Proceedings of the Seventh International Conference on
Railway Technology: Research, Development and Maintenance
Edited by: J. Pombo
Civil-Comp Conferences, Volume 15, Paper 17.2
Civil-Comp Press, Edinburgh, United Kingdom, 2026
ISSN: 2753-3239, doi: 10.4203/ccc.15.17.2

Onboard Monitoring of Vehicle and Track Condition Using Sparse Vibration Fingerprints

D. Prifer, A. Qazizadeh, S. Stichel and J. Leung

Rail Vehicles, Department of Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, Sweden

Abstract

Sparse vibration fingerprints, constructed from physics-informed spectral peaks of short-time Fourier transform spectrograms, provide an interpretable representation for railway vehicle condition monitoring. The method eliminates amplitude dependence and emphasise vehicle eigenmodes through frequency-band weighting. Using multibody simulations with suspension stiffness and track quality variations, vibration fingerprints from carbody, bogie and axlebox vertical accelerations are evaluated across three experiments. A simple matching score achieves clear separation between baseline, vehicle-related and track-related conditions, with axlebox fingerprints showing the strongest discrimination capabilities and minimal cross-confusion between vehicle and track changes. Gradual stiffness variations are detected with sensitivity to the direction of change, particularly at wheelset level where approximate severity levels are also distinguished. Noise robustness analysis reveals sensor-dependent degradation: axlebox fingerprints develop track-change bias, bogie fingerprints remain unchanged, and carbody fingerprints gradually lose their ability to detect track changes and develop a bias towards vehicle change. The proposed fingerprinting algorithm thus proves to be a promising method for physics-informed condition monitoring that balances data sparsity, interpretability and accuracy.

Keywords: vibration fingerprinting, condition monitoring, railway vehicle, multibody simulation, physics-informed, suspension stiffness, track irregularities.

1 Introduction

Railway operators increasingly rely on condition monitoring to maintain safety, availability, and reliability of vehicles and infrastructure. Condition monitoring (CM) can be done with on-board or wayside equipment. Traditionally, on-board monitoring uses vehicle mounted sensors to detect changes in the state of infrastructure, while wayside equipment uses track side sensors to monitor the passing vehicle condition. However, current on-board monitoring is also used for detecting changes in the state of suspension [1,2,3,4,5], bearing [6], wheelsets [7], and other major vehicle components. The goal of CM is to detect and monitor faults prior to unexpected service disruptions, thereby to improve the reliability of railway transportation. Most on-board CM approaches collect acceleration measurements from sensors located on the carbody, bogie frame and axleboxes. These collected acceleration signals are subsequently processed using time-domain statistics, frequency spectra, and time-frequency representation of the signal to obtain features that indicate the possible condition of the rail vehicle or component. These outputs from the sensor signals are used as input features for supervised learning algorithms [3,8,9,10]. While such methods are popular, they face challenges in interpretability of the results, sensitivity to operating conditions, and data sparsity issues.

The first issue with supervised learning algorithms, including machine-learning-based methods, e.g. neural nets or ensemble models, is their limited physical insight to how the classification was made. Therefore, they are also referred to as "black box" methods where there is no traceable physics-based connection between the input conditions and predicted output results. Thus, the relationship between fault detection and fault remediation is unclear, limiting the application of these methods to primarily detection roles. Second, the collected vibration features are strongly affected by operating conditions such as speed, load, track quality, and seasonal changes in e.g. track stiffness, which can imitate damage. This is commonly addressed by signal normalisation to make features comparable across runs, e.g. amplitude scaling, and designing problem-specific indicators [11, 12]. A third issue is that most existing approaches are tuned to specific fault cases, such as wheel flats or rail corrugation, and require re-training or changing the whole model architecture to accommodate new or concurrent degradation patterns. These limitations motivate the ongoing search for better alternatives of signal representations that are both sparse and interpretable, yet robust to noise and operational variability.

In a different domain, commercial music identification systems have demonstrated that such sparse, robust representations are possible in practice. The well-known Shazam algorithm [13] converts short audio snippets into unique identifiers that can be matched to a database of music, thus allowing for the rapid identification of an audio signal or songs. This is done by constructing an audio fingerprint by converting a 3 to 6 second recording into a spectrogram, detecting local spectral peaks, and forming combinatorial peak pairs within a pre-defined time-frequency window. A song is identified from this short snippet by computing its fingerprints and searching for matches in a large database – true matches produce a dense cluster of consistent time-offsets, while false matches remain scattered. This approach is remarkably robust to varying audio volume and background noise and can identify music from just a few seconds of corrupted audio. Although Shazam was developed for static audio signals, its core ideas suggest a potential for vibration-based condition monitoring.

Directly transferring the application field of this algorithm from static audio to dynamic sensor measurement of railway vibrations is, however, non-trivial. Audio fingerprinting assumes that the underlying content is essentially static during the query, whereas vehicle dynamics is non-stationary and depend on the vehicle-track interaction and operating conditions. The relevant frequency range is lower than that for audio signals and linked to the eigenmodes of the vehicle system, which is excited by stochastic track irregularities. Thus, by simply implementing the audio-oriented parameters, such as frequency bands and time windows tuned to musical content would therefore ignore the important aspects of the railway dynamics domain. Moreover, the goal shifts from identifying a unique event to differentiating between multiple condition states: healthy vehicle, degraded suspension components, degradation levels of track quality, and so on.

The challenges with the implementation of an audio-based signal fingerprinting methods into adjacent fields motivates a physics-informed adaptation of the fingerprinting concept to the vehicle dynamics vibration domain. In this work we explore whether sparse vibration fingerprints, constructed from short-time Fourier transform (STFT) spectrograms using peak detection and peak pairing, can capture condition dependent changes in a simulated environment. Instead of treating all frequencies and times equally, we bias the fingerprint formation toward frequency bands known to contain carbody and bogie modes and choose time and frequency quantisation scales compatible with the characteristic frequencies of railway vehicles. Using this railway focused approach, we aim to understand (1) the extent to that fingerprints can distinguish between different vehicle and track conditions, (2) the ability for fingerprints to be sensitive to gradual changes in the vehicle conditions, and (3) the extent to that fingerprints can detect defect signals from sensor noise and environmental disturbances.

For this conceptual study, results from multibody system (MBS) simulations of a passenger vehicle running at constant speed over measured track irregularities with varying operational scenarios are considered. These scenarios involve singular, discrete changes in vehicle suspension components, such as the reduction or increase in primary suspension spring stiffness, and variations in the track quality. Operational speed variation and more complex fault combinations are not considered at this stage. Within these limited operational scenarios, the aim of the paper is to assess the feasibility of physics-informed sparse vibration fingerprints as a viable and interpretable representation for condition differentiation in rail vehicles.

2 Methods

The methodology consists of eight main steps as shown in Figure 1: MBS simulation of vehicle dynamics, acceleration signal generation, segmentation of acceleration time histories, time-frequency representation of signals, sparse peak extraction, fingerprint generation, and finally database construction for subsequent matching experiments.

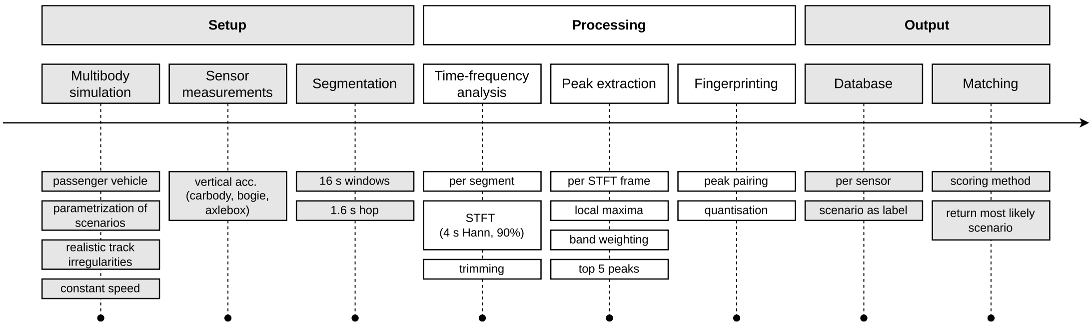

Figure 1: Overview of the proposed methodology

2.1 MBS modelling and simulation scenarios

Synthetic vehicle responses are generated using the Manchester benchmark [14] passenger vehicle model running on ballasted track, developed in SIMULIA Simpack. The wheel-rail contact was modelled using Simpack elements with non-linear wheel & rail profiles, normal forces evaluated according to the Hertzian model and tangential forces evaluated according to Kalker's Fastsim model [15]. Two types of measured track irregularities, shown in Figure 2, were applied as track-related excitations to excite the vehicle in a realistic manner, while the vehicle speed is held approximately constant with a PI controller that applies an appropriate magnitude of driving torque on all wheelsets.

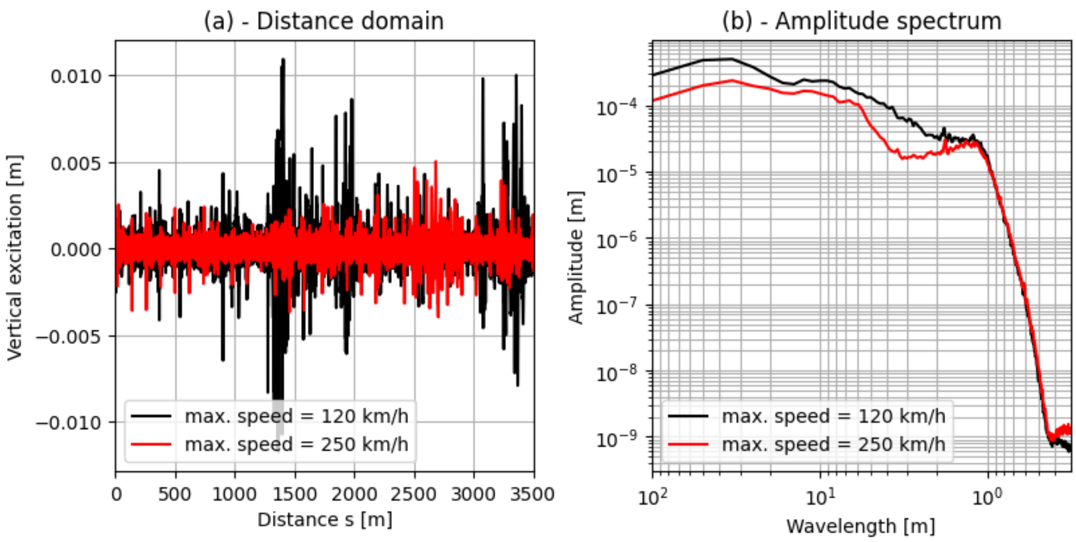

Figure 2: Comparison of applied track irregularities

Five scenarios are defined (see Table 1): one baseline on a good-quality track allowing 250 km/h, three with varying stiffness of the right-hand primary suspension spring at the leading wheelset, and one representing a low-quality track section allowing train operation only up to 120 km/h. Each scenario is simulated for 120 seconds at 50 Hz, and accelerations were measured at four locations: carbody floor above the leading bogie, cab, right side of the leading bogie frame, and axlebox at the leading right wheelset (Figure 3).

Scenario name	Speed [km/h]	Primary spring stiffness [N/mm]	Track quality
Baseline	100	1220	High
Ageing rubber spring	100	854	High
Degraded rubber spring	100	610	High
Broken coil spring	100	1647	High
Worse track quality	100	1220	Low

Table 1: Simulated scenarios.

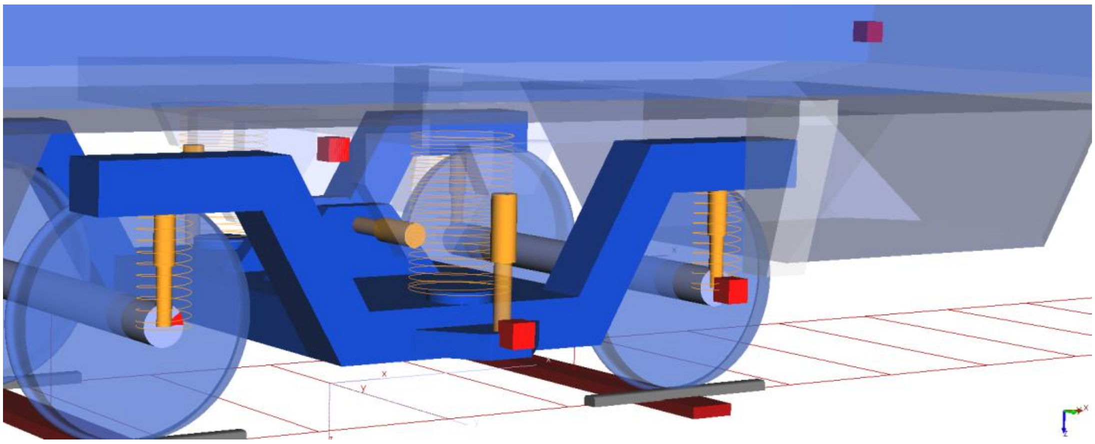

Figure 3: Sensor positions (red cubes) used for vibration signal generation

2.2 Signal segmentation

For each scenario and sensor, the acceleration time history is segmented into overlapping windows of 16 s. Segments are advanced with a hop of 1.6 s, corresponding to 90 % overlap.

2.3 Time-frequency analysis

Each 16 s segment was processed with a 4 s Hann window STFT (90% overlap) using an FFT length of 256 points at 50 Hz, yielding 0.195 Hz frequency bins and 400 ms temporal resolution.

The magnitude spectrogram is calculated for each segment and 6 STFT frames are discarded at each edge before peak detection to reduce spectral leakage and edge effects introduced by zero-padding at the beginning and end of each segment. This leaves 29 frames per segment for subsequent physics-informed weighting and peak extraction.

2.4 Physics-informed spectral peak extraction

After each segment (Figure 4a) is converted into a spectrogram (Figure 4b), local maxima are identified within each 4 s STFT frame. To emphasise frequency bands sensitive to vehicle and track conditions, the local peak amplitudes are weighted (Figure 4b, red rectangle). The typical carbody bounce mode frequency range is multiplied by 1.5, bogie modes by 1.2, and peaks above 12 Hz by 0.5. The 5 most prominent peaks, each characterised by its frequency and STFT frame centre time, are stored as a sparse set of coordinates (Figure 4d), eliminating amplitude-dependency, and emphasising modes of vehicle-track dynamics over high-frequency noise.

The resulting constellation map is a sparse representation of the segment spectrogram and is a unique representation of the given vehicle and track condition. The pattern of the peak constellation should be consistent for similar vehicle and track conditions, i.e. changes in the patterns should indicate a certain type of change in system dynamics reflecting the probability of a certain vehicle or track condition.

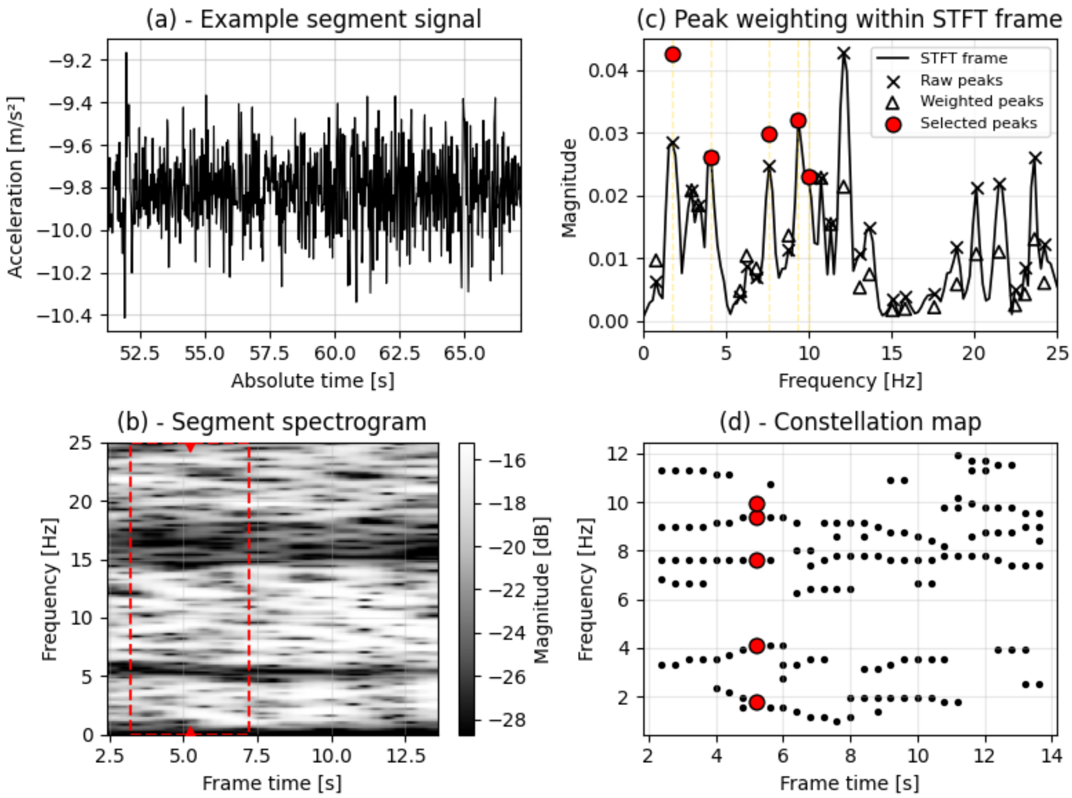

Figure 4: Signal processing and peak extraction steps for an example segment signal (a), example STFT frame (c) within the segment, shown with red rectangle in (b). The peaks extracted from the example frame are also highlighted with red in (d)

2.5 Vibration fingerprint construction

The characteristic sequence of spectral peaks, called vibration fingerprints of a segment, are constructed by forming combinatorial pairs within a time target zone as shown in Figure 5a. For each peak taken as an anchor, peaks occurring within a forward time window of 2 seconds are considered as target peaks. Each anchor - target peak pair defines a descriptor triplet (f_1, f_2, Δ_t), called a combinatorial hash, where f_1 and f_2 are the anchor and target frequencies and Δt is their time difference.

To improve robustness, the frequency and time values are quantised into discrete bins of 1 Hz, which is significantly larger than the frequency bin spacing, and 500 ms, which is slightly coarser than the temporal sampling interval of the spectrogram. The resulting integer triplet after quantization (f_1^*, f_2^*, Δt^*) serves as hash key for the fingerprint. For each unique hash key, the distinct anchor times t_1 at which this pattern occurs in the 16 s segment are stored, yielding a sparse dictionary. The quantisation and fingerprinting method are shown in Figure 5b.

The same hash key may appear in multiple scenarios. The discriminative power arises from the constellation of the hashes and their relative frequencies across scenarios, rather than from any single hash isolation.

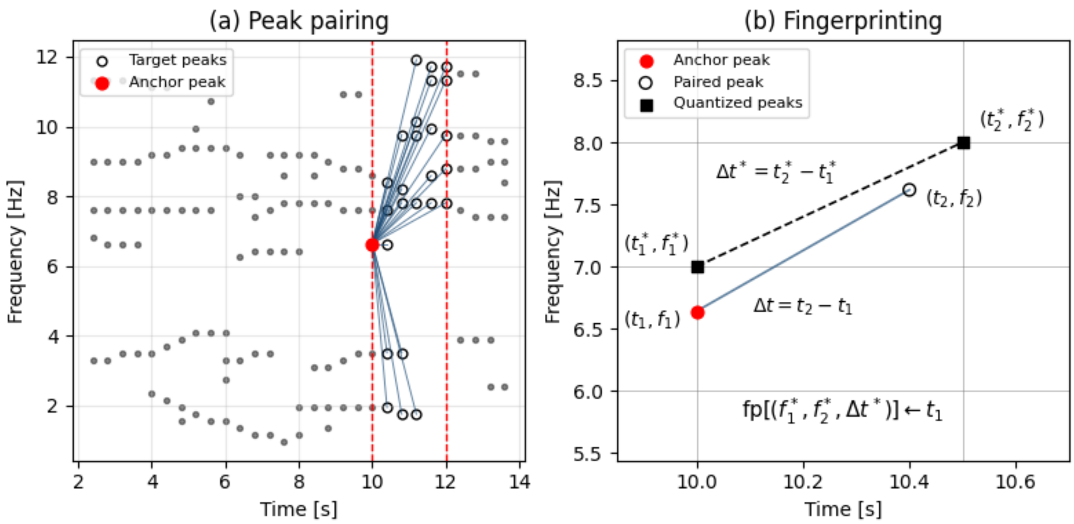

Figure 5: An example peak is acting as an anchor is shown in (a), together with all its combinatorial pairs within a 2 s forward window. Each pairwise connection is subsequently quantised and transformed into a fingerprint as shown in (b)

2.6 Fingerprint database

Simulation results from Section 2.1 are segmented as described in Section 2.2. To evaluate matching performance, segments are split into database and query sets by segment index. Even-indexed segments are processed according to Sections 2.3-2.5 to generate 2 128 500 fingerprints for the database. The database is organised per sensor location, and each hash key is associated with a list of occurrences across scenarios and segments for each sensor. For sensor s, the database dictionary is defined according to Equation (1), where $scenario$ is a condition label from Table 1, and t_{seg} is the segment start time.

$$DB_s[(f_1^*, f_2^*, \Delta t^*)] = \{(scenario, segment\ idx, t_1, t_{seg}), \ldots\} \tag{1}$$

This structure enables counting how often a given hash appears in each scenario and to analyse the alignment of anchor times between a query segment and database segments. Odd-indexed segments are reserved for experiments designed to address the research goals introduced in Section 1. The subsequent analysis phase defines a matching score between query fingerprints and database and uses this score to perform three experiments introduced in the Results section.

3 Results

This section summarises the performance of the sparse vibration fingerprints across three cases. In Section 3.1, general condition separability is assessed, quantifying whether fingerprints can be used to distinguish between changes in the vehicle condition, and changes in the track quality compared to the baseline conditions. In Section 3.2, the detailed condition separability is assessed, quantifying whether the severity and direction of suspension stiffness change can be identified from the fingerprints. In Section 3.3, the robustness of the fingerprint to noise is tested through progressively degrading the acceleration signals with increasing levels of white noise and observing how classification performance deteriorates.

Each query segment (see Section 2.6) is treated as an independent query. The best candidate scenario is selected based on the diversity of matching patterns, i.e. the number of unique hash keys shared between the query and each scenario. The scenario with the highest score is chosen as prediction.

3.1 Experiment 1: general condition separability

Three condition classes are defined: baseline, vehicle change and track change. The query segments are classified into one of the three labels using fingerprints from the different sensors. Figure 6 shows the confusion matrices, one per sensor.

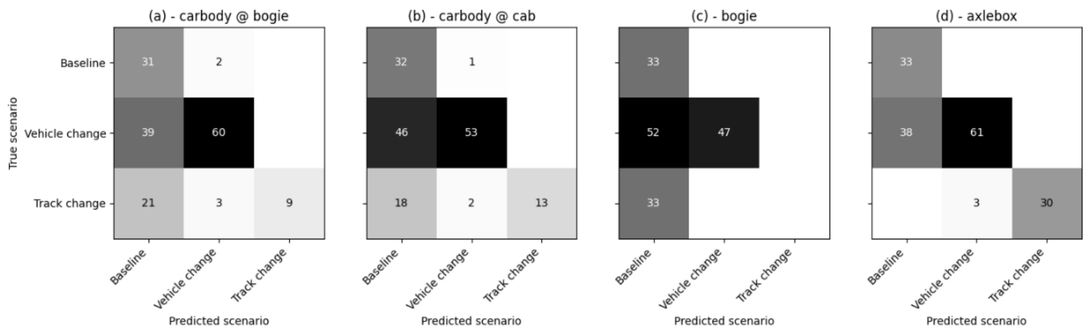

Figure 6: Confusion matrices for condition separability test for the four sensors

The carbody fingerprints, shown in Figure 6a and Figure 6b yielded a pronounced diagonal, but with a noticeable bias towards the baseline class. A substantial fraction, 55% and 58% of the segments, for the above the bogie and cab, respectively, across all conditions for each case were classified as baseline, which is consistent with the filtering effects of primary and secondary suspensions. Despite this tendency, over half of the vehicle-change segments are correctly identified, with none confused with track changes. For the 33 track-change segments, most non-baseline predictions are correct, with only a few vehicle-change confusions. This pattern indicates that carbody fingerprints can reliably distinguish vehicle issues from track issues, even if they tend to over-report the baseline class.

The confusion matrix of the bogie sensor, shown in Figure 6c, showed a stronger bias towards the baseline class compared to the carbody mounted sensor. This sensor identified correctly the suspension changes, but with worse recall than the carbody-mounted sensors. However, an important observation is that all the predictions made other than baseline, are correct predictions, i.e. the bogie-mounted sensor is very good at picking up vehicle suspension changes, but the algorithm can not distinguish between baseline and track-change fingerprints.

The axlebox sensor results, shows in Figure 6d, produced the cleanest diagonal without a systematic confusion between baseline and track change conditions. This indicates that the wheelset level fingerprints capture the track excitation sufficiently well to separate track-related changes from suspension-related changes and baseline conditions. Thus, this sensor is the most informative for the simulated sensors for condition separability in the present framework.

3.2 Experiment 2: sensitivity to gradual stiffness change

The second experiment tests whether the vibration fingerprints can represent gradual changes in primary suspension stiffness. As presented in Table 1, four cases are distinguished: baseline, ageing rubber spring (30% stiffness reduction), degraded rubber spring (50% stiffness reduction) and broken coil spring (35% stiffness increase).

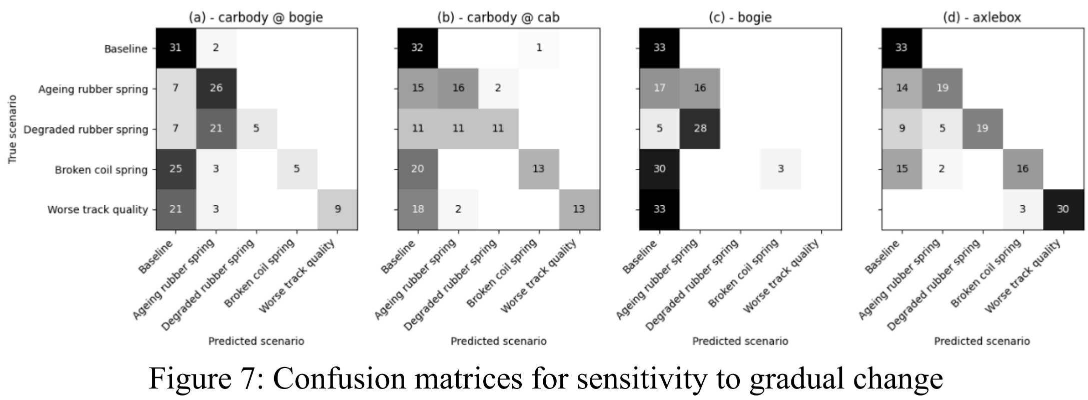

Figure 7: Confusion matrices for sensitivity to gradual change

Experiment 1 showed the systematic bias towards baseline condition for all sensors. Experiment 2 examines a more detailed level of condition separability to understand the correctness of true vehicle-related changes. The sensors mounted on the carbody exhibit a clear diagonal, shown by Figure 7a and Figure 7b. However, there are differences in how well they distinguish between stiffness conditions. The sensor above the bogie reliably detects "something is softer than baseline" and correctly assigns the direction of change towards reduced stiffness but tends to underestimate severity by misclassifying the 50% stiffness reduction as 30% stiffness reduction. In contrast, broken springs are mostly classified as baseline indicating difficulties in the carbody sensor above the bogie to detect increases in suspension stiffness. The carbody sensor located in the cab, which is longitudinally offset from the bogie centre and thus more sensitive to carbody pitch, shows a different pattern. Here, broken coil springs are detected more clearly, and the two reduced-stiffness cases are also better separated compared to the above bogie carbody sensor, although some confusion remains between 50% and 30% reductions.

The tendency for the carbody sensors to more accurately identify reduction in primary suspension spring stiffness compared to increases suggests that the direction of stiffness change, is also a significant factor in the effectiveness of the algorithm. This behaviour suggest that the direction of stiffness change (increase or decrease in stiffness) shifts the dominant frequencies in a way that the changes in the peak constellation map are more prominent for accelerations including carbody pitch content.

The bogie-mounted sensor systematically detects "stiffness has reduced" and does not confuse stiffness changes with track changes for these simulated cases, but it is essentially blind to the magnitude of that change, as shown in Figure 7c. This suggests that bogie-level fingerprints respond to the deviation from nominal stiffness, but the resulting constellation is not discriminative enough to separate different types or levels of degradation.

The axlebox sensor continues to provide the clearest separation, with condition discrimination across different stiffness changes. Its confusion matrix retains a strong diagonal, as shown in Figure 7d, and the misclassifications are modest and predictable; when the classification is wrong, it tends to favour the ageing rubber spring case. Overall, these results indicate that for the tested experimental conditions, wheelset-level fingerprints compared to the other carbody and bogie fingerprints, are the most successful to separate vehicle-related and track-related changes and may also carry useful information about the direction and approximate magnitude of suspension stiffness change.

3.3 Experiment 3: robustness to added noise

This experiment investigates how white noise injection impacts the ability of the fingerprints to distinguish between baseline, vehicle-related and track-related changes. Rather than focusing on overall accuracy, which is dominated by baseline bias discussed earlier, we consider the recall of each class – defined as the proportion of true scenarios correctly predicted – as a function of noise level and sensor location.

To control the noise level, white noise was injected into each acceleration signal at different signal-to-noise ratios (SNR) according to Equation (2),

$$\sigma_n^2 = \frac{\sigma_s^2}{10^{SNR_{db}/10}}, \tag{2}$$

where σ_s^2 is the power of the mean-centred signal and σ_n^2 is the noise power.

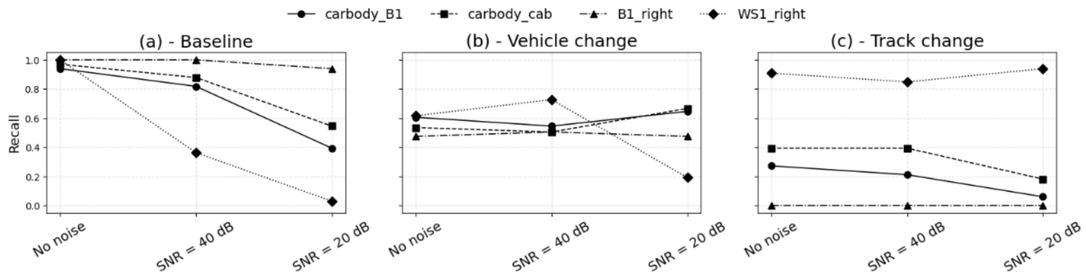

Figure 8: Recall of each condition class per sensor as a function of SNR

For the baseline condition, shown in Figure 8a all sensors start with recall close to one in the noise-free case, which is consistent with the strong baseline bias observed in Experiment 1. However, all sensors drop in recall with increasing noise content, except for the bogie-mounted sensor, which only changes marginally. Since the main purpose is to show how well the fingerprints retain their capabilities of distinguishing between vehicle-related and track-related changes in conditions, the focus is on the other two subfigures, Figure 8b and Figure 8c.

The axlebox sensor, indicated by markers WS1_right, shows a very pronounced shift to a track-related change bias. A steep drop in both the baseline and vehicle change classes, and virtually no change in track change indicates that most segments are confused with track change under noisy conditions.

The bogie-mounted sensor behaves differently. As discussed in Section 3.1, in the noise-free case it is strongly biased towards baseline, many times detecting vehicle changes and never confusing vehicle and track changes. With increasing noise levels, only a few baseline segments are misclassified into vehicle change, but for track changes we can observe a flat response with respect to noise, as shown in Figure 8c, indicating that the property of not confusing vehicle changes with track change remains intact under noisy conditions.

For the carbody sensors, recall curves evolve more smoothly with increasing noise levels. Baseline and track change recall degrades gradually as SNR is reduced, while the recall for vehicle change slightly improves, as shown in figure 8b, which suggests a slight increase of bias towards vehicle change. Overall, it can be surmised that noise modifies the peak constellation in a way that emphasizes vehicle changes but makes track changes even harder to detect.

The observations above indicate that the robustness of the sparse fingerprints to white noise is strongly sensor dependent. The axlebox retains good track change recall at the expense of vehicle change recall, effectively converging to a “track change or not” detector. Conversely, the bogie sensor remains dominated by the baseline condition regardless of noise, while the carbody sensors exhibit a more nuanced shift in bias with noise.

4 Conclusions

This paper explores whether sparse vibration fingerprints, constructed from physics-informed spectral peaks, can serve as a viable and interpretable representation for condition differentiation in rail vehicles and track, using numerical simulations.

Wheelset-level fingerprints delivered the clearest separation between baseline, vehicle changes and track changes, while bogie-level fingerprints only indicated “something changed in the suspension”. Carbody-level fingerprints tended to over-predict the baseline configuration, but confusion between vehicle suspension-related and track-related changes was rare. The placement of the sensor within the carbody also influenced the detection capability where a longitudinally offset sensor in the cab, rather than directly above the bogie centre, improved the ability to detect suspension changes.

For gradual suspension stiffness changes, the fingerprints were sensitive to the direction of stiffness variation, but the magnitude of change was only captured reliably at wheelset-level. Under additive white noise, robustness proved strongly sensor-dependent: wheelset fingerprints increasingly mislabelled vehicle changes as track changes as SNR decreased, bogie-level fingerprints remained essentially unchanged, and carbody-level fingerprints became slightly more sensitive to vehicle change while losing some track change recall.

These conclusions are drawn from an exploratory study based on multibody simulations of a single vehicle at one speed, with isolated parameter changes and idealised sensors. Real-world data will introduce additional variability from speed, loading, sensor placement and combined faults. In addition, the parameters of time-frequency analysis and fingerprinting pipeline (segmentation length, sampling frequency, physics-informed frequency weighting, quantisation and the number of peaks retained per STFT frame) were chosen heuristically in this work and require systematic optimisation.

The agreement between the fingerprints and the experimental scenarios indicate that sparse fingerprints are a good candidate for onboard condition detection of vehicle and track scenarios. The results also suggest that simply counting the number of unique matching hashes per condition is sufficient to obtain interpretable separation between baseline, vehicle-related and track-related changes. Nevertheless, identifying alternative scoring methods that reduce baseline bias and improve quantitative severity estimation remains an important direction for future research.

References

[1] F. Liu et al., “Correlation Signal Subset-Based Stochastic Subspace Identification for an Online Identification of Railway Vehicle Suspension Systems,” Vehicle System Dynamics, vol. 58, no. 4, pp. 569–589, 2019. doi:10.1080/00423114.2019.1589534

[2] M. Dumitriu, “Fault detection of damper in railway vehicle suspension based on the cross-correlation analysis of bogie accelerations,” Mechanics and Industry, vol. 20, no. 1, 2019. doi:10.1051/meca/2018051

[3] Y. Zhang et al., “Fault Diagnosis of High-speed Train Bogie Based on Deep Neural Network,” IFAC-PapersOnLine, vol. 52, no. 24, pp. 135–139, 2019. doi:10.1016/j.ifacol.2019.12.395

[4] Q. Rao et al., “Condition recognition of high-speed train bogie based on multi-view kernel FCM,” Big Data Mining and Analytics, vol. 2, no. 1, pp. 1–11, 2019. doi:10.26599/BDMA.2018.9020027

[5] D. Lebel et al., “High-speed train suspension health monitoring using computational dynamics and acceleration measurements,” Vehicle System Dynamics, vol. 58, no. 6, pp. 911–932, 2019. doi:10.1080/00423114.2019.1601744

[6] Z. Zheng et al., “A Fault Diagnosis Method of Bogie Axle Box Bearing Based on Spectrum Whitening Demodulation,” Sensors, vol. 20, 2020. doi:10.3390/s20247155

[7] M. Hassan et al., “Experimental and numerical investigation of the possibilities for the structural health monitoring of railway axles based on acceleration measurements,” Structural Health Monitoring, vol. 18, no. 3, pp. 902–919, 2019. doi:10.1177/1475921718786427

[8] J. Li et al., “Rail Corrugation Detection of High-Speed Railway Using Wheel Dynamic Responses,” Shock and Vibration, vol. 2019, 2019. doi:10.1155/2019/2695647

[9] J. Liu et al., “Detecting anomalies in longitudinal elevation of track geometry using train dynamic responses via a variational autoencoder,” Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems, vol. 10970, 2019. doi:10.1117/12.2513711

[10] M. Sysyn et al., “Common crossing condition monitoring with on-board inertial measurements,” Acta Polytechnica, vol. 59, no. 4, pp. 423–434, 2019. doi:10.14311/AP.2019.59.0423

[11] R. Sánchez et al., “Evaluation of Time and Frequency Condition Indicators from Vibration Signals for Crack Detection in Railway Axles,” Applied Sciences, vol. 10, no. 12, 2020. doi:10.3390/app10124367

[12] X. Wei et al., “Squats and corrugation detection of railway track based on time-frequency analysis by using bogie acceleration measurements,” Vehicle System Dynamics, vol. 58, no. 8, pp. 1167–1188, 2019. doi:10.1080/00423114.2019.1610181

[13] A. L.-C. Wang, “An Industrial-Strength Audio Search Algorithm,” in Proceedings of the 4th International Conference on Music Information Retrieval (ISMIR 2003), Baltimore, Maryland, USA, October 27–30, 2003, pp. 7–13.

[14] S. Iwnicki, “Manchester Benchmarks for Rail Vehicle Simulation,” Vehicle System Dynamics, vol. 30, no. 3–4, pp. 295–313, 1998. doi:10.1080/00423119808969454
[15] J. J. Kalker, “A fast algorithm for the simplified theory of rolling contact,” Vehicle System Dynamics, vol. 11, no. 1, pp. 1–13, 1982. doi:10.1080/00423118208968684