DETECTION OF MAINTENANCE ACTIONS BASED ON RAIL WEAR ANALYSIS

Alexander Plotho
Graz University of Technology, Institute of Railway Engineering and Transport Economics, Austria

Abstract

Predictive maintenance is becoming increasingly important for identifying maintenance requirements at an early stage and making efficient use of pre-planned track assignments. The attractiveness of railway transport increased due to growing demand for passenger rail services, while that same demand has also reduced the availability of time windows for maintenance activities. The main problem is that rail maintenance activities have often not been sufficiently documented in Austria. For this reason, there is a need to develop methods to close these data gaps and enable robust, meaningful trend analyses. It is hypothesized that abrupt structural changes in the combined vertical and lateral wear signal correspond to executed rail maintenance actions, specifically rail surface treatments and rail replacements. A dual change-detection algorithm with regression-based post-filtering can identify and temporally/spatially localize these events with high precision. The methodology relies on time series representing the combined wear for each track curve element. Therefore, the preparation of the data involves implementing processes such as averaging, outlier detection, and filtering to identify structural changes in the wear signal. The results demonstrate high detection performance: structural changes can be reliably identified and classified. In addition to rail surface treatments, rail replacements can be detected both spatially and temporally with high accuracy. Consequently, the proposed methodology has demonstrated that undocumented maintenance actions can be detected. However, there are limitations. Maintenance actions cannot be fully detected in sections where wear is insignificant. Future studies on the rail surface signal could improve the detection performance. In conclusion, the methodology provides a robust foundation for the developing and implementing predictive maintenance strategies in the railway sector.

Keywords: maintenance, rail wear, data analysis, railway

1 Introduction

Maintaining the track is crucial to ensure the quality and safety of railway operations [1, 2]. This is becoming more important as rail transport is growing in popularity and train traffic is expected to increase in the future [3]. However, as the number of trains increases, the available maintenance windows, the periods during which work can be carried out on the tracks, are becoming shorter and less frequent [4]. This makes it necessary to plan maintenance measures more efficiently. This approach is known as predictive maintenance. While predictive maintenance is already established for certain track components, it has not yet been adopted to the same extent for rails. The aim of this work is therefore to contribute to predictive maintenance for rails. Forecasting models are important for estimating future maintenance requirements.

DOI: https://doi.org/10.5592/CO/CETRA.2026.1756

To specify maintenance times more precisely, reliable information on previous maintenance work is also required. However, this is where a key problem arises: in Austria, past maintenance measures have sometimes not been documented or have been documented inaccurately. For this reason, this work focuses on the automated identification of undocumented maintenance work, particularly rail surface treatments and rail replacement. Once these measures can be detected, a data-driven, demand-oriented maintenance strategy can be developed. This strategy will enable more targeted and efficient planning of future maintenance measures and will contribute to long-term quality assurance and operational safety. Time series are required to detect maintenance actions. These time series are derived from rail wear data. Maintenance actions can be identified by structural breaks in the time series, such as abrupt changes in the trend that cannot be explained by normal wear. To identify structural breaks more reliably, the time series is defined in terms of cumulative wear ("added wear"). Cumulative wear is calculated as the sum of height wear and lateral wear. Height wear describes the vertical reduction of the rail head measured along the rail axis [5]. It occurs mainly on the inner rail in curves [5]. Lateral wear describes the lateral reduction of the rail head measured at an angle of 45° to the rail axis [5, 6]. It is caused by wheel-flange contact and mainly affects the outer rail in curves [5]. Rail wear directly influences the need for rail replacement, but it has no immediate effect on rail surface treatments [7]. Nevertheless, wear is operationally relevant because it affects wheel–rail contact, running smoothness, and the loads on the overall track system [7]. To measure wear, a track recording car travels along Austria's main railway lines about four to six times per year. The car is equipped with several measurement systems, including a rail profile measurement system that is central to wear assessment. During measurement, the rail is illuminated by a laser operating at a frequency outside the visible range. Cameras then capture the rail, from the rail head to the rail foot, as grayscale images. These images are used to reconstruct the rail profile and compare it with a target profile. The resulting profile deviations are then used to compute wear values. The rail is scanned at 1 m intervals, with a maximum measurement frequency of 400 Hz. The measurement system and the resulting data form the basis for the method used in this study [8, 9].

2 Methodology

The method aims to identify maintenance measures as structural breaks in the time series of cumulative wear. To achieve this, the wear data are first processed and checked for outliers. A multi-stage analysis procedure is then applied to evaluate the time series. AI-assisted tools supported the implementation and writing. ChatGPT was used to support the programming of the algorithm. DeepL was used to revise formulations and to check spelling. All AI-generated outputs were reviewed and validated before use.

2.1 Data preparation

To use the wear data for further analysis, the data must first be processed. The first step is to define track sections. For this purpose, the track is divided into segments of 100–200 m. Sections that are too short increase computational effort, whereas sections that are too long smooth local effects and may mask relevant changes. Segmentation is based on curvature and on maintenance measures that have already been carried out. Sections are selected such that the curvature within a section is either constant (straight track or full curve) or changes systematically (transition curve). In addition, maintenance windows must always cover a section completely. After segmentation, a mean value is calculated for each section and measurement run. This yields a time series per section, into which the known maintenance measures are entered. The next step is to remove faulty measurement runs.

Measurement failures occur, for example, when an entire run or individual sections do not provide valid data. Such cases are detected using the standard deviation. If the standard deviation of a section in a given measurement run is zero, the corresponding measurement is classified as faulty and removed. In addition, values of zero or less are removed from individual measurements because they are not physically plausible. The resulting gaps are then filled by interpolation between the remaining values. Finally, outliers are removed using Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [10]. Due to large measurement fluctuations in the wear signal, only the strongest 1% of outliers are removed. The result is an adjusted time series, as shown in figure 1.

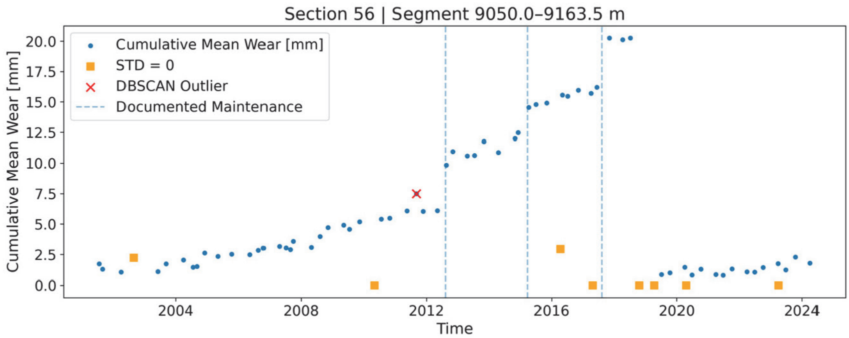

Figure 1 Cumulative mean wear overtime for a 113.5 m track section, documented maintenance and data points removed during preprocessing are indicated

Figure 1 shows a 113, 5 m track section. It plots cumulative mean wear values over time, where each value corresponds to one measurement run for this section. The documented maintenance measures around 2013, 2015 and 2017 are also indicated. In addition, the data points removed during data preparation, either as outliers or due to measurement errors, are shown.

2.2 Detection of maintenance actions

After data preparation, the actual analysis is performed. Depending on the availability of documented maintenance actions, a multi-stage procedure is applied. If only one or no maintenance actions are documented, Kernel Change Point Detection (KernelCPD) [11] is applied first to pre-structure the time series. KernelCPD identifies the largest structural breaks in the cumulative mean wear trend. In the next step, either for time series with sufficient documented maintenance measures or for time series pre-processed by KernelCPD, the Cross-Section- and RANSAC-Based-Algorithm (CRAB) [12] is used. It detects additional relevant structural breaks and thereby identifies further potential maintenance events, as presented in figure 2.

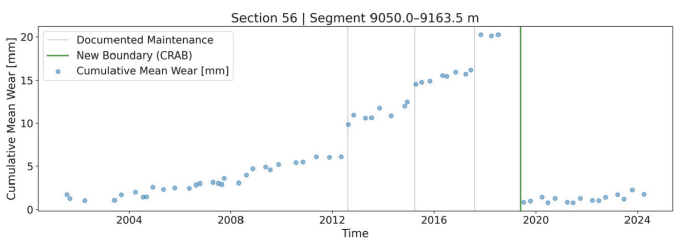

Figure 2 Cumulative mean wear overtime for the section from data preparation, CRAB detects an additional structural break in 2019

Figure 2 shows the results for the same section used in the data-preparation step. It again plots cumulative mean wear values over time. Three maintenance measure were documented for this section, meaning KernelCPD was not applied. However, the CRAB algorithm did detect an additional structural break in 2019. Once all structural breaks have been detected, each break must be validated. The validation serves two purposes: (i) to determine the maintenance time as accurately as possible and (ii) to retain only breaks that correspond to maintenance events that occurred. For validation, four data points before and four data points after each detected break are examined. Within this window, the largest adjacent change in the mean value is identified and analyzed in more detail. Specifically, the two affected measurement runs (before and after the suspected event) are compared. For both measurement runs, the median is calculated, due to possible measurement errors in the signal itself. If the median change is less than 15% of the larger of the two medians, the change is classified as too small. The 15% threshold is based on empirical considerations. In this case, the structural break is considered invalid and is removed. If the structural break is classified as valid, the measurement run after the validated event is used to classify the maintenance action. If the post-event values are higher than the pre-event values, the event is interpreted as a rail surface treatment. If they are lower, it is interpreted as a rail replacement. Finally, the post-event measurement run is checked for consistency across the entire section. A within-section change indicates that the maintenance action starts or ends within the section, as shown in figure 3.

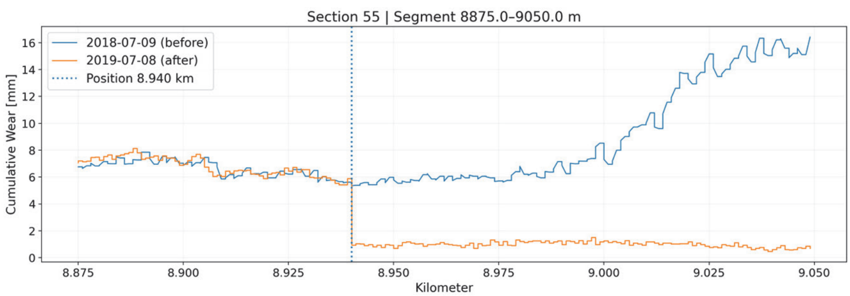

Figure 3 Comparison of two measurements runs for a 175 m section plotted against route kilometer, divergence from km 8.94 and lower wear in the later run indicate a rail replacement starting at this location

Figure 3 compares two measurement runs for a different section. The section is 175 m long, and cumulative wear is plotted against route kilometer. Two runs are shown between which a maintenance action likely occurred. In the first part of the section, the rate of wear is almost constant. This pattern changes at km 8.94: from this point onward, the two runs diverge. Because the newer measurement run shows lower wear values, which indicates a rail replacement starting at this location. Once all sections have been validated, sections that are close in time and location can be merged to represent continuous maintenance measures. This accounts for the fact that a measure detected in one section may appear shifted within the time interval between the two measurement runs used for detection. The merging step therefore combines only sections for which a joint maintenance measure is plausible. Finally, the detected events can be used to supplement and complete previously incomplete maintenance records.

3 Result

The algorithmic process developed in this work is the main result. It provides a robust and more precise basis for subsequent analyses because it integrates the required steps for systematically identifying maintenance events in the wear data. The process comprises (1) data preparation, including outlier detection and the removal of faulty measurement runs, (2) a multi-stage procedure for detecting structural breaks in the time series, and (3) validation of the detected breaks to ensure a plausible temporal and spatial assignment. Figure 4 provides a schematic overview of the detection process and summarizes the main steps. First, the data are prepared: height and lateral wear are combined into cumulative wear, the route is divided into sections, faulty measurement runs are removed, and outliers are corrected. Next, previously documented maintenance measures are entered into the time series. Depending on data availability, the subsequent analysis is performed in one or two stages. KernelCPD is used for pre-structuring, and the CRAB algorithm then identifies additional relevant structural breaks in the signal. All detected events are subsequently validated. Finally, validated events are merged across sections that are adjacent in time and location to represent continuous maintenance actions.

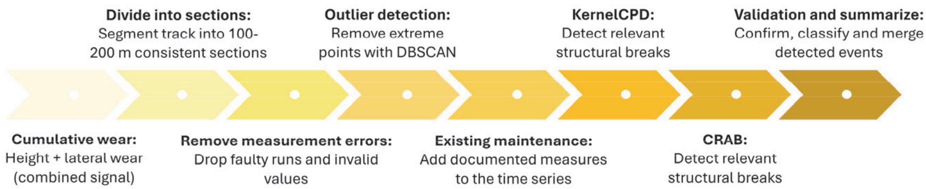

Figure 4 Overview of the detection process: data preparation, maintenance annotation, structural-break detection (KernelCPD, CRAB), validation, and merging of adjacent sections into continuous measures

4 Discussions

Since predictive maintenance for rails is still in its infancy, reliable information about past maintenance measures is essential. However, these measures have often been documented incompletely or inaccurately. The algorithm developed in this work supports this process by helping to systematically close these knowledge gaps. The results indicate that the algorithm can detect a large proportion of previously undocumented maintenance actions. The multi-stage validation also enables a precise temporal and spatial assignment of the detected events. At the same time, the analysis reveals limiting factors that prevent maintenance actions from always being identified unambiguously. Low in a section over a long period of time, structural breaks may be weak and can lead to false detections. In the present methodology, this risk is mitigated by a strict validation procedure that largely excludes potential false positives. As a result, detected maintenance actions are highly likely to be correct, even if not all events are captured. To address this limitation, the rail surface signal is expected to provide additional information and improve detection, especially in low-wear sections. Future work will therefore extend the approach by incorporating this signal, with the aim of detecting maintenance measures more completely and robustly. In summary, this work contributes to predictive maintenance for railways. The developed methods and the obtained results provide a basis for integrating predictive maintenance into future maintenance strategies and for planning measures in a more targeted, efficient, and data-driven way.

References

[1] Sharma, S., Cui, Y., He, Q., Mohammadi, R., Li, Z.: Data-driven optimization of railway maintenance for track geometry, Transportation Research Part C, 2018., DOI: 10.1016/j.trc.2018.02.019

[2] Keeping trains, track and infrastructure in tip-top condition, Europe's Rail, https://rail-research.europa.eu/latest-news/keeping-trains-track-and-infrastructure-in-tip-top-condition/, 30.01.2026.

[3] Rail passenger transport increased by 5.8% in 2024, Eurostat, https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251031-1, 30.01.2026.

[4] Ivina, D., Palmqvist, C.W.: Railway maintenance windows: Discrepancies between planning and practice in Sweden, Transportation Research Interdisciplinary Perspectives, 2023., DOI: 10.1016/j.trip.2023.100927

[5] Sadeghi, J., Akbari, B.: Field investigation on effects of railway track geometric parameters on rail wear, J. Zhejiang Univ. Sci. A, 7 (2006) 11, pp. 1846–1855, DOI: 10.1631/jzus.2006.A1846

[6] Daniyan, I., Mpofu, K., Nwankwo, S.: Design of a robot for inspection and diagnostic operations of rail track facilities, International Journal of Quality & Reliability Management, 40 (2023) 3, pp. 653–673, DOI: 10.1108/IJQRM-03-2020-0083

[7] Yang, Y., Wang, J., An, W., Cai, J., Bai, H.: Advances in prediction methods of wear of rails, Wear, 564–565 (2025) 205676, DOI: 10.1016/j.wear.2024.205676

[8] Auer, F.: Zur Verschleißreduktion von Gleisen, TU Graz, Wien, 2010.

[9] Analysis focused on rails, Plasser&Theurer, https://www.plassertheurer.com/en/infrastructure/modular-measuring-and-inspection-packages/modular-measuring-and-inspection-package-for-the-rail, 30.01.2026.

[10] Garreau, D., Arlot, S.: Consistent change-point detection with kernels, Electron. J. Statist, 2018., pp. 4440- 4486.

[11] Li, S., Xie, Y., Dai, H., Song, L.: M-Statistic for Kernel Change-Point Detection, Advances in Neural Information Processing Systems, 2015.

[12] Schatzl, J., Gerhold, F., Loidolt, M., Marschnig, S.: Enhancing Predictive Maintenance Through Detection of Unrecorded Track Work, Infrastructures, 2024.