Proceedings of the Seventh International Conference on
Railway Technology: Research, Development and Maintenance
Edited by: J. Pombo
Civil-Comp Conferences, Volume 15, Paper 7.5
Civil-Comp Press, Edinburgh, United Kingdom, 2026
ISSN: 2753-3239, doi: 10.4203/ccc.15.7.5

A CNN-Based Rail Classification Platform for Explainable Fault Detection in Weigh-In-Motion and Rolling Stock Monitoring Systems

N. S. Vyas

Department of Mechanical Engineering, Indian Institute of Technology Kanpur, India

Abstract

Wheel Impact Load Detection (WILD) and Weigh-In-Motion (WIM) systems have been deployed at trackside locations on the Indian Railways network for over two decades, providing reliable detection of severe wheel defects through threshold-based Impact Load Factor (ILF) analysis. While this represents a significant operational capability, the diagnostic reach of conventional WILD systems is fundamentally constrained: they are effectively blind to the broader class of rolling stock health conditions — suspension degradation, axle box misalignment, uneven loading, tread polygonisation — that precede and predict catastrophic failure. This paper presents a CNN-based rail classification platform, designated WILDPulse, deployed as a software augmentation layer above existing WILD hardware, extending the diagnostic pipeline through three interconnected capabilities: (i) physics-informed multi-domain signal analysis in time, frequency, and time-frequency domains; (ii) a ResNet-18-based CNN classification layer mapping STFT spectrogram images of axle events to rolling stock health states; and (iii) a Train Health Index (THI) providing longitudinal, rake-level health records for condition-based maintenance. Results are presented for a real field WILD dataset with fault-modified signals synthesised via physics-informed simulation for five fault types and their combination. Explainability is provided through Grad-CAM attribution maps.

Keywords: wheel impact load detector, wheel-rail contact forces, digital signal processing, deep learning, fault diagnosis, rolling stock maintenance.

1 Introduction

Railway safety and operational reliability are critically dependent on the continuous health of two interacting systems: the rolling stock fleet and the track infrastructure over which it operates. Wheel–rail contact forces govern derailment risk, track degradation rates, and passenger ride quality, and their dynamic characteristics at the contact point are the primary source of information about the condition of both systems [1]. The detection and characterisation of abnormal dynamic loads — arising from wheel defects, suspension degradation, or track irregularities — is therefore fundamental to condition-based maintenance (CBM) in railway operations. Weigh-In-Motion (WIM) and Wheel Impact Load Detection (WILD) systems address this need from the wayside, measuring the dynamic load history of every axle through a fixed strain-gauged track section, providing a non-intrusive, continuous measurement capability without modification to the rolling stock fleet [2,3]. The diagnostic metric conventionally derived from WILD data is the Impact Load Factor (ILF) — the ratio of the dynamic peak load to a reference quasi-static level — against which threshold-based alarms are set to detect wheel flats and severe tread defects.

WILD systems have been deployed on the Indian Railways network since approximately 2005, with the IIT Kanpur system operational at approximately 25 locations. However, the ILF threshold represents only the most visible fraction of the diagnostic information encoded in WILD signals [4]. The strain-gauge data at 10,000 samples per second per channel contains rich spectral information sensitive to a wide range of rolling stock conditions, most invisible to ILF-threshold logic. This paper presents WILDPulse, a CNN-based rail classification platform that installs as a software intelligence layer above existing WILD hardware — requiring no new sensors, no site disruption, and no modification to existing acquisition systems — transforming the WILD platform from a threshold-based defect detector into a comprehensive rolling stock condition assessment system through three technical contributions: (i) physics-informed characterisation of how fault conditions encode their signatures in time-domain, FFT, and STFT representations of WILD signals; (ii) a CNN-based classification layer using a pretrained ResNet-18 architecture mapping STFT spectrogram images to rolling stock health-state categories; and (iii) a Train Health Index (THI) — a composite, rake-level scalar updated on every train passage.

The paper is presented as work in progress. The WILD platform and ILF analytics are fully field-validated and operational; multi-domain signal analysis is demonstrated on real field data (LHB coach passage, 10 kHz). The CNN has been validated for train-configuration classification on real data. Individual fault signature characterisations in Sections 4 and 5 are based on physics-informed numerical simulation, clearly labelled throughout, and require validation on labelled field fault datasets. The paper is organised as follows: Section 2 reviews the WILD system; Section 3 describes ILF analytics; Section 4 presents fault signal signatures; Section 5 describes the WILDPulse architecture and CNN; Section 6 introduces the THI; Sections 7 and 8 discuss limitations and conclude.

2 WILD System: Instrumentation and Deployment

2.1 Instrumentation Principles

The WILD system measures rail deflections caused by wheel–rail contact forces through strain gauges bonded to the rail web, foot, and/or head. The system employs twelve measurement channels per rail (left and right), providing spatial redundancy across approximately 1.2 m of track. Signals are acquired at 10 kHz, providing a Nyquist frequency of 5,000 Hz sufficient to capture wheel-flat impact transients (energy to several kilohertz), suspension resonances (10–200 Hz), and wheel rotation harmonics (6–100 Hz). Triggering is achieved by an inductive wheel sensor upstream, initiating acquisition approximately 2 seconds before the first axle arrives. A LabVIEW-based front-end system performs signal conditioning, axle segmentation, speed estimation, ILF calculation, and alarm management in real time, with data stored in CSV/XLSX format for subsequent analysis.

2.2 System Deployment and Salient Features

The WILD system technology developed at IIT Kanpur from 2005 onwards has been deployed at approximately 25 locations across the Indian Railways network, covering both passenger and freight corridors — typically at major junctions, near maintenance depots, and on high-traffic freight corridors. Figure 4 shows the system architecture from rail-mounted strain gauges through the LabVIEW real-time platform to the WILDPulse intelligence layer and its THI output. The platform provides on a 24×7 basis: automatic train identification; dynamic load and speed measurement with bidirectional traffic support; self-calibration and self-diagnostics; automatic alarm transmission via GSM/GPRS; and -real-time operation with offline replay capability.

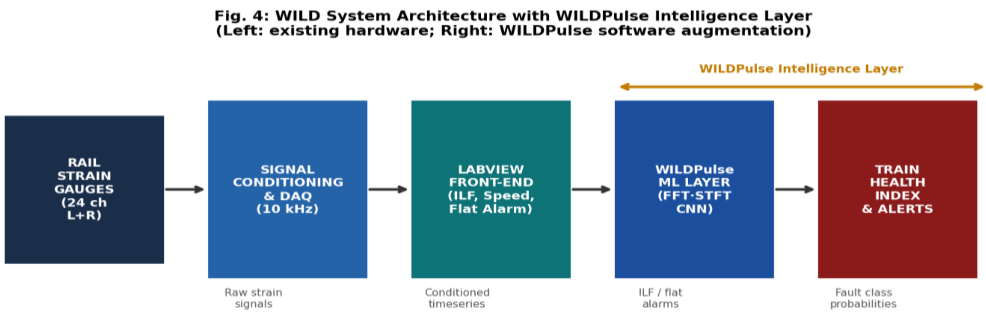

Figure 4: WILD System Architecture with WILDPulse Intelligence Layer. Left: existing field-deployed hardware (strain gauges, conditioning, LabVIEW DAQ). Right: WILDPulse software augmentation layer (signal processing, CNN, THI). No hardware modification required.

3 Impact Load Factor and Conventional Wheel Flat Detection

3.1 Impact Load Factor Definition

The Impact Load Factor (ILF) is the central diagnostic metric of conventional WILD systems, defined as the ratio of the peak dynamic strain response during an axle event to a reference quasi-static strain level:

$$ILF = \varepsilon\_peak / \varepsilon\_static \qquad (1)$$

where ε_peak is the maximum strain magnitude and ε_static is the quasi-static reference strain corresponding to the nominal static axle load. The ILF is dimensionless, enabling direct comparison across different train types, axle loads, and operating speeds. Signal-to-noise performance is enhanced by the 12-channel spatial redundancy, with ILF computed independently for each channel pair and averaged.

3.2 Threshold-Based Alarm Logic

Wheel flat detection is implemented through a multi-level threshold logic applied to the computed ILF. Three alarm levels are defined in the deployed system, as shown in Table 1.

Alarm Level	ILF Threshold	Operational Response
Advisory	1.5 – 2.0	Flag for inspection at next scheduled maintenance opportunity; monitor trend
Warning	2.0 – 3.0	Prioritise inspection within defined distance/time limit; increase monitoring frequency
Critical	> 3.0	Remove vehicle from service at next available stop; immediate inspection required

Table 1: ILF threshold levels deployed WILD platform.

Crucially, each alarm is traceable to the underlying strain signal, ensuring regulatory auditability. Figure 5 shows the baseline healthy wheel passage signal from the field dataset, along with its FFT and STFT representations, which serve as the reference for all subsequent fault-signature comparisons.

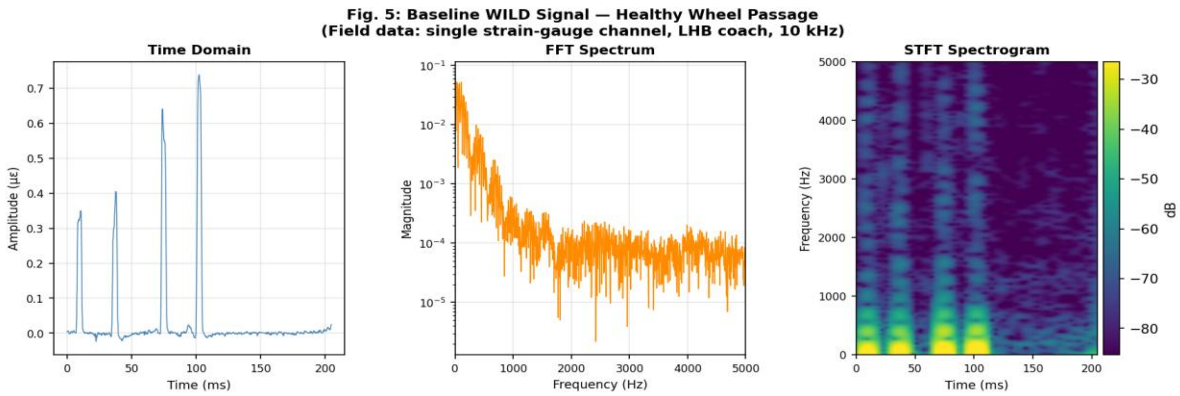

Figure 5: Baseline WILD Signal — Healthy Wheel Passage (field data). Left: time domain (10 kHz, LHB coach, 204.8 ms window). Centre: FFT (log scale). Right: STFT spectrogram (256-point window, 87.5% overlap). No impact transients; quasi-periodic loading at wheel rotation frequency.

3.3 Limitations of Conventional ILF-Based Diagnostics

By reducing the full information content of the WILD signal to a single scalar, ILF discards the vast majority of the diagnostic information encoded in the signal's temporal structure, frequency content, and time-frequency characteristics. ILF-based diagnostics are: insensitive to progressive degradation (ILF responds to impact-generating defects but poorly to gradual conditions evolving over weeks without discrete impact events); unable to discriminate fault types (an ILF spike provides no information about its origin); blind to dynamic imbalance and asymmetry (uneven loading and axle box canting are invisible to single-channel vertical ILF); and incapable of health trending. These limitations motivate the WILDPulse framework.

4 Strain-Gauge Signal Signatures: Rolling Stock Dynamics and Fault Physics

4.1 Physical Basis: How Rolling Stock Conditions Encode in Rail Strain

The strain at a fixed rail point is the direct mechanical response to wheel–rail contact forces. For a wheel rolling at velocity v, the strain time-series during a single axle passage decomposes as:

$$\varepsilon(t) = \varepsilon\_static(P_0) + \varepsilon\_dynamic(Fd(t)) + \varepsilon\_rotation(Fflat, \theta(t)) + \varepsilon\_suspension(k, c, t) + \varepsilon\_noise(t) \quad (2)$$

where P_0 is the quasi-static axle load, Fd(t) is the dynamic contact force, Fflat and θ(t) are the impulsive force and rotation angle for wheel defects, k and c are suspension stiffness and damping, and ε_noise is measurement noise. Table 2 summarises the physical decomposition and diagnostic significance of each signal component. The wheel rotation frequency $fr = v/(2\pi r)$ ($r = 0.46$ m for LHB coach) gives $fr \approx 6.9$ Hz at 72 km/h; its harmonics are among the most diagnostic features in the FFT of WILD data.

Signal Component	Frequency Range	Diagnostic Information
Quasi-static load (ε_static)	DC – 2 Hz	Static axle load; overload condition; gross left-right asymmetry
Dynamic contact force (ε_dynamic)	2 – 30 Hz	Track irregularities; wheel profile geometry; suspension dynamics; kinematic oscillations
Wheel rotation harmonics (ε_rotation)	fr × 1,2,3… Hz	Wheel roundness; polygonisation order; flat chord frequency content
Flat/defect impact (ε_flat)	10 Hz – 3 kHz	Wheel flat severity; tread shelling pit size; polygonisation amplitude
Suspension dynamics (ε_suspension)	10 – 300 Hz	Primary/secondary spring and damper condition; anti-roll bar stiffness
Axle box dynamics (ε_axlebox)	20 – 1000 Hz	Bearing race condition; axle box canting; lubrication degradation

Table 2: Physical decomposition of WILD strain signal components and their diagnostic significance.

4.2 Fault Case 1 — Wheel Flat

A wheel flat generates a periodic impulse at fr with an exponentially-decaying oscillation. In the frequency domain, energy appears at fr and all harmonics with a broadband pedestal to several kilohertz. In the STFT, the flat signature appears as vertically-oriented bright streaks at periodic time intervals of 1/fr (Figure 6a).

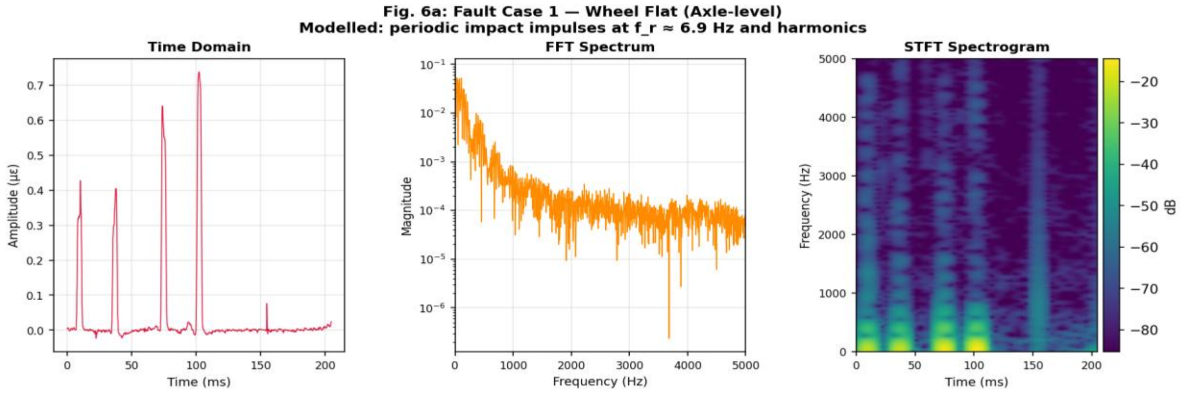

Figure 6a: Fault Case 1 — Wheel Flat. Periodic sharp impulses at fr intervals (time domain); harmonic series at fr and multiples in FFT; vertically-oriented bright streaks in STFT.

4.3 Fault Cases 2 and 3 — Suspension Deterioration

Primary suspension deterioration (spring fatigue, damper leakage) manifests as elevated energy in the 30–130 Hz band corresponding to bogie natural frequencies, appearing as persistent horizontal bands in the STFT throughout the passage event (Figure 6b). Secondary suspension degradation (air spring pressure loss, damper wear) appears primarily as elevated low-frequency content below 15 Hz, manifesting as a slow amplitude modulation at 4–12 Hz and a low-frequency STFT band spanning the full passage duration (Figure 6c). The STFT distinguishes secondary suspension degradation from wheel polygon signatures by the persistence of low-frequency energy across the entire passage.

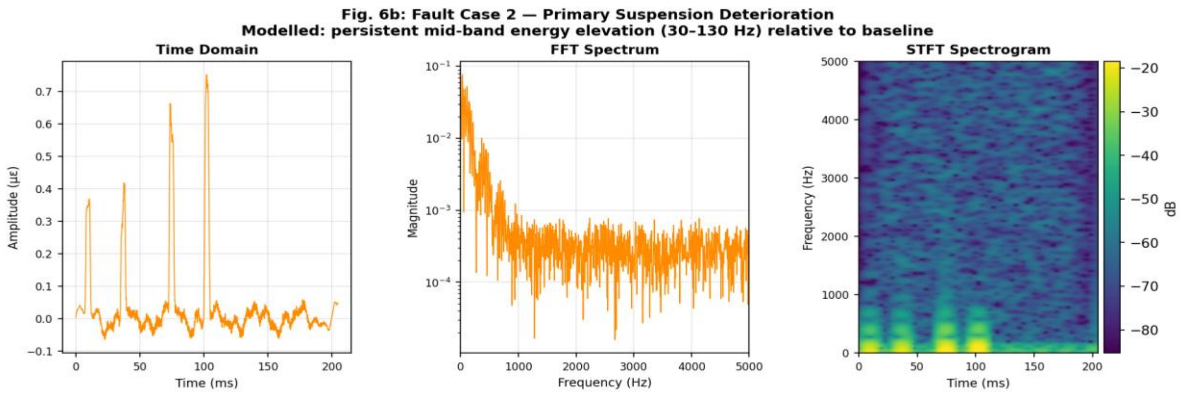

Figure 6b: Fault Case 2 — Primary Suspension Deterioration. Persistent mid-band energy at 30–130 Hz in FFT; horizontal energy bands in STFT.

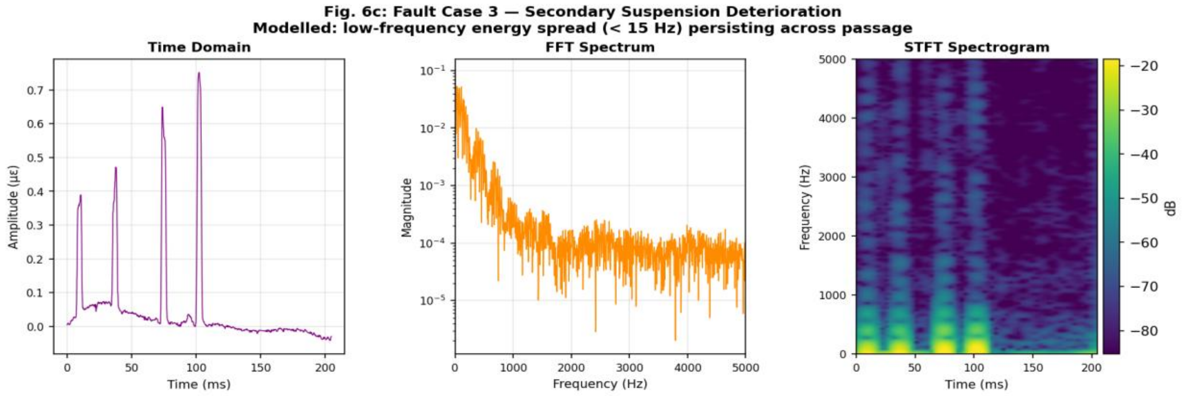

Figure 6c: Fault Case 3 — Secondary Suspension Deterioration. Dominant energy below 15 Hz in FFT; low-frequency energy band spanning full passage duration in STFT.

4.4 Fault Cases 4 and 5 — Axle Box Canting and Tread Shelling

Axle box canting produces a lateral load imbalance and excites wheelset hunting oscillations at 2–6 Hz. In the strain signal this appears as an amplitude modulation of the carrier at the hunting frequency, with sideband components flanking fr harmonics in the FFT and a slow amplitude modulation of STFT energy bands (Figure 6d). Wheel tread shelling/polygonisation produces a dense harmonic series at fr × n (n = 1, 2, 3…) extending to several hundred hertz, identifiable in the FFT by harmonics beyond the 5th plus a broadband noise floor elevation, and as multiple thin vertical streaks in the STFT (Figure 6e).

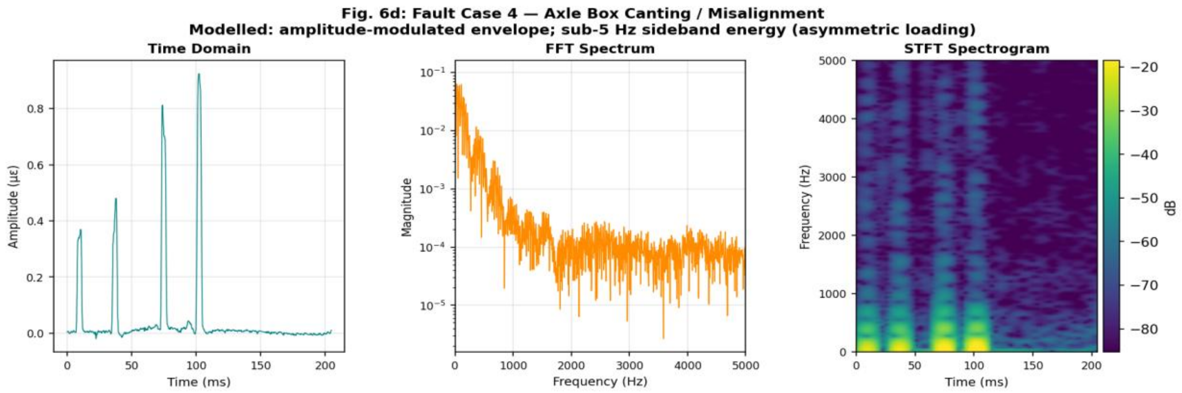

Figure 6d: Fault Case 4 — Axle Box Canting. Amplitude-modulated signal envelope at ~3 Hz; sideband components flanking fr; slow STFT energy modulation.

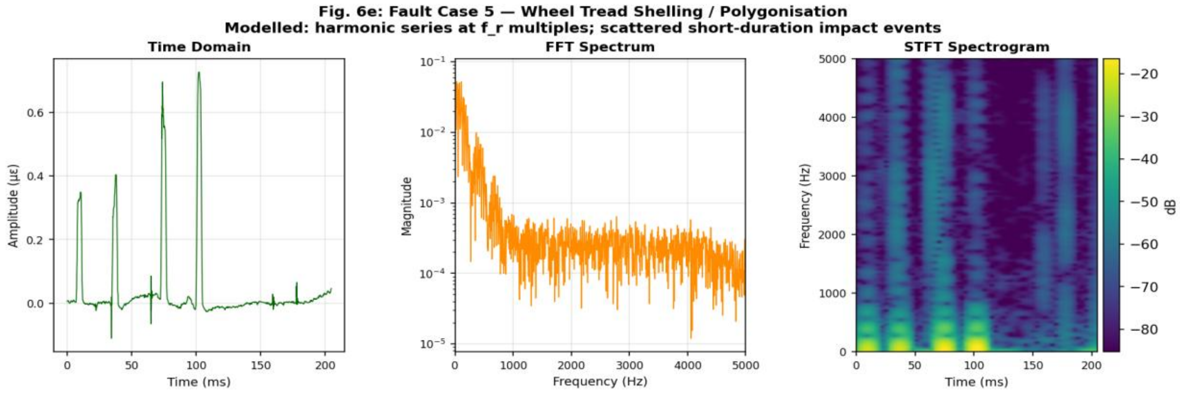

Figure 6e: Fault Case 5 — Tread Shelling / Polygonisation. Dense harmonic series at fr × 1–7+ in FFT; multiple thin vertical STFT streaks

4.5 Combined Fault Case and Spectrogram Comparison

In operational service, multiple fault conditions frequently coexist. Figure 6f demonstrates that superimposing the wheel flat, primary suspension degradation, and axle canting modifications on the baseline produces a combined signal in which all three signature types remain present and individually recognisable — a critical additivity property for multi-class CNN classification. Figures 7, 8, and 9 present the full STFT spectrogram comparison across all seven cases at identical colour scale limits, the overlaid FFT spectra with annotated diagnostic frequency bands, and the Fault-Feature Atlas mapping each fault to its time-domain, FFT, and STFT characteristics respectively. The visual separability of the cases in both the spectrogram and FFT domains directly supports the feasibility of CNN-based classification.

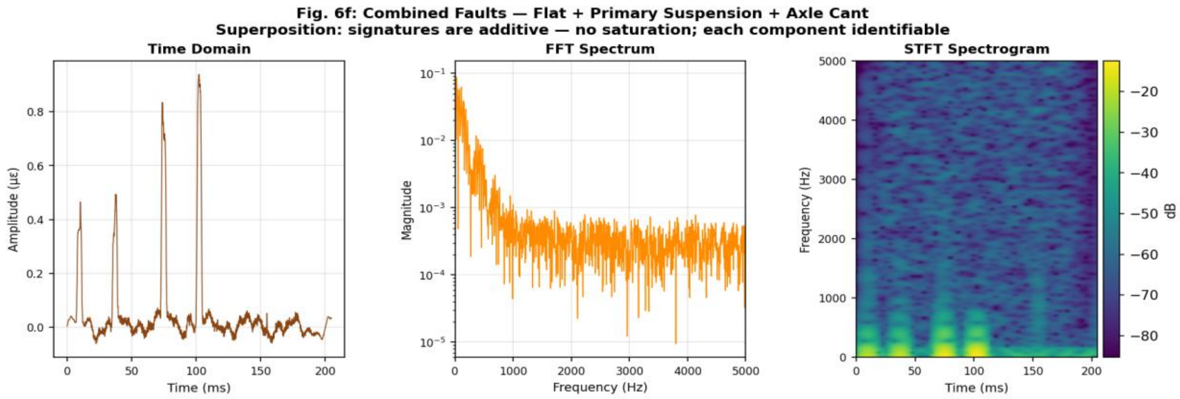

Figure 6f: Combined Faults — Wheel Flat + Primary Suspension Degradation + Axle Canting. All three fault signatures are present and identifiable; no saturation or mutual suppression.

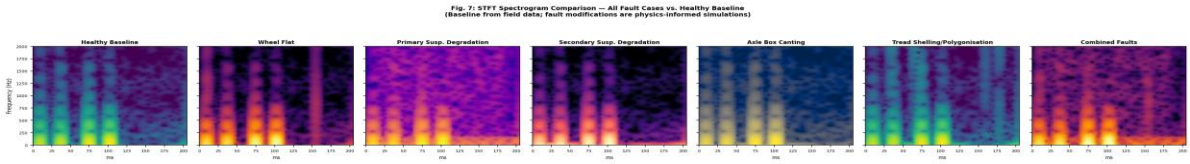

Figure 7: STFT Spectrogram Comparison — All Cases (identical scale). Left to right: Healthy; Wheel Flat; Primary Susp. Degradation; Secondary Susp. Degradation; Axle Box Canting; Tread Shelling; Combined Faults.

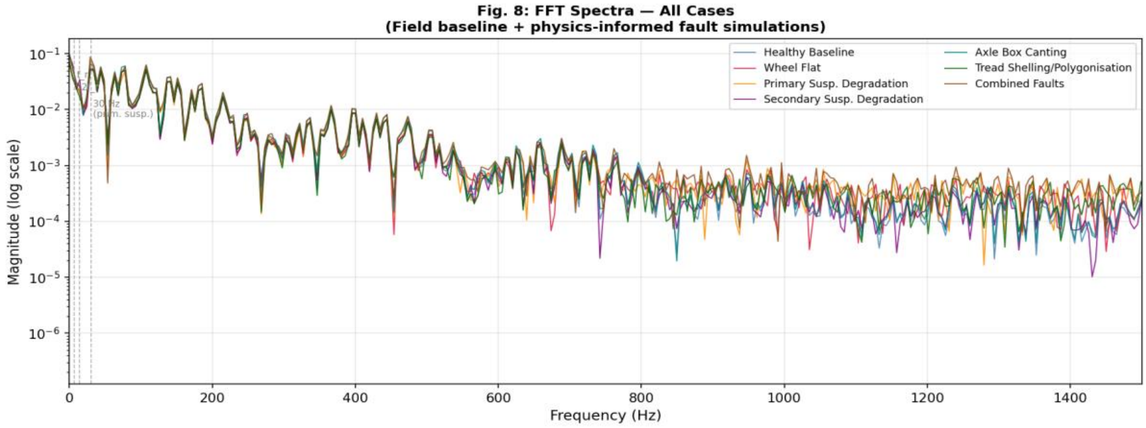

Figure 8: FFT Spectra — All Cases Overlaid. Vertical dashed lines mark fr, 2fr and 30 Hz (primary suspension natural frequency). Each fault case produces distinct spectral elevation patterns.

Fig. 9: Fault-Feature Atlas — Frequency and STFT Patch Signatures
Key diagnostic markers in FFT and STFT for each rolling stock fault condition

Fault Type	Time Domain	FFT Signature	STFT Patch to look for
Wheel Flat	Periodic sharp impulses @ T_r intervals	f_r harmonics (≈7, 14, 21...Hz) +broadband bursts	Vertical bright streaks periodically spaced in time
Primary Susp. Deg.	Elevated amplitude vibration throughout	Elevated 30–130 Hz band peak at 30 Hz nat. freq.	Persistent mid-band horizontal energy bands
Secondary Susp. Deg.	Slow low-freq oscillation modulation	Dominant 4–15 Hz components	Low-freq energy (base of spectrogram)
Axle Box Canting	AM-modulated envelope @ ~3 Hz	Sideband pair around f_r; < 5 Hz component	Vertically uniform bands with slow modulation
Tread Shelling / Polygon.	Multiple small impacts per revolution	Dense harmonic series f_r × 1–7+	Multiple thin vertical streaks; distributed
Combined Faults	All above superimposed	Composite spectrum; clearly additive	Multiple distinct patch types present

Figure 9: Fault-Feature Atlas — Diagnostic Signatures in Time, FFT, and STFT Domains.

5 WILDPulse Framework: Architecture and CNN Classification

5.1 Framework Overview

WILDPulse is organised as a four-layer modular pipeline above the existing WILD acquisition platform. Layer 1 — Signal Acquisition: existing WILD infrastructure acquires synchronised left- and right-rail signals at 10 kHz in CSV/XLSX format; no modification required. Layer 2 — Signal Processing: each segmented axle event is processed through time-domain conditioning (calibration, band-pass filtering, ILF computation), FFT magnitude computation, and STFT spectrogram generation (256-sample window, 87.5% overlap, Hanning taper), with the STFT rendered as a 224×224×3 RGB image for CNN input. Layer 3 — CNN Classification: the ResNet-18 CNN classifies each axle-event spectrogram into one of N health-state categories. Layer 4 — Explainable Analytics: the Streamlit-Python interface presents training curves, confusion matrices, class-wise metrics, and Grad-CAM attribution maps [9] revealing which spectral STFT regions drove each classification decision — essential for regulatory acceptance in safety-critical railway operations.

5.2 ResNet-18 CNN Architecture and Physical Interpretation

The CNN classification layer is based on a pretrained ResNet-18 architecture [7] (~11M parameters), selected for its established image classification capability and robustness to overfitting on moderate-sized datasets. Crucially, each stage corresponds to a physically meaningful level of abstraction in the STFT spectrogram (Figure 10, Table 3): Conv1 provides coarse localisation of high-energy time-frequency patches and detects sharp wheel-flat transients; Residual Stages 1–4 progressively encode impact shape and frequency spread, left–right strain asymmetry, composite multi-frequency patterns from combined conditions, and bogie-level suspension stiffness/damping descriptors; Global Average Pooling produces a compact 512-d condition embedding mapped via softmax to class probabilities feeding the THI.

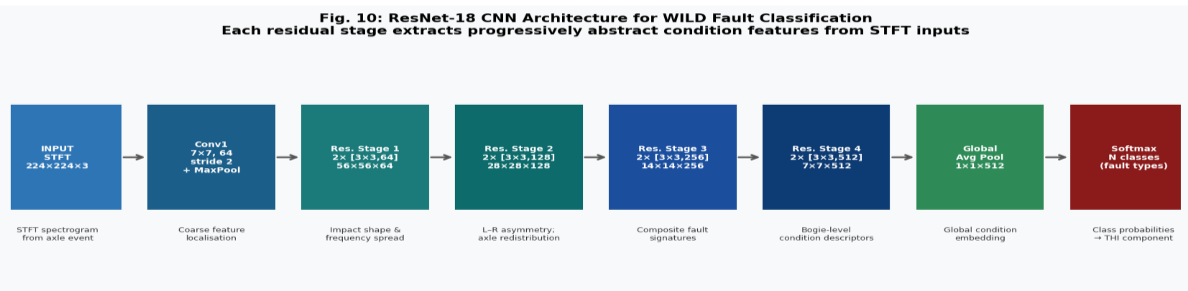

Figure 10: ResNet-18 CNN Architecture for WILDPulse Fault Classification. Five stages of progressively abstract feature extraction from 224×224×3 STFT spectrogram images to a 512-d condition embedding and softmax class probabilities.

CNN Stage	Output	Physical Interpretation for WILD Data
Conv1 (7×7, 64, stride 2) + MaxPool	56×56×64	Coarse localisation of high-energy time-frequency patches; detection of sharp wheel-flat transients
Residual Stage 1 (2×[3×3, 64])	56×56×64	Impact shape, frequency spread, and onset/decay characteristics of wheel-rail contact events
Residual Stage 2 (2×[3×3, 128])	28×28×128	Left–right strain asymmetry; axle-level load redistribution patterns indicative of canting
Residual Stage 3 (2×[3×3, 256])	14×14×256	Composite fault signatures; simultaneous multi-frequency energy patterns from combined conditions
Residual Stage 4 (2×[3×3, 512])	7×7×512	Bogie-level condition descriptors; suspension stiffness/damping ratio encoded in frequency-band ratios
Global Average Pooling → Softmax	512-d → N	Compact condition embedding; class probabilities mapped to Train Health Index components

Table 3: Physical interpretation of each ResNet-18 stage in the context of WILD spectrogram classification.

5.3 Training, Transfer Learning, and Validation Status

Transfer learning from an ImageNet-pretrained ResNet-18 initialisation accelerates training and improves generalisation. The final fully-connected layer is retrained from scratch; earlier layers are fine-tuned at a reduced learning rate. Validation status: train-configuration classification on field data achieves high accuracy over 30 training epochs with rapidly converging training loss and low overfitting; individual fault classification on physics-informed simulated data achieves >95% accuracy on a 6-class confusion matrix. Individual fault classification from labelled field WILD data — the critical next validation step — is being actively pursued through collaboration with Indian Railways maintenance depots and RDSO.

6 Train Health Index: Transformative Fleet Health Intelligence

6.1 Concept, Computation, and Longitudinal Tracking

The Train Health Index (THI) is the operationally transformative output of WILDPulse: a composite scalar score computed for each identified rake on each passage, encoding the overall rolling stock health state in a single trackable value. The conceptual contrast between pre- and post-WILDPulse operation is summarised in Table 4.

Before WILDPulse	After WILDPulse
Binary alarm: threshold exceeded or not	Continuous THI score per rake, per passage
No memory of prior train visits	Day-by-day degradation trend, visually trackable
Silent on suspension, cant, uneven loading	Suspension, canting, loading health all encoded in THI
Reactive maintenance; no trend awareness	Predictive alerts before failure; CBM-ready
One scalar (ILF) per axle per pass	Rich health vector per bogie and vehicle per pass

Table 4: Before/After WILDPulse: operational capability contrast.

The THI for rake r at passage time t is defined as:

$$THI(r, t) = w_1 \cdot ILF\_norm(r, t) + w_2 \cdot P\_healthy(r, t) + w_3 \cdot Asym\_norm(r, t) + w_4 \cdot Trend(r, t) \qquad (3)$$

where ILF_norm is the normalised axle-level ILF statistics; P_healthy is the CNN-computed healthy-class probability averaged across all axle events; Asym_norm is a normalised left-right load asymmetry index from paired channel ILF ratios; and Trend is the slope of the THI time-series over the preceding N_hist passages, capturing progressive degradation velocity. The THI is bounded to [0, 100]: values above 75 indicate Healthy operation; 50–75 trigger Watch (enhanced monitoring); below 50 triggers Action (priority inspection). Crucially, the THI's value lies in the longitudinal trend — a rake declining from 85 to 70 over 30 days may not have triggered any ILF alarm, yet the trend is a leading indicator warranting early scheduling of inspection. Figure 11 illustrates THI tracking for three rakes over a 90-day period: Rake A maintains healthy THI throughout; Rake B shows systematic decline from progressive suspension degradation; Rake C shows an acute event on Day 55 — immediately apparent in the THI history, invisible to ILF until failure.

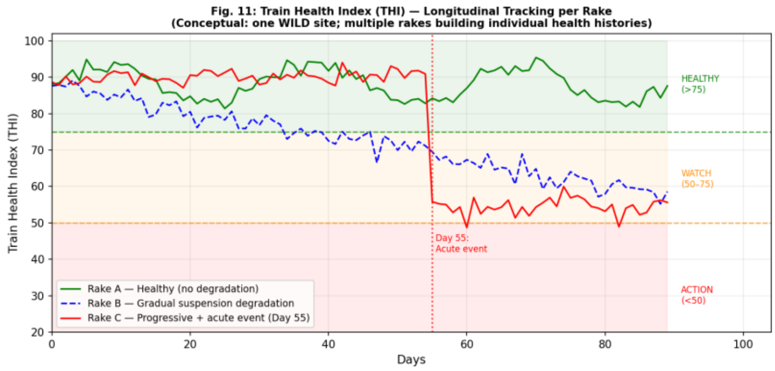

Figure 11: Train Health Index — Longitudinal Tracking per Rake (conceptual illustration). Rake A: stable healthy operation. Rake B: gradual degradation trending below Watch threshold. Rake C: acute event on Day 55. Green / orange / red zones = Healthy / Watch / Action THI ranges.

One WILDPulse-enabled site, tracking dozens of rakes daily at zero additional hardware cost, provides a fleet-wide health surveillance capability that transforms the maintenance planning function from reactive inspection scheduling to proactive, data-driven condition-based maintenance.

7 Limitations and Future Work

Field-validated components include the WILD hardware, ILF computation, wheel flat threshold detection, LabVIEW front-end (operationally deployed at ~25 locations), multi-domain signal processing pipeline (demonstrated on real field data), and CNN architecture (validated for train-configuration classification). Simulation-derived components requiring field validation include the individual fault signature characterisations of Sections 4.2–4.5, which are physics-informed modifications applied to the measured baseline signal — grounded in established railway dynamics literature, but not measurements of defective rolling stock. Field validation requires labelled datasets from verified fault cases, being actively pursued through collaboration with Indian Railways maintenance depots.

Key outstanding tasks include: collection and labelling of field WILD data for known fault conditions; multi-site CNN model validation across different WILD locations; optimisation and field calibration of THI sub-index weights against ground-truth inspection outcomes; incorporation of multi-channel (left-right) signal fusion; extension to freight wagon fault classification; integration with CRIS operational databases (FOIS, COA); and assessment of transfer learning strategies for other rail networks.

8 Conclusions

This paper has presented WILDPulse, a CNN-based rail classification platform for explainable fault detection in Weigh-In-Motion and rolling stock monitoring systems. Three core contributions have been described: (1) Physics-informed multi-domain signal analysis characterising how five distinct rolling stock fault conditions — wheel flats, primary and secondary suspension deterioration, axle box canting, and tread shelling/polygonisation — encode distinctive and visually separable signatures in time-domain, FFT, and STFT representations of WILD strain-gauge signals, providing an interpretable Fault-Feature Atlas. (2) CNN-based fault classification using a ResNet-18 convolutional neural network mapping STFT spectrogram images to rolling stock health-state categories, with each residual architecture stage physically interpreted in the WILD spectrogram domain; validated for train-configuration classification on field data and for fault classification on physics-informed simulated data. (3) Train Health Index providing longitudinal rake-level health tracking at zero additional hardware cost, converting the WILD system from a point-in-time defect detector into a continuous fleet health surveillance platform. The WILDPulse platform is immediately deployable as a software upgrade to any existing WILD installation, transforming binary alarm outputs into continuous fleet health intelligence and constituting a critical component of the pathway to condition-based maintenance across railway networks.

Acknowledgements

This work was supported by the Technology Mission for Indian Railways (TMIR), Ministry of Railways, and the Department of Science and Technology (DST), Government of India. The author acknowledges the technical collaboration of RDSO Lucknow in the field deployment and validation of the WILD system.

References

[1] K. Knothe, S. Stichel, "Rail Vehicle Dynamics", Springer, Berlin, 2017. https://doi.org/10.1007/978-3-319-45376-7

[2] M.J.M.M. Steenbergen, "The role of the contact geometry in wheel–rail impact due to wheel flats", Vehicle System Dynamics, 45(12), 1097-1116, 2007. https://doi.org/10.1080/00423110601143173

[3] A.M. Remennikov, S. Kaewunruen, "A review of loading conditions for railway track structures due to train and track vertical interaction", Structural Control and Health Monitoring, 15(2), 207-234, 2008. https://doi.org/10.1002/stc.227

[4] N.S. Vyas, A. Gupta, "Modeling Rail Wheel-Flat Dynamics", Proceedings of the World Congress on Engineering Asset Management (WCEAM 2006), Gold Coast, Australia, Springer Verlag, 2006.

[5] A. Alemi, F. Corman, G. Lodewijks, "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(8), 961-981, 2017. https://doi.org/10.1177/0954409716656218

[6] C. Li, S. Luo, C. Cole, M. Spiryagin, "An overview: modern techniques for railway vehicle on-board health monitoring systems", Vehicle System Dynamics, 55(7), 1045-1070, 2017. https://doi.org/10.1080/00423114.2017.1296963

[7] K. He, X. Zhang, S. Ren, J. Sun, "Deep residual learning for image recognition", Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, USA, 770-778, 2016. https://doi.org/10.1109/CVPR.2016.90

[8] A. Mosleh, P.A. Montenegro, P.A. Costa, R. Calcada, "Railway train-track monitoring and assessment with machine learning: a comprehensive review", Sensors, 24(8), 2451, 2024. https://doi.org/10.3390/s24082451

[9] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, "Grad-CAM: visual explanations from deep networks via gradient-based localization", International Journal of Computer Vision, 128(2), 336-359, 2020. https://doi.org/10.1007/s11263-019-01228-7

[10] R. Gadhave, N.S. Vyas, "Rail-wheel contact forces and track irregularity estimation from on-board accelerometer data", Vehicle System Dynamics, 60(6), 2145-2166, 2022. https://doi.org/10.1080/00423114.2021.1899253

[11] O.S. Ropalkar, O.P. Yadav, C. Zambare, N.S. Vyas, "Improved estimation of rail-wheel contact forces from instrumented wheel-set data through higher harmonic cancellation and a back-propagation neural network scheme", Engineering Applications of Artificial Intelligence, 116, 105458, 2022. https://doi.org/10.1016/j.engappai.2022.105458

[12] C.P. Ward, R.M. Goodall, R. Dixon, "Contact force estimation in the railway vehicle wheel-rail interface", IFAC Proceedings Volumes, 44(1), 4398-4403, 2011. https://doi.org/10.3182/20110828-6-IT-1002.01869
[13] M. Mohammadi, A. Mosleh, C. Vale, D. Ribeiro, P. Montenegro, A. Meixedo, "An unsupervised learning approach for wayside train wheel flat detection", Sensors, 23(4), 1910, 2023. https://doi.org/10.3390/s23041910
[14] N.S. Vyas, G.V.L.S. Kumar, "Data-based cross-functional digital twinning framework for rail transport efficiency enhancement", submitted to Rail 26 Conference, Budapest, Aug. 2026.