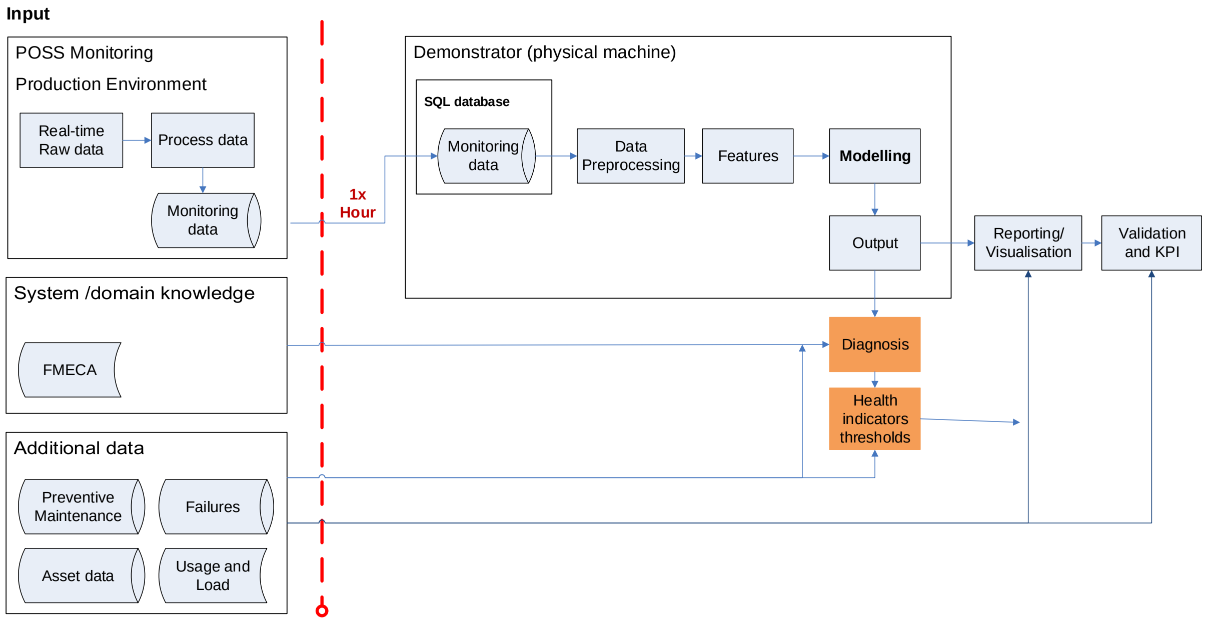

Fig. 1 Schematic representation of the workflow under real operation conditions.

With regard to the anomaly detection results, the model output is reported and visualized as time series of the anomaly score as well as of the most significant features, where a data point represents a CC. The graphical representation of the switch functional condition is made available on the server, and is accompanied by historical and planned maintenance actions, which for instance, can be helpful in understanding sudden changes in the switch behaviour. Moreover, further development of the graphical representation is planned to include time series projections into the future resulting from predictive modelling. In this workflow, on the one hand, the visualization of results, together with alarms and failure diagnosis are analysed and provide support to specialists at the SCC in decision making. On the other hand, annotated lists of real anomalies and diagnostic results (ground truth) will become available over time, serving for validation of the models performance and assessment of informative KPIs.

The workflow presented in this paper is required in order to integrate the anomaly detection, diagnosis and predictive models into the daily practice at the SCC.

3.2. Current implementation and results

In the current demonstrator pre-processed data from hundreds of monitored switches is made available every hour. More than eighty features have been defined and twenty of them are currently being computed for all monitored switches in the demonstrator. In order to consider or not a given switch in the workflow, its ID has to be included in or excluded from a configuration file. Switches included so far are characterized by a relatively high frequency of failures and preferentially by long-term degradation (not all switches present a clear degradation process).

A first version of the anomaly detection model was applied to historical data of seven double-slip switches gathered over more than five years; the model was shown to be temperature-robust in Narezo Guzman et al. (2018b). A more temperature-sensitive version of the model has since been developed and is now embedded in the workflow. This version of the model described in Narezo Guzman et al. (2019a) is applied every hour to eight (out of the twenty) computed features derived for tens of switches. The most useful model output is the anomaly score, which is computed for every CC. The anomaly score ranges from -0.5 (for data points that are extreme outliers) to +0.5 (for explicitly normal/average data points), as described by Sharova (2018). The latest version of the model was applied to the same historical data as in Narezo Guzman et al. (2018b). The results in Narezo Guzman et al. (2019b) showed that, provided the model is trained with a set of CC that present a consistent and narrowly distributed response to small changes in temperature, the sensitivity, precision and specificity of the model for detecting anomalies (based on expert assessment of nearly 600 CC) are high: 0.95, 0.96 and 0.85, respectively.

Anomaly detection results obtained in the workflow are displayed on a user interface together with other additional information (e.g. maintenance actions and switch failures). Fig. 2 shows one of the windows on this interface, where the anomaly score of CC measured for switch S4730 is plotted as a function of time. This window provides the user with a view of the switch behaviour over time, making evident any trend in the anomaly score. The user can modify the time-window of the displayed timeseries. The threshold value below which a CC is identified as anomalous is user-defined and can be modified at any time. In this example the threshold is set to around -0.05. Several anomalies can be identified within the time ranging between Feb. 1st 2018 and Oct. 1st 2019. Special attention is drawn to the end of year 2018, where over the course of a couple of weeks the anomaly score increasingly became negative. In fact a failure was registered on Dec. 12th but it is evident that it was not repaired right away, which is confirmed by the maintenance log, as anomalies continued to be detected for more than a week. Moreover in the first half of Jan. 2019 a maintenance action took place, likely restoring the normal switch functioning, as it is observed from the anomaly score timeseries.

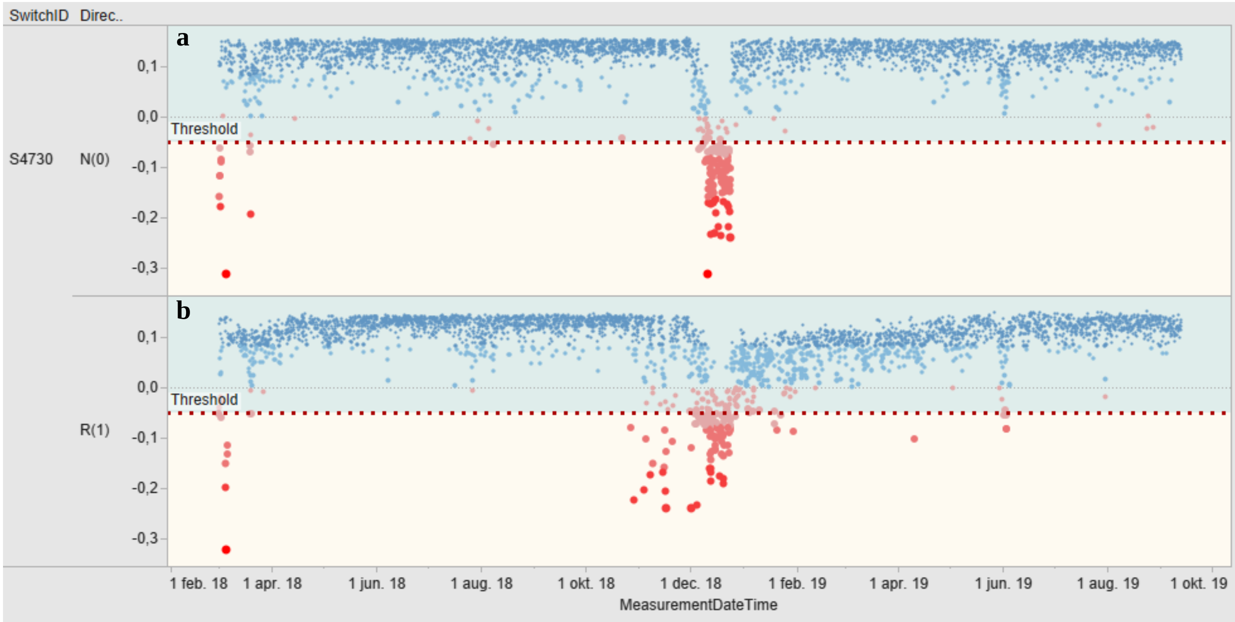

Fig. 2 User interface displays the anomaly score of CC as a function of time for each blades moving direction separately: (a) N(0) and (b) R(1). The user-defined threshold (horizontal dotted line) separates anomalous from normal data points.

Another window of the user interface, called operational view, is shown in Fig. 3. In it the user has a more detailed overview of recent anomalies detected for the selected switch (S4730); it includes temperature at the time the CC was measured as well as corresponding feature values. Furthermore, the user interface is linked to POSS®, which can easily be accessed with a single click for visualizing the corresponding CC and any other information stored in that system.

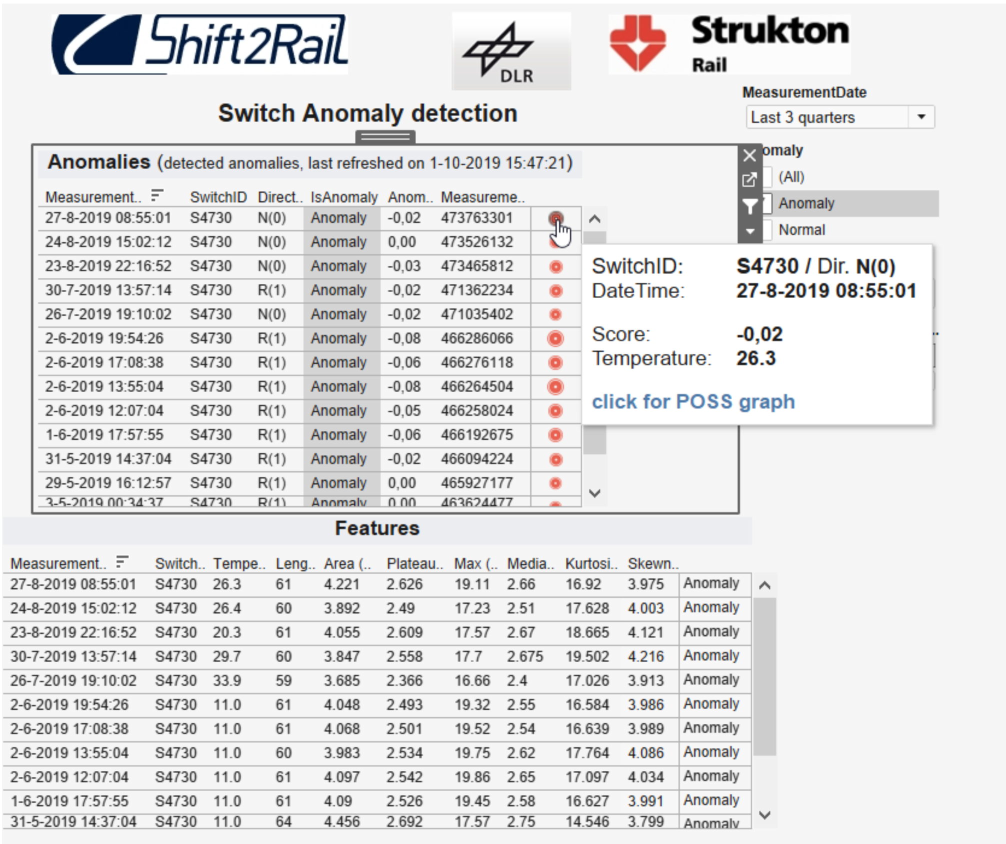

Measurement..	SwitchID	Direct..	IsAnomaly	Anom..	Measureme..
27-8-2019 08:55:01	S4730	N(0)	Anomaly	-0,02	473763301
24-8-2019 15:02:12	S4730	N(0)	Anomaly	0,00	473526132
23-8-2019 22:16:52	S4730	N(0)	Anomaly	-0,03	473465812
30-7-2019 13:57:14	S4730	R(1)	Anomaly	-0,02	471362234
26-7-2019 19:10:02	S4730	N(0)	Anomaly	-0,02	471035402
2-6-2019 19:54:26	S4730	R(1)	Anomaly	-0,08	466286066
2-6-2019 17:08:38	S4730	R(1)	Anomaly	-0,06	466276118
2-6-2019 13:55:04	S4730	R(1)	Anomaly	-0,08	466264504
2-6-2019 12:07:04	S4730	R(1)	Anomaly	-0,05	466258024
1-6-2019 17:57:55	S4730	R(1)	Anomaly	-0,06	466192675
31-5-2019 14:37:04	S4730	R(1)	Anomaly	-0,02	466094224
29-5-2019 16:12:57	S4730	R(1)	Anomaly	0,00	465927177
3-5-2019 00:34:37	S4730	R(1)	Anomaly	0,00	463624477

Measurement..	Switch..	Tempe..	Leng..	Area (..	Plateau..	Max (..	Media..	Kurtosi..	Skew..	
27-8-2019 08:55:01	S4730	26.3	61	4.221	2.626	19.11	2.66	16.92	3.975	Anomaly
24-8-2019 15:02:12	S4730	26.4	60	3.892	2.49	17.23	2.51	17.628	4.003	Anomaly
23-8-2019 22:16:52	S4730	20.3	61	4.055	2.609	17.57	2.67	18.665	4.121	Anomaly
30-7-2019 13:57:14	S4730	29.7	60	3.847	2.558	17.7	2.675	19.502	4.216	Anomaly
26-7-2019 19:10:02	S4730	33.9	59	3.685	2.366	16.66	2.4	17.026	3.913	Anomaly
2-6-2019 19:54:26	S4730	11.0	61	4.048	2.493	19.32	2.55	16.584	3.986	Anomaly
2-6-2019 17:08:38	S4730	11.0	61	4.068	2.501	19.52	2.54	16.639	3.989	Anomaly
2-6-2019 13:55:04	S4730	11.0	60	3.983	2.534	19.75	2.62	17.764	4.086	Anomaly
2-6-2019 12:07:04	S4730	11.0	61	4.097	2.542	19.86	2.65	17.097	4.034	Anomaly
1-6-2019 17:57:55	S4730	11.0	61	4.09	2.526	19.45	2.58	16.627	3.991	Anomaly
31-5-2019 14:37:04	S4730	11.0	64	4.456	2.692	17.57	2.75	14.546	3.799	Anomaly

Fig. 3 User interface displays a list of recent anomalies detected for switch S4730. Each anomaly is accompanied by date/time, direction, anomaly score and measurement ID (columns in top table). The interface is linked with POSS®: by clicking on the red button the user can open the corresponding CC in that system. Feature values derived from anomalous CC are listed in the bottom table.

3.3. Future developments

The approach under consideration for predicting the evolution of switch incipient failures i.e. failures which develop gradually over a period of time, is ARIMA modelling of timeseries (see Aggarwal (2017)). Here an ARIMA model is separately developed for the timeseries of log(T^2) and log(SPE), where a data point in each of these timeseries represents a CC. T^2 and SPE are parameters that summarize the information contained in the features (see Narezo Guzman et al. (2018a) and Yue and Qin (2001) for more details), and in our approach they are closely related to the anomaly score since they define the parameter space on which the anomaly detection algorithm (Isolation Forest) is applied. ARIMA model parameters must be computed for the logarithm of both T^2 and SPE separately, and for single incipient failure events given that every case is different. This involves training the model with the degrading trend observed/detected in the timeseries that precedes complete failure. First results found in Narezo Guzman et al. (2019a) are encouraging however further research is required before the model can be validated and implemented in a real condition environment. Several aspects are still open, including the definition of rules that initiate model training (e.g. how many consecutive anomalies indicate a persistent failure mode that can be modelled), the applicability and generality of ARIMA models for predicting switch failures evolution, the increase in prediction confidence in the presence of more accurate weather variables measurements, etc.

With regard to the diagnostic model presented in Neumann and Narezo Guzman (2019), the basic structure has been derived from current expert knowledge. The novelty of the network structure is that it consists of dedicated modules that can for instance be switched off or adjusted, allowing to consider various characteristics concerning the specific construction of a given switch. Moreover, the model principally incorporates the probabilities that certain weather conditions or other influencing factors induce on the observation of different failures types. Further investigation is necessary in order to calibrate the model, identify other relevant external factors and possible causes of switch failure, and establish a link between the diagnostic model and the switch anomaly detection output, e.g. by considering features as evidences in a (extended) version of the Bayesian network model.

All models presented here may benefit from considering a larger number of features than done until now. Therefore more features will be computed and exploited by the anomaly detection, predictive and diagnostic models for railway switches in the future.

4. Final remarks and outlook

The described workflow implements the previously validated anomaly detection model in an operational environment; and it is the base for running the - currently under development - predictive and diagnostic models in the future. This work represents an essential step towards condition-based and predictive maintenance enabled by machine-learning approaches, since it provides the results to the end users (domain experts) via an interface in (nearly) real-time. The TRL (according to the definition by NASA) of the switch anomaly detection model embedded in the workflow is that between technology development and demonstration, i.e. TRL between 4 and 5. Nevertheless, many challenges remain and research is necessary with regard to the workflow in order to include all mentioned models and reach TRL 9. These efforts will be continued within the IN2SMART-2 project.

Acknowledgements

Research conducted within the Shift2Rail project IN2SMART (EU Horizon 2020 research and innovation program, grant agreement 730569).

5. References

Aggarwal, C. C.: Outlier analysis, Second edition, Springer International Publishing Switzerland, op. 2017.

Alemi, A., Corman, F., and Lodewijks, G.: Condition monitoring approaches for the detection of railway wheel defects, Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 231, 961–981, 2016.

INNOTRACK: Deliverable 3.3.2 - Available sensors for railway environments for condition monitoring.

Koller, D. and Pfeffer, A.: Object-Oriented Bayesian Networks, in: Proceedings of the Thirteenth Annual Conference on Uncertainty in

Artificial Intelligence, Providence, Rhode Island, 1-3.8.1997, 302–313, 1997.

Narezo Guzman, D., Hadzic, E., Baasch, B., Heusel, J., Neumann, T., Schrijver, G., Buursma, D., and Groos, J. C.: Anomaly Detection and Forecasting Methods Applied to Point Machine Monitoring Data for Prevention of Switch Failures, in: Proceedings of the 32nd International Congress and Exhibition on Condition Monitoring and Diagnostic Engineering Management (COMADEM 2019), Huddersfield, 3-5 September, Springer, 2019a.

Narezo Guzman, D., Hadzic, E., and Groos, J. C.: Anomaly detection model based on switch monitoring data for condition-based maintenance, in: 12th World Congress on Railway Research, Tokyo, Japan, October 28 - November 1, https://www.sparkrail.org/Pages/SparkWelcome.aspx, 2019b.

Narezo Guzman, D., Hadzic, E., Schuil, R., Baars, E., and Groos, J. C.: Data-driven condition now- and forecasting of railway switches for improvement in the quality of railway transportation, in: Proceedings of the European Conference of the PHM Society, Kulkarni, C. S., and Tinga, T. (Eds.), 4th European Conference of the PHM Society, Utrecht, The Netherlands, July 3-6, 2018a.

Narezo Guzman, D., Hadzic, E., Schuil, R., Baars, E., and Groos, J. C.: Turning data driven condition now- and forecasting for railway switches into maintenance actions, in: Proceedings of 7th Transport Research Arena TRA 2018, Vienna, Austria, April 16-19, 2018b.

Neumann, T. and Narezo Guzman, D.: Bayesian network design for fault diagnostics of railway switches, in: 29th European Safety and Reliability Conference, Beer, M., and Zio, E. (Eds.), Hannover, Germany, September 22-26, Research Publishing Services, Singapore, 1117–1124, 2019.

Ngigi, R. W., Pislaru, C., Ball, A., and Gu, F.: Modern techniques for condition monitoring of railway vehicle dynamics, J. Phys.: Conf. Ser., 364, 012016, doi:10.1088/1742-6596/364/1/012016, 2012.

Sharova, E.: Unsupervised Anomaly Detection with Isolation Forest, PyData, London, U.K, April 28 2018.

Silmon, J. A. and Roberts, C.: Improving railway switch system reliability with innovative condition monitoring algorithms, Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, 224, 293–302, doi:10.1243/09544097JRRT313, 2010.

Tsunashima, H.: Condition Monitoring of Railway Tracks from Car-Body Vibration Using a Machine Learning Technique, Applied Sciences, 9, 2734, doi:10.3390/app9132734, 2019.

Vinberg, E. M., Martin, M., Firdaus, A. H., and Tang, Y.: Railway Applications of Condition Monitoring, Stockholm, 2018.

Yue, H. H. and Qin, S. J.: Reconstruction-Based Fault Identification Using a Combined Index, 40, 4403–4414, 2001.