



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

RiahiSamani, R., Núñez, A., & De Schutter, B. (2026). WaveletInception networks for on-board vibration-based infrastructure health monitoring. *Engineering Applications of Artificial Intelligence*, 168, Article 113976. <https://doi.org/10.1016/j.engappai.2026.113976>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.



Research paper

WaveletInception networks for on-board vibration-based infrastructure health monitoring

Reza Riahi Samani ^a*, Alfredo Núñez ^b, Bart De Schutter ^a

^a Delft Center for Systems and Control (DCSC), Delft University of Technology, Mekelweg 2, 2628 CD, Delft, The Netherlands

^b Section of Railway Engineering, Department of Engineering Structures, Delft University of Technology, Stevinweg 1, 2628 CN, Delft, The Netherlands

ARTICLE INFO

Keywords:

On-board vibration signal
Wavelet inception
Feature fusion
Bidirectional gated recurrent unit
Infrastructure health monitoring

ABSTRACT

This paper presents a deep learning framework for analyzing on-board vibration response signals in infrastructure health monitoring. The proposed WaveletInception-BiGRU network uses a Learnable Wavelet Packet Transform (LWPT) for early spectral feature extraction, followed by one-dimensional Inception-Residual Network (1D Inception-ResNet) modules for multi-scale, high-level feature learning. Bidirectional Gated Recurrent Unit (BiGRU) modules then integrate temporal dependencies and incorporate operational conditions, such as the measurement speed. This approach enables effective analysis of vibration signals recorded at varying speeds, eliminating the need for explicit signal preprocessing. The sequential estimation head further leverages bidirectional temporal information to produce an accurate, localized assessment of infrastructure health. Ultimately, the framework generates high-resolution health profiles spatially mapped to the physical layout of the infrastructure. Case studies involving track stiffness regression and transition zone classification using real-world measurements demonstrate that the proposed framework significantly outperforms state-of-the-art methods, underscoring its potential for accurate, localized, and automated on-board infrastructure health monitoring.

1. Introduction

Efficient maintenance is fundamental to the sustainability, safety, and cost-effectiveness of transportation infrastructure. In 2020, for example, maintenance and renewal for railway infrastructure in Europe amounted to €10.47 billion and €11.24 billion, respectively, accounting for 52% of total infrastructure-related expenditures (European Commission, 2023). An effective maintenance strategy heavily relies on its underlying monitoring techniques. Early and accurate fault detection is crucial for minimizing downtime, extending service life, and reducing overall life cycle and maintenance costs. Traditionally, railway infrastructure health monitoring has relied on expert visual inspections or deploying sensing devices in strategic locations. However, these methods are often time-consuming, labor-intensive, and limited in scale. By overcoming such challenges, vehicle-based monitoring tools have gained significant attention (Fernández-Bobadilla and Martin, 2023). On-board vibration response analysis is a powerful and widely adopted vehicle-based monitoring technique that enables cost-effective and frequent health assessments across the entire infrastructure network under operational conditions. In particular, in the context of railway track and bridge infrastructures, vibration signals collected by accelerometers mounted on operational vehicles offer a cost-effective

and scalable means of assessing infrastructure conditions (Sensiñena et al., 2025; Phusakulkajorn et al., 2025a).

Several approaches have been reported in the literature for the analysis of on-board vibration response signals (Phusakulkajorn et al., 2023). At the core of this analysis lies feature extraction, a crucial step that involves identifying damage-sensitive features that can reflect the health status of the infrastructure. These features may be derived from the time and frequency domain Lederman et al. (2017), time-frequency domain Shen et al. (2023), Lamprea-Pineda et al. (2024), and space-frequency domain Unsiwilai et al. (2023), with the latter often utilizing wavelet transforms to capture spectral information. Once extracted, these features are commonly fed into machine learning classifiers for classifying vibration signals or employed for developing track quality indices (Shen et al., 2023; Unsiwilai et al., 2023). While these approaches offer physical interpretations, they are often limited by the complexity of selecting the most relevant features and requiring experts' experience in defining handcrafted features, which can hinder an automated process with high performance.

In recent years, deep learning methods for analyzing on-board vibration responses have gained significant attention. However, they still encounter key challenges. Most notably, existing approaches often

* Corresponding author.

E-mail address: r.riahisamani@tudelft.nl (R.R. Samani).

struggle to effectively capture the multi-scale, non-stationary characteristics of on-board vibration signals. Wavelet-inspired neural networks can decompose signals into time-frequency components, presenting a promising solution for extracting physically meaningful features from this data. Nonetheless, there is currently no fully automated, supervised learning framework that integrates these capabilities. Moreover, existing deep learning frameworks overlook the incorporation of operational factors, such as measurement speed, rail profile, or track alignment, when analyzing on-board vibration data, failing to account for the influence of these factors. Specifically, current methods typically have difficulty accommodating varying measurement speeds, which limits their generalization in real-world applications under diverse operational conditions. To address these research gaps, this paper introduces a novel deep learning methodology for monitoring the health of railway infrastructure using on-board vibration responses. The key contributions of this work include:

- A WaveletInception network is proposed for extracting features from vibration signals. It combines a learnable wavelet transform with one-dimensional Inception-Residual Network (1D Inception-ResNet) blocks to capture multi-scale information. The stem module directly incorporates spectral information, allowing the model to learn meaningful kernels in the early stages of learning. In the deeper layers, 1D Inception-ResNet blocks extract more nuanced features. This approach improves both accuracy and computational efficiency compared to existing feature extraction methods.
- A feature fusion strategy is proposed to integrate operational conditions, such as measurement speed, rail profile, and track alignment, into vibration signal analysis. These conditions strongly influence excitation levels, signal length, and spatial frequency content. To explicitly model their interaction with learned features, a bidirectional Gated Recurrent Unit (BiGRU) module is introduced within the feature extraction network, enabling mid-level fusion of operational information and bidirectional temporal dependencies, thereby enabling analysis under different operational conditions.
- The model processes time-domain signals of varying speeds and lengths without preprocessing, making it broadly applicable and suitable for automated monitoring.
- A bidirectional sequential modeling strategy is employed in the estimation stage to exploit contextual information from neighboring signal segments. By leveraging bidirectional temporal dependencies in on-board vibration responses, the proposed framework achieves more accurate infrastructure condition estimation and finer localized assessment compared with unidirectional recurrent architectures.
- The WaveletInception-BiGRU network improves estimation accuracy and computational cost over existing deep learning models for on-board vibration analysis.

The remainder of this paper is organized as follows. Section 2 reviews related work on on-board vibration signal analysis and recent deep learning advances. Section 3 presents the preliminaries. Section 4 describes the proposed model architecture. Section 5 evaluates the model using two case studies, including one based on real-world railway measurements. Sections 6 and 7 discuss the results, summarize key findings, and outline future research directions.

2. Related work

Infrastructure health monitoring using on-board vibration response requires methodologies for extracting meaningful information from the vibration signals. Modal parameters, such as natural frequencies, mode shapes, and damping ratios, are widely recognized as damage-sensitive features (Yang and Yang, 2018). In addition, stiffness is often estimated for structural health monitoring purposes. Papers like (Quirke et al.,

2017; Zhu et al., 2018) have used optimization techniques, including cross-entropy and adaptive regularization, to estimate railway track stiffness by matching vehicle-track interaction models to vibration data. However, these inverse identification methods are computationally intensive. For instance, estimating track stiffness for a 140-meter segment can take up to 11 h (Quirke et al., 2017), which makes them impractical for (pseudo) real-time monitoring. Although advancements in computational power have reduced simulation times, solving inverse problems in railway infrastructure remains time-consuming due to the numerous possible parameter combinations, structural configurations, operational scenarios, and the distributed nature of railway tracks. This complexity significantly limits the feasibility of using inverse methods for (pseudo) real-time monitoring.

Other approaches have focused on extracting informative features from vibration signals for infrastructure monitoring. These features, often compact representations of the original signals, are mostly spectral features derived from Fourier transforms (Caprioli et al., 2007), wavelet transforms (Shen et al., 2021, 2023; Unsiwilai et al., 2023), Kalman filters (Hoelzl et al., 2023), or Hilbert transforms (Malekjafarian et al., 2019a), which are sensitive to structural health conditions. For example, features can be extracted in the time domain Mao et al. (2025), time-frequency (Shen et al., 2023), or space-frequency domain Unsiwilai et al. (2023) by employing synchro-squeezed wavelet transforms, and continuous wavelets transform to construct localized wavelet power spectra. These approaches leverage the strength of wavelets to localize features along the signal. Moreover, statistical features derived from the power spectrum in the space domain are also robust to spatial frequency variations, making them practical for signals obtained at diverse measurement speeds. These features are used as health condition indicators, such as track quality indices (Unsiwilai et al., 2023), or as inputs to machine learning models like Gaussian process regression (Shen et al., 2023) for estimating stiffness parameters. Papers like (Zhang et al., 2025, 2024, 2022a; Peng et al., 2025; Khodadoost et al., 2025) further discuss the challenges and opportunities of data-driven and machine learning approaches in vibration-signal-based infrastructure monitoring.

While frameworks based on feature engineering effectively capture the spectral characteristics of signals, they heavily rely on carefully selected hyperparameters, such as window type, length, stride, and transformations (e.g., using log-mel spectrograms, which apply a logarithmic scale to Mel-frequency features). This reliance can pose challenges, as these hyperparameters significantly impact analysis accuracy, complicate full automation, and require time-consuming tuning (Cho et al., 2021; Zhang et al., 2022b). Additionally, the frequency components extracted do not directly provide physical insights and typically necessitate further modeling or expert interpretation to connect them to specific structural behaviors (Unsiwilai et al., 2024; Phusakulkajorn et al., 2023). However, supervised deep learning methods allow models to learn patterns from labeled data by minimizing prediction errors, leading to superior performance. These models can directly learn features and generalize more effectively than fixed wavelet coefficients, without the need for extensive fine-tuning and manual feature engineering.

Recent studies have increasingly investigated the use of supervised deep learning methodologies for monitoring the health of infrastructure through on-board vibration responses. These frameworks have been applied to various tasks related to infrastructure monitoring, including track stiffness estimation (Shen et al., 2023, 2021; Riahi Samani et al., 2025; Huang et al., 2022), as instantaneous decreases in moment of inertia for bridge elements (Locke et al., 2020), model bridge's dynamic changes due to additional weight (Hajializadeh, 2023), identification of seized bearings and cracking in the bridge beams (Corbally and Malekjafarian, 2024), and detecting transition zone characteristics (Phusakulkajorn et al., 2025b). Due to the high performance of deep learning methods in terms of computational efficiency and prediction accuracy, the literature indicates a growing application of these

techniques for on-board vibration-based monitoring (Phusakulkajorn et al., 2023; Zhang et al., 2025; Hajializadeh, 2023).

Deep learning frameworks for the on-board vibration signal analysis are typically CNN-based (Huang et al., 2022; Locke et al., 2020; Hajializadeh, 2023; Corbally and Malekjafarian, 2024; Phusakulkajorn et al., 2025c) or LSTM-based and hybrid architectures (Phusakulkajorn et al., 2025b; Samani et al., 2024). For instance, Locke et al. (2020) utilized a multi-layer CNN to classify bridge damage from simulated vehicle vibration signals. The model input includes the frequency response spectrum, adjusted to a spectral resolution of 0.1 Hz, combined with operational conditions like ambient temperature, speed, and vehicle mass. Locke et al. (2020) also demonstrated the robustness of CNNs to environmental and operational noise. Huang et al. (2022) employed a dilated CNN to analyze simulated axle-box acceleration (ABA) signals for predicting railway track dynamic stiffness. By training models with three input window sizes in the time domain and with measurement speed variations (60–120 km/h), the approach estimates track stiffness over different infrastructure lengths. The dilated CNN required much less computation time compared to standard CNNs without any significant enhanced accuracy (Huang et al., 2022). Transfer learning was explored in Hajializadeh (2023), which adapted GoogLeNet to classify vibration response spectrograms and to identify the location and intensity of damage. The model uses continuous wavelet transforms to generate time-frequency spectrograms from train-borne acceleration signals at varying speeds, followed by image resizing to standardize input dimensions for the 2-D GoogLeNet architecture, enabling the classification of vibration responses across different speeds. In Corbally and Malekjafarian (2024), an auto-calibrated vehicle-bridge interaction model is proposed to generate training data for damaged bridge scenarios and to train a CNN to classify damage type and location. The frequency spectrum from a laboratory-scale vehicle-bridge interaction model was interpolated and resampled at 0.05 Hz intervals using a cubic spline function to standardize input sizes. The algorithm accurately detected seized bearings and cracks in bridge beams, though crack location accuracy decreases at lower damage levels. In our previous work (Riahi Samani et al., 2025), we developed an LSTM-BiLSTM network to estimate track stiffness parameters from on-board vibration signals. The model was trained on simulated vehicle-track interactions at a single speed in the time domain and demonstrated accurate stiffness parameter estimation. It was also highlighted that the importance of incorporating temporal relationships in feature extraction and health condition estimation architectures of deep learning models for on-board vibration response analysis (Riahi Samani et al., 2025).

Most deep learning frameworks, particularly those based on convolutional neural networks (CNNs), require fixed input sizes. However, in on-board measurements, varying operational speeds result in signals of different lengths and spatial sampling frequencies for the same infrastructure segment. This necessitates preprocessing steps such as spectrogram generation, followed by spectrogram resizing or resampling, which can introduce limitations. Generating a spectrogram involves careful selection of hyperparameters, such as window type, length, and stride, all of which can influence the performance of the model. Additionally, the resizing or resampling of spectrograms may lead to information loss or the creation of synthetic data points, which can degrade the resolution of spectral features and affect overall accuracy. To address these challenges, the current paper proposes a fully automated methodology that processes data directly in the time domain. This approach effectively handles the variable-length inputs caused by different measurement speeds. Papers (Michau et al., 2022; Frusque and Fink, 2024) already introduced DeSpaWN and learnable wavelet packet transform (LWPT) networks for unsupervised high-frequency signal reconstruction, which can accommodate variable-length inputs. However, these methods do not specifically address the root cause of the variations, measurement speed. In contrast, our supervised learning framework not only incorporates wavelet-based analysis but also explicitly integrates measurement speed into the analysis.

Wavelet-inspired deep learning architectures are increasingly utilized to integrate spectral information into models across various fields. For example, Liu et al. (2019) introduced a multilevel Wavelet-CNN for image restoration, leveraging the Discrete Wavelet Transform for down-sampling and the Inverse Wavelet Transform for up-sampling. In Li et al. (2022), WaveletKernelNet CNN architectures were proposed for industrial diagnosis, where a continuous wavelet convolutional layer replaces the first standard CNN convolutional layer. This can provide a physically meaningful input layer, improving training performance and model accuracy (Li et al., 2022). Wavelet-informed deep learning frameworks have been applied to a variety of tasks in several application fields, such as ECG signal classification (bouny et al., 2020), image classification (Yao et al., 2022), and time series analysis (Wang et al., 2018a). These papers demonstrate that incorporating spectral analysis techniques, such as wavelet transforms, can enhance deep learning performance by providing physically meaningful information derived from wavelet decompositions. In the current paper, we propose a novel WaveletInception which integrates spectral information in the early stages via an LWPT module and ensures multi-scale feature extraction with 1D Inception-ResNet modules. We further employ BiGRU networks to account for operational conditions, particularly measurement speed, and to capture temporal dependencies in the analysis of vibration response signals. Sections 3 and 4 provide methodological preliminaries and further details.

3. Preliminaries

3.1. Discrete wavelet transform

The wavelet transform is a powerful mathematical tool for analyzing signals in the time and frequency domains (Mallat, 1999). Unlike the Fourier transform, which provides only a frequency-domain representation, the wavelet transform offers a multi-resolution analysis by decomposing signals into components at various scales. This capability makes it particularly effective for analyzing non-stationary signals, such as on-board vibration signals.

The Discrete Wavelet Transform (DWT) is a computationally efficient version of the wavelet transform, widely used for signal and image processing. The basic operation of DWT is the convolution operation of a low-pass and a high-pass filter, followed by a downsampling by a factor of two, which can be denoted as

$$\begin{cases} (x^l)_n = (f^l * x)_{2n} \\ (x^h)_n = (f^h * x)_{2n} \end{cases} \quad (1)$$

where x^l and x^h correspond to the low-pass (approximation) and high-pass (details) filtered input data with a downsampling rate of 2, f^l and f^h are low-pass and high-pass filters, $*$ is the convolution operation, x is the input signal with the signal length of $2n$. The perfect reconstruction and orthogonality property of the DWT ensure that all information is retained in its coefficients despite the downsampling (Liu et al., 2019; Mallat, 1999).

3.2. Discrete wavelet packet transform and LWPT

The Discrete Wavelet Packet Transform (WPT) provides a multi-resolution analysis by decomposing a signal into uniform frequency bands through a hierarchical, tree-like structure (Mallat, 1999). At each level L , the signal is convolved with low-pass and high-pass filters (Eq. (1)), splitting the spectrum into finer sub-bands. While this process doubles the frequency resolution at each step, it reduces the temporal resolution by half. The result is a comprehensive decomposition where each leaf node in the WPT tree (see Fig. 1) represents a specific frequency interval.

The Learnable Wavelet Packet Transform (LWPT), originally introduced in Michau et al. (2022), Frusque and Fink (2024), initializes its filters using standard DWT bases, such as Haar or Daubechies ($db4$),

but refines them dynamically via backpropagation. Unlike traditional WPT, which relies on fixed coefficients that may not align with the complex transients of railway vibrations, LWPT learns the optimal filter bank directly from the data. This approach leverages the perfect reconstruction and orthogonality of wavelets to downsample high-frequency vibration signals without the information loss typical of standard pooling layers (Michau et al., 2022; Frusque and Fink, 2024). Furthermore, the LWPT incorporates an adaptive denoising mechanism that exploits coefficient sparsity; by utilizing a differentiable double sharp sigmoid activation function, the model applies learnable thresholds to suppress stochastic noise while preserving salient structural transients (Michau et al., 2022). Compared to the DWT and WPT, this adaptive capability ensures that refined, noise-reduced mechanical signatures are effectively propagated through the network.

In the current paper, the LWPT module is coupled with a 1×1 channel-projection layer to adapt the spectral output to a supervised learning framework effectively. Furthermore, the model is evaluated both with and without the denoising function to empirically assess the impact of adaptive thresholding on health-condition estimation.

3.3. 1DCNN

A Convolutional Neural Network (CNN) is a deep learning architecture designed to extract spatial and temporal features from input data using convolutional operations with filter banks (Aggarwal et al., 2018). While CNNs are commonly associated with image data, their 1D variants are particularly effective for processing sequential or time-series data, such as vibration and audio signals. A 1D CNN operates on one-dimensional signals, with convolutional kernels in a filter bank sliding along the temporal dimension. These networks are designed to capture local patterns or dependencies in the input signal. This process generates feature maps that capture abstract representations of the input and highlight different aspects of the input data. The formula for a typical convolutional layer is

$$y_f = \text{conv1D}(W_f, x) + b_f \quad (2)$$

where y_f , W_f , and b_f denote the output vector, weight vector, and bias parameter of filter f , respectively; x is the input time sequence; and conv1D represents the one-dimensional convolution operator. The i th output of the convolution is computed as:

$$[\text{conv1D}(W_f, x)]_i = [W_f * x]_i = \sum_{j=0}^{N_f-1} w_{f,j} \cdot x_{i+j} \quad (3)$$

where N_f is the length of the filter f , $w_{f,j}$ is the j th element of the vector W_f , and i is the starting index in the input for each convolution window.

3.4. Gated Recurrent Unit (GRU)

Gated Recurrent Units (GRUs) are a streamlined variant of Recurrent Neural Networks designed to capture temporal dependencies while mitigating the vanishing gradient problem (Cho et al., 2014; Chung et al., 2014). Unlike the LSTM, the GRU achieves an efficient information flow using a simplified architecture that merges the forget and input gates into a single update gate and utilizes a reset gate to regulate the integration of previous hidden states.

Given an input sequence $\{x_t\}_{t=1}^T$, the GRU updates its hidden state h_t at each time step t by combining the current input x_t with the previous hidden state h_{t-1} . The update gate determines how much of the previous information is carried over to the current hidden state, while the reset gate decides how much of the past context should be discarded when computing the new candidate state. This mechanism allows the model to retain long-term dependencies effectively with fewer parameters than traditional LSTMs.

While a standard GRU processes data in a forward direction, a Bidirectional GRU (BiGRU) adds a backward pass, capturing both past

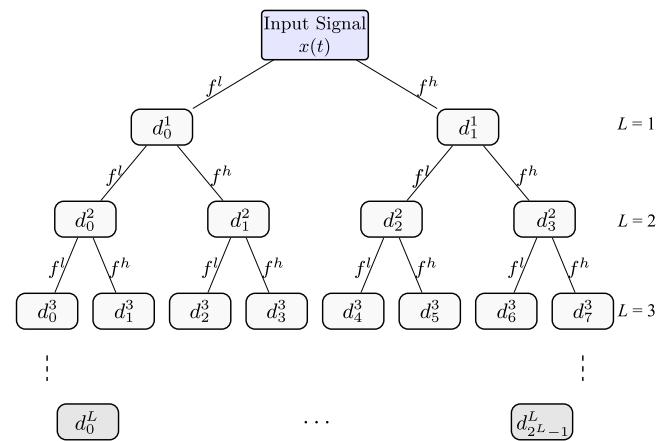


Fig. 1. Tree-like structure of the Wavelet Packet Transform (WPT).

and future signal context. This comprehensive temporal perspective is especially valuable for vibration-based infrastructure health monitoring (Riahi Samani et al., 2025).

3.5. Multi-modal data fusion

Multi-modal data fusion combines data from different modalities, such as vibration signals, LiDAR, audio, and speed measurements, each of which provides complementary insights into the same phenomenon. Integrating multiple data sources can result in more robust, accurate, and generalizable models (Lahat et al., 2015). Existing multi-modal fusion methods can be categorized into three main types: feature-level fusion, decision-level fusion, and hybrid fusion (Lahat et al., 2015). Feature-level fusion entails extracting features from each modality and merging them into a unified representation. In contrast, decision-level fusion combines the outputs of separate models trained on individual modalities. Hybrid fusion integrates feature and decision levels, leveraging the strengths of each approach (Lahat et al., 2015).

Feature fusion is widely used across various domains to exploit correlations among multiple input modalities, sensors, or extracted features. This method leverages inherent correlations among features across modalities, thereby enhancing training. However, it does face challenges related to temporal synchronization, as features from closely related modalities may be sampled at different time instants (Zhang et al., 2019). For instance, in Dang et al. (2021), a hybrid 1D CNN-LSTM network is proposed that combines autoregressive features, discrete wavelet features, and empirical mode decomposition features from multiple sensors into a unified representation for model training. Similarly, Mou et al. (2021) utilized an attention mechanism to integrate features extracted by a CNN-LSTM network from three distinct input sources: environmental data, vehicle data, and driver eye data.

In the current paper, we propose a feature-fusion approach that combines features from vibration-response signals and operational-condition embeddings. We use a BiGRU layer to capture temporal and interrelated dependencies within the feature-extraction network, as explained below.

4. WaveletInception-BiGRU networks

4.1. Problem statement

Effective Structural Health Monitoring (SHM) via on-board vibration signals requires multi-resolution analysis to identify defects across

various frequency bands (Shen et al., 2023; Lamprea-Pineda et al., 2024). Standard CNNs, while excellent at spatial pattern recognition, lack an inherent mechanism for multi-resolution spectral decomposition, and their effectiveness in performing spectral analysis is still not well understood (bouny et al., 2020). Furthermore, real-world measurement conditions — specifically varying measurement speeds — induce non-stationarity in vibration signals, leading to inconsistent excitation levels and varying signal lengths. Conventional preprocessing, such as fixed-ratio resampling or rigid resizing, often leads to information loss or fails to account for the physical relationship between measurement speed and vibration intensity. Additionally, most existing infrastructure health monitoring frameworks provide aggregated health indices over long segments (e.g., 30–200 m) (Unsilwilai et al., 2023, 2024). Such coarse resolution hinders localized maintenance and component-level defect identification.

To address these challenges, we propose the WaveletInception-BiGRU framework, which integrates three key components: a Wavelet-Inception module for learnable spectral decomposition and multi-scale feature extraction from vibration signals; a feature fusion technique to incorporate operational conditions; and a BiGRU network for sequential modeling and capturing temporal relationships. This combination enables high-resolution, component-level health estimation by leveraging both physical and operational insights.

4.2. WaveletInception-BiGRU

The proposed architecture consists of two main stages: (I) a Wavelet-Inception stage for multi-scale spatial feature extraction, and (II) a BiGRU-based temporal head for multi-modal fusion and sequential health conditions estimation. This design handles variable-length inputs while maintaining a physically meaningful relationship between the signal features and the infrastructure's spatial layout. Fig. 3 shows an overview of the proposed WaveletInception-BiGRU networks

4.2.1. WaveletInception feature extraction networks

The WaveletInception stage serves as the primary feature extractor, designed to handle the high sampling rates characteristic of vibration signals. It comprises an LWPT followed by a channel projection layer, and Inception-ResNet networks.

The WaveletInception network begins with an LWPT stem that hierarchically decomposes the raw vibration signal into 2^L sub-bands (where L denotes the decomposition level). This process utilizes learnable high- and low-pass filters to extract multi-resolution features while simultaneously downsampling the temporal dimension by a factor of two at each level. Thanks to its orthogonality, LWPT enables high-frequency signals to be downsampled without information loss (Michau et al., 2022). The filters are initialized with Discrete Wavelet Transform (DWT) bases, such as Haar or Daubechies db4 (Frusque and Fink, 2024), and are then refined through backpropagation. This initialization provides a physically grounded starting point for time-frequency decomposition, where db4 filters, in particular, offer smoother signal boundary handling for zero-padded regions compared to Haar filters.

Since the diagnostic relevance of wavelet sub-bands varies significantly across different operational conditions, a 1×1 convolution layer is implemented as a learned channel projection. This layer replaces manual sub-band selection by serving as an automated weighting mechanism that identifies and prioritizes the most informative spectral components while suppressing noisy channels. By learning these projections during training, the network adaptively optimizes the integration of disparate frequency information. Following this projection, two sequential 1D convolutional layers further refine the spatial-spectral features and downsample the temporal dimension. This process optimizes the feature density and ensures that only the most salient local patterns are propagated to the subsequent 1D Inception-ResNet backbone.

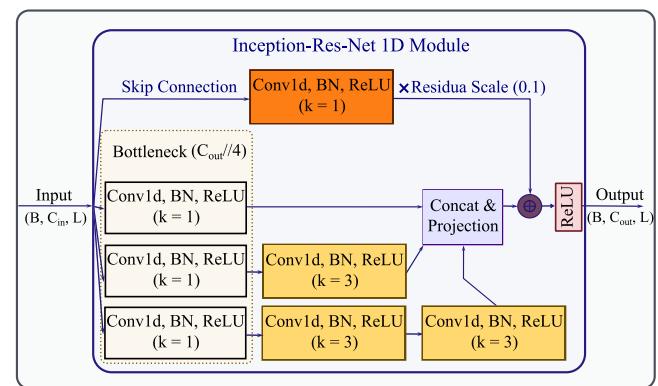


Fig. 2. 1D Inception block structure.

The 1D Inception-ResNet network efficiently extracts high-level features from on-board vibration signals. As shown in Fig. 2, these blocks utilize parallel convolutional branches to capture varied signal patterns while leveraging residual connections to stabilize training in deep architectures. To minimize computational overhead, bottleneck layers are integrated, which effectively refine fine-grained features without excessive parameterization. Unlike conventional 2D models that require complex signal-to-image transformations, these 1D modules operate directly on the temporal axis. This not only improves efficiency but also ensures physical interpretability, as the learned filters act as adaptive digital signal processors that correspond to the time-domain characteristics (Borghesani et al., 2023) of the on-board vibration response.

The hyperparameters of the WaveletInception networks consist of the number of decomposition levels (L) for the LWPT, the number of Inception-ResNet blocks, and the number of output channels of the convolution layers. Table 2 shows the fine-tuned WaveletInception networks after hyperparameter optimization for the two case studies conducted.

4.2.2. BiGRU networks, feature fusion and temporal modeling

Our proposed BiGRU networks consist of two BiGRU modules as shown in Fig. 3. First, the BiGRU module extracts bidirectional temporal features from the concatenated vibration signal features and the operational-condition embedding, in particular, the speed embedding. The vibration features are the outputs of the WaveletInception network, and the measurement speed is processed through a lightweight Multi-Layer Perceptron (MLP). This transforms the scalar speed into a high-dimensional speed embedding that matches the temporal length of the vibration features.

The first BiGRU layer fuses the vibration features with the speed embedding. This layer integrates temporal dependencies into the feature-extraction process and captures interrelationships in the concatenated features of the vibration signal and operational conditions, such as signal length variation and excitation level. Additionally, the BiGRU layer accounts for the relationships among measurement speed, signal length, and padding length, thereby further refining the learned features.

To achieve component-level resolution, we implement a distributed sampling strategy designed to extract N feature vectors from the BiGRU outputs for sequential modeling. In this approach, both the value of N and the specific indices for feature extraction are determined by the physical layout of the infrastructure. For instance, if an infrastructure segment contains 30 equidistant sleepers, N is set to 30, and feature vectors are extracted at uniform intervals from the temporal sequence to align with the sleeper positions. This physical-temporal mapping ensures that each network output corresponds to a distinct structural component. By incorporating prior knowledge of the infrastructure configuration, the model maintains a one-to-one correspondence with

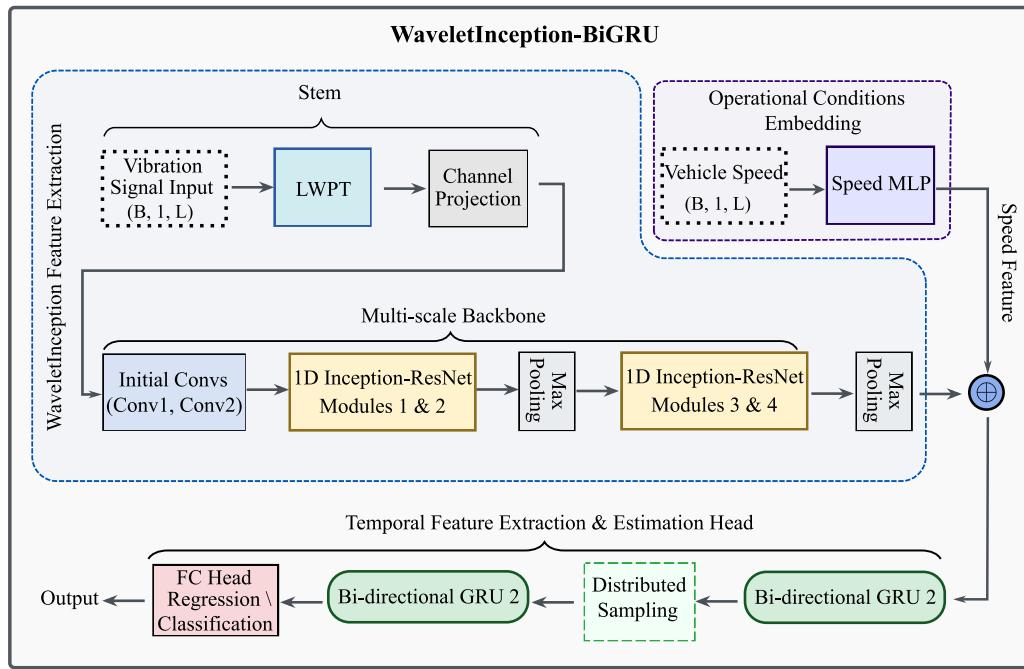


Fig. 3. Overview of the WaveletInception-BiGRU architecture, including WaveletInception vibration signal feature extraction, operational conditions feature fusion, and BiGRU sequential health conditions estimation.

the physical layout, thereby significantly improving the accuracy of high-resolution estimation. The selected feature vectors pass through a dropout layer before being fed into the final temporal modeling in the health estimation head.

Finally, these selected vectors pass through a dropout layer prior to temporal modeling in the health estimation head.

The second BiGRU module serves as a sequential model for the estimation head in monitoring infrastructure health conditions. In this context, sequential modeling captures spatial dependencies across a sequence of beams along the infrastructure, with each beam corresponding to a single time step. The BiGRU layer processes the input at each time step and captures temporal dependencies along the sequences. The BiGRU outputs are passed through fully connected layers to map the learned representations onto specific health-condition indicators, such as regression (as in Case Study I) and classification (as in Case Study II) tasks. This sequential modeling enables detailed predictions at the component level, corresponding to the structure's physical layout. Consequently, this method improves the resolution of health condition estimations and facilitates comprehensive assessment of individual components, providing valuable insights for maintenance planning.

The proposed framework utilizes BiGRU to capture comprehensive temporal context by concatenating hidden states from both forward and backward passes. Physically, the health condition of infrastructure elements — like railway sleepers or bridge beams — is characterized not only by their local dynamic response but also by mechanical signatures propagated through adjacent components. This mechanism enables the model to identify complex, long-range dependencies between fused vibration features and operational conditions, such as the measurement speed. By integrating multi-modal inputs from both preceding and succeeding spatial steps, the BiGRU layers effectively capture the global context inherent in on-board vibration data. Consequently, this bidirectional awareness markedly improves estimation performance relative to unidirectional models, which fail to incorporate future signatures from downstream vibration signals to refine the current health-condition assessment.

The hyperparameters of the temporal modeling and estimation head network include the size of the operational condition embedding module, the number of BiGRU units in the BiGRU layer, the size of the fully

connected layers, and the dropout ratio. **Table 2** shows the parameters fine-tuned in the two case studies.

5. Illustrative case study of railway track health monitoring

The Axle-Box Acceleration (ABA) measurement system is a cost-effective drive-by inspection method used on operational trains to capture vibration signals through accelerometer sensors mounted on axle boxes (Phusakulkajorn et al., 2025d; Molodova et al., 2014). This system allows for the development of a fully automated framework to analyze ABA vibrations, enhancing railway infrastructure monitoring over time and space.

To demonstrate this methodology, we present two case studies: (I) railway track stiffness estimation and (II) railway transition zone identification. In both cases, time-domain ABA signals serve as inputs to our deep learning architecture, which integrates preprocessing, feature extraction, and estimation.

5.1. Case study I: Railway track stiffness estimation

Railway track stiffness is a key indicator of infrastructure health, primarily influenced by the stiffness of the ballast substructure and fastening components like bolts, clamps, and rail pads. Estimating the stiffness of these components offers insights into track condition. For instance, defects in railway sleepers, such as hanging sleepers or ballast crushing, can significantly reduce track stiffness (Shi et al., 2023).

In this paper, we use simulated ABA vibration responses to estimate track stiffness parameters, characterized by rail pad stiffness and ballast stiffness, represented as the vector $k = [k_p, k_b]^\top$.

5.1.1. Data

We simulate vehicle-track interactions over segments of 10 sleepers at speeds of 35 km/h, 50 km/h, 65 km/h, 80 km/h, and 90 km/h, under four stiffness variation scenarios. The dataset is generated using a finite-element vehicle-track interaction model implemented in MATLAB, where the rail and sleepers are modeled as Timoshenko beam elements supported by rail pads and ballast with stiffness and damping properties. The vehicle is represented as a wheelset with Hertzian

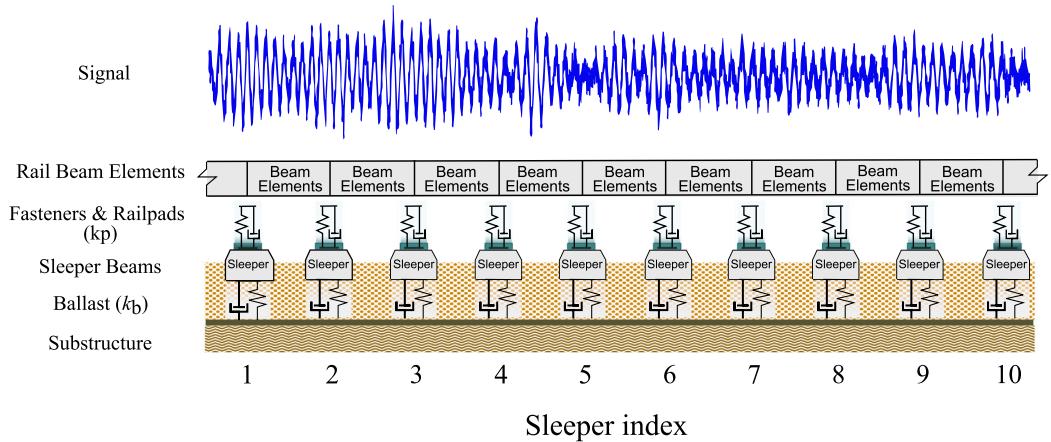


Fig. 4. The layout of the 10-sleeper track segment.

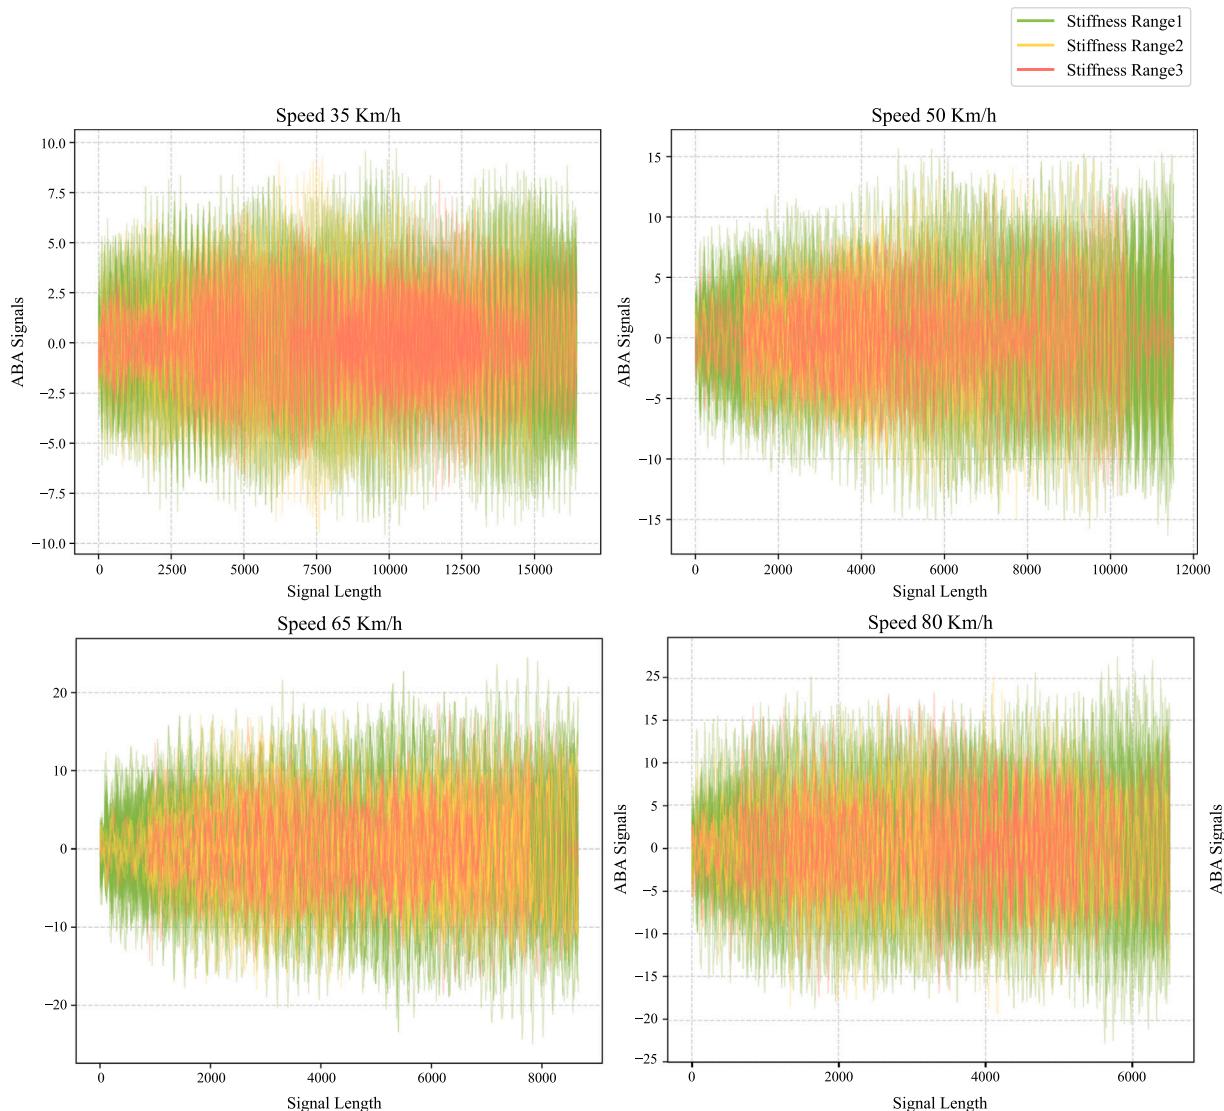


Fig. 5. Representation of the simulated ABA signals in four different measurement speeds. ABA signals obtained over the stiffness range sets R1, R2, and R3 are plotted in green, yellow, and red, respectively.

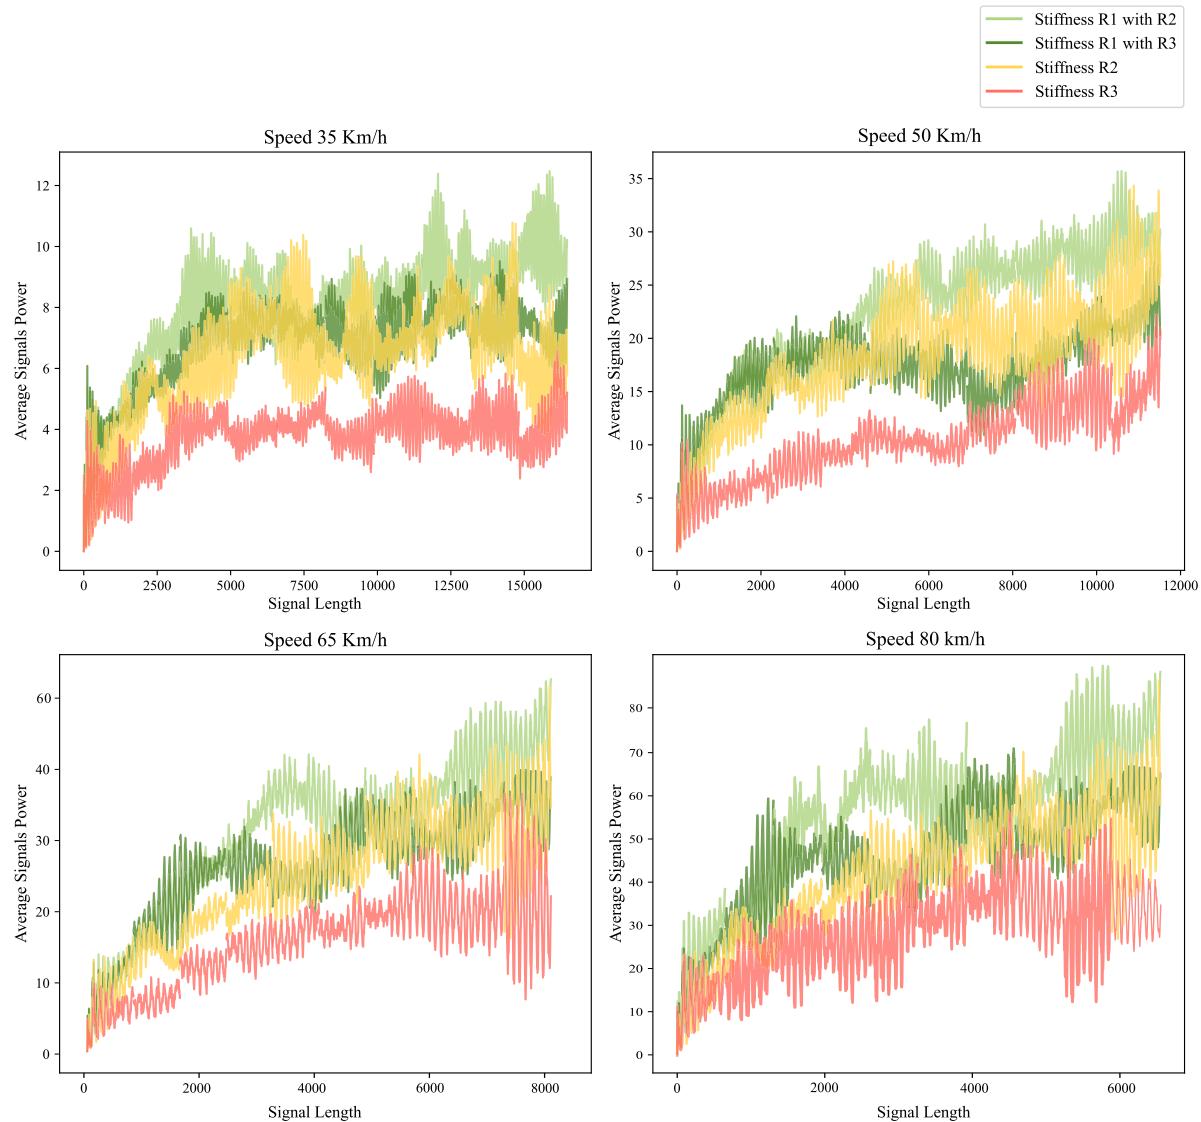


Fig. 6. Impact of local stiffness reduction on long-range ABA signal power, showing average signal power for three sleepers with stiffness range R1 reduced to ranges R2 (light green) and R3 (dark green), alongside original ranges R2 (yellow) and R3 (red).

contact for wheel-rail interaction. Further details on the model can be found in [Shen et al. \(2023\)](#). [Fig. 4](#) shows a schematic of the track segment and the corresponding ABA measurements for each combination of speed and stiffness scenario.

The four stiffness scenarios include:

(I) Uniform track stiffness: Normal conditions with stiffness values randomly sampled from either the nominal (R1) or moderate (R2) stiffness ranges in [Table 1](#).

(II) Local stiffness reduction in one sleeper: Simulates a defect like a hanging sleeper, with stiffness values set to R3 representing varying fault severities.

(III) Local stiffness reduction in three sleepers: Models reduced support over a longer section, with values sampled uniformly from R1 to R2 or R3.

(IV) Transition zones: Simulates stiffness variations between the two halves of segments, with values randomly sampled from R1 and R2.

[Fig. 5](#) illustrates the ABA signals across stiffness ranges (R1, R2, R3) and speeds. Lower stiffness correlates with lower signal amplitude, whereas higher speeds increase signal amplitude. [Fig. 6](#) illustrates the average signal power across different stiffness ranges, specifically in the context of localized stiffness reductions at three sleepers. The

figure focuses on stiffness range R1 while examining adjacent reductions in either range R2 (light green) or range R3 (dark green). The figure shows that a local reduction in stiffness affects the ABA signals recorded in nearby sections, even when those sections remain within the stiffness range R1. This observation supports our choice of methodology: employing a sequential BiGRU-based model to capture long-range dependencies within the data effectively.

The dataset comprises 12500 records, with 2500 for each measurement speed across different scenarios. To simulate real-world conditions, Gaussian noise, modeled as $\sim N(0, \sigma^2)$, is added to the ABA signals, achieving a noise-to-signal ratio of 5%. This noise variance accounts for 5% of the ABA signal power, replicating measurement noise from environmental factors and vehicle disturbances ([González et al., 2023](#); [Malekjafarian et al., 2019b](#)). The dataset is divided into training (70%), validation (15%), and testing (15%) sets ([Bishop and Bishop, 2023](#)).

5.2. Case study II: Transition zones identification using real-world data

Transition zones near viaducts, bridges, tunnels, and level crossings are critical parts of railway infrastructure. They often experience abrupt changes in track stiffness, e.g., when shifting from ballast to

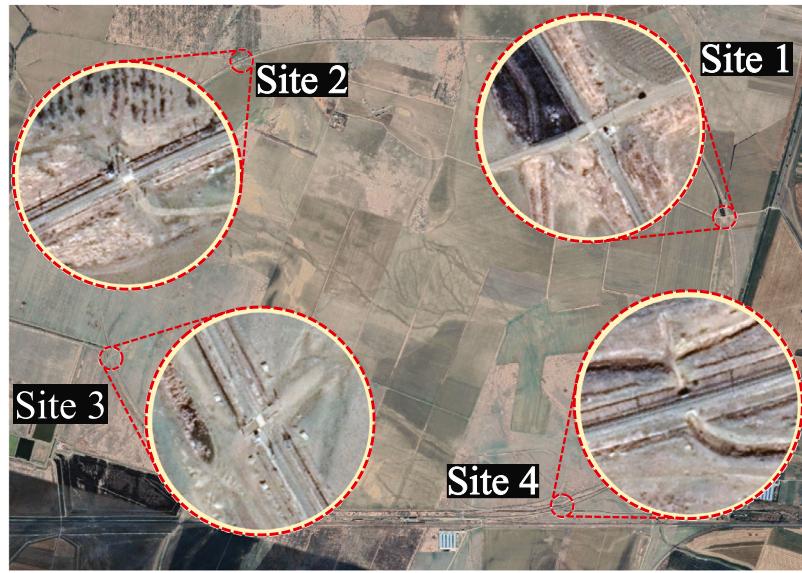


Fig. 7. Case Study II map, four transition zone sites selected from the Romania test ring. Maps are adapted from Google Maps 2025.

Table 1
Range of parameter values used for railpad and ballast stiffness.

| Range sets | k_p (N/m) | k_b (N/m) |
|------------|-----------------------------------|-----------------------------------|
| R1 | $2.0 \cdot 10^8 - 3.0 \cdot 10^8$ | $1.6 \cdot 10^7 - 2.2 \cdot 10^7$ |
| R2 | $1.0 \cdot 10^8 - 2.0 \cdot 10^8$ | $1.0 \cdot 10^7 - 1.6 \cdot 10^7$ |
| R3 | $0.1 \cdot 10^8 - 1.0 \cdot 10^8$ | $0.4 \cdot 10^7 - 1.0 \cdot 10^7$ |

slab track (Phusakulkajorn et al., 2025b; Wang et al., 2018b; Wang and Markine, 2018). These stiffness variations can lead to increased dynamic loads, differential settlements, and accelerated track degradation, posing challenges for maintenance and safety (Wang et al., 2018b). Accurate identification and monitoring of these zones are vital for the reliability of railway networks. In this case study, we apply our proposed framework to pinpoint regions within transition zones at a sleeper-level resolution, using real-world measurements of ABA signals collected at different speeds.

5.2.1. Data

This case study uses data from the Făurei testing ring in Romania, operated by the Romanian Railway Authority (AFER). ABA measurements are collected over transition zones near crossing levels and viaducts during dynamic testing at speeds ranging from 20 km/h to 140 km/h. Fig. 7 shows the test ring map and the four selected transition zones for this paper.

The ABA system on the measurement train collects 16 vibration responses from the two wheelsets during each round. These signals are extracted over 18-m track segments, focusing specifically on transition zones. Data from sites 1–3 are utilized for model training (85%) and validation (15%), totaling 7000 ABA samples. To rigorously assess generalization, site four is reserved as an entirely unseen test set. This set, comprising 107 transition-zone and 110 normal-track samples, is used exclusively for final evaluation and deployment. Such a separation ensures an unbiased estimate of model performance (Bishop and Bishop, 2023).

Fig. 8 presents the average short-time Fourier transform (STFT) spectrograms of the ABA signals from three consecutive segments, before, during, and after the viaduct, at Site 4. The left spectrogram covers the interval from 1.2102 km to 1.2282 km. The middle segment includes the viaduct located between 1.2348 km and 1.2402 km. The right segment extends from 1.2462 km to 1.2642 km.

Each spectrogram illustrates the changes in vibration frequency content below 100 Hz. The plots indicate that the dominant frequency components associated with variations in substructure stiffness are primarily concentrated around 10 Hz. This frequency range in the figures representing conditions before and after the viaduct suggests that changes in track stiffness can influence ABA responses over a distance, highlighting the presence of long-range spatial dependencies.

These findings highlight the necessity of modeling ABA signals as sequential data with long-range dependencies, which allows the learning framework to capture both localized features and their broader impact on the characteristics of ABA signals along the track.

5.3. Training, validation, and test design

Fixed training, validation, and test sets are used to evaluate the models. The mean squared error (MSE) loss for Case Study I and the binary cross-entropy with logits loss (BCEWithLogitsLoss) for Case Study II are minimized using the Adam optimizer (Kingma and Ba, 2017). To enhance loss convergence, we implemented a learning rate warm-up, followed by a scheduler with a factor of 0.5 and a patience of 6. For example, Fig. 9 shows the loss convergence of the models on the training and validation sets in Case Study I. We tuned hyperparameters — including the number of layers, units, dropout rates, and learning rates — during the validation process. In particular, random search in Weights& Biases (Biewald, 2020) is used to explore 50 hyperparameter configurations for each architecture, efficiently covering high-dimensional spaces without the combinatorial cost of exhaustive grid search (Bergstra and Bengio, 2012). Table 2 shows the configuration of the fine-tuned architectures for both case studies.

A systematic comparison of the proposed WaveletInception-BiGRU model is conducted across external baselines and its own ablations. The analysis mainly consists of three key components: preprocessing (stem), speed-data fusion, and temporal feature extraction. In the preprocessing phase, several techniques are evaluated to handle variable-length input signals. These techniques include interpolation-based representations using the Continuous Wavelet Transform (CWT) and Short-Time Fourier Transform (STFT), as well as zero-padding the signals followed by the LWPT, STFT, and length-aware modules. For speed integration, two mid-level data fusion techniques, early and late fusion, are examined. Temporal feature extraction is compared across LSTM, GRU,

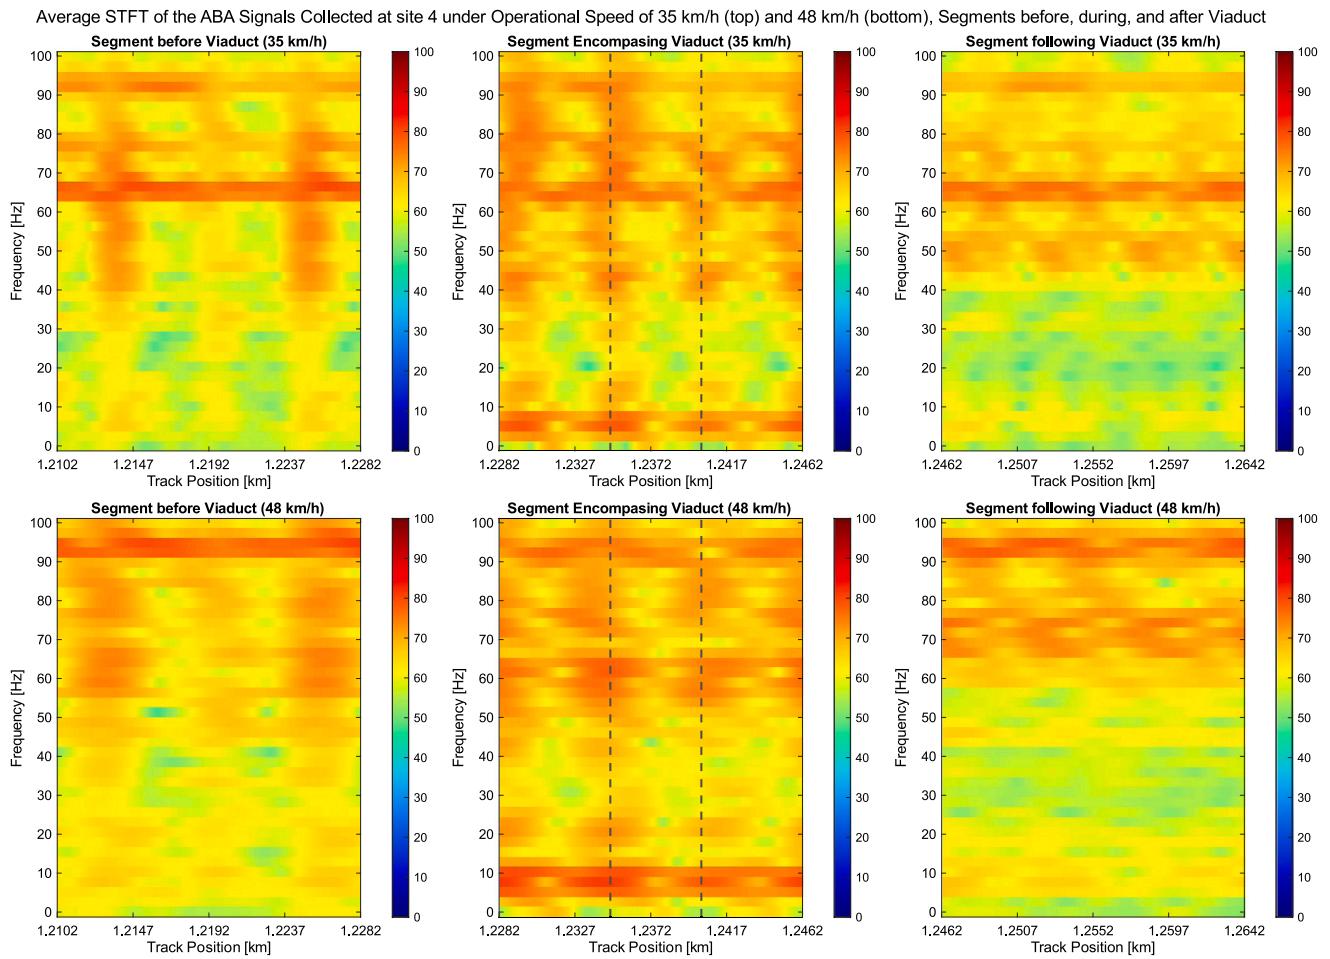


Fig. 8. Influence of the local change in track stiffness on a longer range of STFT spectrogram in the low frequency. The viaduct is located between 1.2348 km and 1.2402 km.

Table 2
Wavelet-Inception-BiGRU architecture fine-tuned for Case Studies I and II.

| Network stage | Output tensor size (Batch \times Channels \times Length) | |
|-----------------------------------------|--------------------------------------------------------------|------------------------------|
| | Case Study I | Case Study II |
| Input vibration signal | 32 \times 1 \times 16460 | 32 \times 1 \times 78150 |
| LWPT stem (level L) | 32 \times 64 \times 258 | 32 \times 128 \times 611 |
| Channel projection (Conv1D, $k = 1$) | 32 \times 96 \times 258 | 32 \times 128 \times 611 |
| Conv1D ₁ | 32 \times 96 \times 258 | 32 \times 96 \times 306 |
| Conv1D ₂ | 32 \times 96 \times 258 | 32 \times 96 \times 306 |
| Inception-ResNet module 1 | 32 \times 64 \times 258 | 32 \times 128 \times 306 |
| Inception-ResNet module 2 | 32 \times 98 \times 258 | 32 \times 128 \times 306 |
| Max pooling | 32 \times 98 \times 128 | 32 \times 128 \times 152 |
| Inception-ResNet module 3 | 32 \times 160 \times 128 | 32 \times 256 \times 152 |
| Inception-ResNet module 4 | 32 \times 256 \times 128 | 32 \times 380 \times 152 |
| Max pooling | 32 \times 256 \times 63 | 32 \times 380 \times 75 |
| Speed fusion (channel concatenation) | 32 \times 320 \times 63 | 32 \times 444 \times 75 |
| Temporal modeling (BiGRU ₁) | 32 \times 256 \times 63 | 32 \times 256 \times 75 |
| Uniform temporal sampling | 32 \times 256 \times 10 | 32 \times 256 \times 30 |
| Temporal modeling (BiGRU ₂) | 32 \times 196 \times 10 | 32 \times 192 \times 30 |
| Fully connected layer 1 | 32 \times 98 \times 10 | 32 \times 96 \times 30 |
| Fully connected layer 2 (output) | 32 \times 2 \times 10 | 32 \times 1 \times 30 |

BiLSTM, and BiGRU architectures. Our architecture — including a variant without auxiliary speed fusion — is benchmarked against literature baselines: VGG (Locke et al., 2020), LSTM-BiLSTM (Riahi Samani et al., 2025), Laplace-Inception (Li et al., 2022), GoogLeNet (Hajializadeh,

2023), and WaveletInception networks. Variants of LWPT, including Haar-WaveletInception and db4-WaveletInception, with and without the denoising module, are explored. Optimal hyperparameters are selected via validation, and test set performance is reported. For reproducibility, code is available on GitHub (Riahi Samani et al., 2025). The Case Study I dataset is available online (Samani, 2025); the Case Study II dataset is proprietary but available upon request from the corresponding author.

5.4. Results

Our analysis of the proposed model architecture begins with the stem module, which handles data preprocessing. This step is essential for analyzing drive-by vibration signals measured at different speeds, as varying signal lengths pose challenges for deep learning methods—particularly for CNNs, which require fixed-size inputs. To address this issue, Table 3 compares four preprocessing strategies: resizing the Continuous Wavelet Transform (CWT) and the Short-Time Fourier Transform (STFT) spectrograms to a standard size, zero-padding followed by the STFT, zero-padding followed by the LWPT stem module, and zero-padding combined with masking and length-aware components. The proposed zero-padding-based LWPT preprocessing achieves the best performance, yielding the lowest overall mean absolute percentage error (MAPE) of 5.56% in Case Study I and the highest classification accuracy of 93.29% in Case Study II. The length-aware LWPT variant ranks second, with an overall MAPE of 6.41% in Case Study I and an accuracy of 91.78% in Case Study II. In contrast, zero-padding followed

Table 3

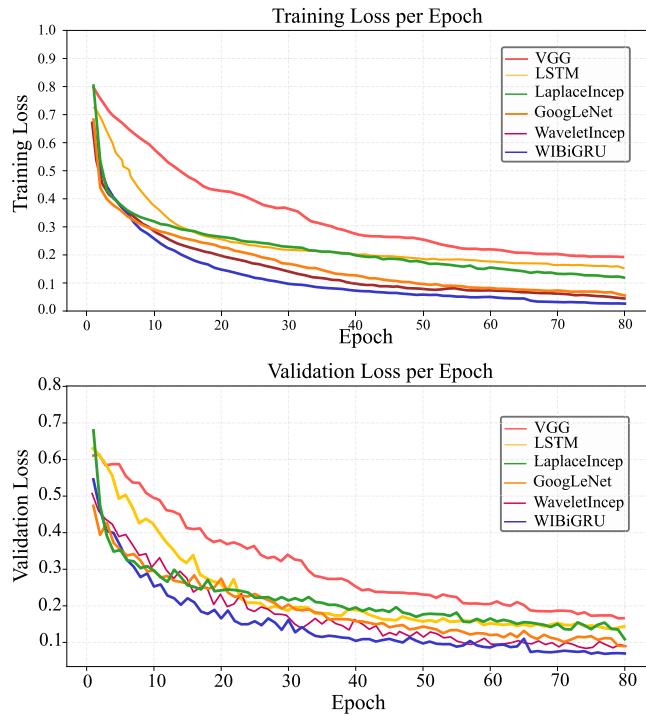
Comparison of data preprocessing strategies for handling variable-length vibration signals in the proposed framework.

| Model types | Case Study 1 (MAPE %) | | | Case Study 2 (Classification %) | | | |
|-------------------------------|-----------------------|-------|---------|---------------------------------|-------|--------|-------|
| | k_p | k_b | Overall | Prec. | F1 | Recall | Acc. |
| CWT – resizing | 8.94 | 3.42 | 6.18 | 78.82 | 85.61 | 95.19 | 87.91 |
| STFT – resizing | 9.85 | 3.62 | 6.73 | 77.28 | 83.31 | 93.73 | 85.96 |
| Padding – STFT | 9.67 | 4.07 | 6.87 | 69.75 | 74.88 | 79.97 | 84.78 |
| Padding – LWPT (Length-aware) | 9.08 | 3.75 | 6.41 | 81.05 | 87.39 | 97.43 | 91.78 |
| Padding – LWPT (ours) | 7.94 | 3.17 | 5.56 | 82.60 | 89.80 | 98.37 | 93.29 |

Table 4

Impact of the learnable signal denoising module on LWPT performance.

| Model type | Case Study 1 (MAPE %) | | | Case Study 2 (Classification %) | | | |
|------------------------|-----------------------|-------|---------|---------------------------------|-------|--------|-------|
| | k_p | k_b | Overall | Prec. | F1 | Recall | Acc. |
| Haar | 9.68 | 3.78 | 6.73 | 74.19 | 80.19 | 88.47 | 87.91 |
| db4 | 8.41 | 3.39 | 5.90 | 74.48 | 82.41 | 92.24 | 88.19 |
| Haar + Denoising | 8.89 | 3.43 | 6.16 | 76.78 | 83.66 | 91.89 | 89.23 |
| db4 + Denoising (ours) | 7.94 | 3.17 | 5.56 | 82.60 | 89.80 | 98.37 | 93.29 |

**Fig. 9.** Training and validation Loss convergence for the proposed model and baselines.

by STFT exhibits the weakest performance, with the highest overall MAPE of 6.87 in Case Study I and the lowest accuracy of 84.78% in Case Study II.

Table 4 analyzes the performance of different LWPT configurations with and without the learnable denoising module. The proposed db4-based LWPT with denoising achieves the best performance, yielding the lowest overall MAPE of 5.56% in Case Study I and the highest accuracy of 93.23% in Case Study II. Removing the denoising module degrades performance, with the db4 wavelet reaching an overall MAPE of 5.90% and an accuracy of 88.19%. Haar-based configurations show consistently higher estimation errors and lower classification accuracy, with the non-denoising Haar variant exhibiting an overall MAPE of 6.73% and an accuracy of 87.91%.

Table 5 compares the proposed framework with existing baseline models across both case studies. For a fair comparison with baseline feature extractors, the WaveletInception–BiGRU model without speed

fusion is considered. This model achieves the best overall performance among all compared methods, with an overall MAPE of 6.09% in Case Study I and an accuracy of 91.53% in Case Study II. The WaveletInception model, corresponding to our variant without recurrent temporal modeling, follows with an overall MAPE of 6.46% and an accuracy of 90.32%. Among other deep learning baselines, GoogLeNet (Hajializadeh, 2023) exhibits the strongest performance, achieving an overall MAPE of 7.47% and an accuracy of 88.93%, followed by the Laplace–Inception model (Lamprea-Pineda et al., 2024) with an overall MAPE of 8.72% and an accuracy of 86.49%. LSTM-based (Riahi Samani et al., 2025) and VGG (Locke et al., 2020) models rank lower, respectively, with overall MAPEs of 10.56% and 13.08% and accuracies of 81.03% and 78.77%, respectively.

Table 6 presents the computational efficiency of the proposed framework compared with baseline models. The WaveletInception–BiGRU model achieves a balanced trade-off between model size and inference speed. It requires only 1.01 M and 1.20 M parameters for Case Study I and Case Study II, respectively, making it substantially lighter than GoogLeNet (Hajializadeh, 2023), which has a comparable accuracy. In terms of inference time per sample, the WaveletInception–BiGRU model ranks in the middle, slightly slower than the Wavelet–Inception and Laplace–Inception (Li et al., 2022) models. However, WaveletInception–BiGRU, with 0.99 and 3.14 ms, respectively, in Case Study I and II, remains significantly faster than GoogLeNet (Hajializadeh, 2023) and LSTM Network (Riahi Samani et al., 2025) models.

To incorporate operational conditions such as measurement speed, two mid-level fusion strategies, early and late, are evaluated. In early fusion, the encoded speed tensor is concatenated with vibration signal features before processing by the 1D Inception–ResNet modules for deep feature extraction. In late fusion, this concatenation occurs after high-level vibration features are extracted, as designed in our proposed architecture. **Table 7** compares six fusion designs, including CNN-, BiLSTM-, and BiGRU-based speed fusion modules. The proposed late BiGRU-based fusion achieves the best overall performance, with the lowest overall MAPE of 5.56% in Case Study I and the highest accuracy of 93.29% in Case Study II. The second-best configuration, achieving an overall MAPE of 5.81% and an accuracy of 91.52%. These results indicate that late fusion consistently outperforms early fusion, while BiGRU-based designs provide additional performance gains. In contrast, CNN-based fusion approaches exhibit the lowest performance across both fusion stages, with overall MAPE exceeding 6.49% and accuracy below 85.41%, highlighting their limited ability to capture temporal dependencies and fuse auxiliary speed information.

We further evaluate the performance of LSTM- and GRU-based modules for capturing temporal dependencies in our model architecture. **Table 8** compares LSTM, GRU, BiLSTM, and BiGRU networks for

Table 5

Performance comparison of the proposed framework against state-of-the-art baseline models.

| Model types | Case Study 1 (MAPE %) | | | Case Study 2 (Classification %) | | | |
|-----------------------------------------------------|-----------------------|-------|---------|---------------------------------|-------|--------|-------|
| | k_p | k_b | Overall | Prec. | F1 | Recall | Acc. |
| VGG (Locke et al., 2020) | 18.88 | 7.28 | 13.08 | 62.82 | 69.27 | 78.19 | 78.77 |
| LSTM Network (Riahi Samani et al., 2025) | 15.10 | 6.02 | 10.56 | 67.19 | 72.06 | 80.29 | 81.03 |
| Laplace-Inception (Li et al., 2022) | 12.33 | 5.11 | 8.72 | 72.38 | 79.79 | 88.88 | 86.49 |
| GoogLeNet (Hajializadeh, 2023) | 11.17 | 3.76 | 7.47 | 77.82 | 82.67 | 92.17 | 88.93 |
| WaveletInception | 9.07 | 3.84 | 6.46 | 76.89 | 86.93 | 97.41 | 90.32 |
| WaveletInception-BiGRU (ours, without speed fusion) | 8.64 | 3.54 | 6.09 | 81.22 | 88.56 | 98.02 | 91.53 |

Table 6

Comparison of computational efficiency and model size between the proposed framework and baseline models from the literature.

| Model types | Case Study 1 | | Case Study 2 | |
|-----------------------------------------------------|--------------|----------------|--------------|----------------|
| | Time (ms) | Model size (M) | Time (ms) | Model size (M) |
| VGG (Locke et al., 2020) | 3.09 | 2.23 | 4.69 | 2.69 |
| LSTM Network (Riahi Samani et al., 2025) | 6.69 | 2.88 | 7.70 | 3.32 |
| Laplace-Inception (Li et al., 2022) | 2.16 | 2.1 | 2.46 | 2.29 |
| GoogLeNet (Hajializadeh, 2023) | 2.93 | 6.45 | 8.27 | 5.61 |
| WaveletInception | 0.93 | 1.25 | 2.24 | 1.65 |
| WaveletInception-BiGRU (ours, without speed fusion) | 0.99 | 1.01 | 3.14 | 1.2 |

Table 7

Ablation study of six fusion modules for incorporating measurement speed as auxiliary information in the proposed WaveletInception-BiGRU model.

| Model types | Case Study 1 (MAPE %) | | | Case Study 2 (Classification) | | | |
|-------------------------|-----------------------|-------|---------|-------------------------------|-------|--------|-------|
| | k_p | k_b | Overall | Prec. | F1 | Recall | Acc. |
| Early CNN-based | 9.63 | 3.86 | 6.74 | 70.86 | 76.02 | 82.57 | 83.84 |
| Early BiLSTM-based | 8.91 | 3.69 | 6.30 | 76.03 | 82.87 | 93.44 | 87.67 |
| Early BiGRU-based | 8.67 | 3.63 | 6.15 | 77.55 | 85.29 | 95.69 | 89.18 |
| Late CNN-based | 9.24 | 3.74 | 6.49 | 71.67 | 77.75 | 84.95 | 85.41 |
| Late BiLSTM-based | 8.28 | 3.35 | 5.81 | 78.96 | 87.61 | 98.09 | 91.52 |
| Late BiGRU-based (ours) | 7.94 | 3.17 | 5.56 | 82.60 | 89.80 | 98.37 | 93.29 |

Table 8

Ablation study comparing unidirectional and bidirectional GRU and LSTM networks for temporal feature extraction in the proposed WaveletInception-BiGRU model.

| Model types | Case Study 1 (MAPE %) | | | | Case Study 2 (Classification %) | | | |
|--------------|-----------------------|-------|---------|----------|---------------------------------|-------|--------|-------|
| | k_p | k_b | Overall | Size (M) | Prec. | F1 | Recall | Acc. |
| LSTM | 9.38 | 4.08 | 6.73 | 1.11 | 68.87 | 76.84 | 90.29 | 85.61 |
| BiLSTM | 8.28 | 3.35 | 5.81 | 1.22 | 78.96 | 87.61 | 98.09 | 91.52 |
| GRU | 8.97 | 3.66 | 6.31 | 1.02 | 73.13 | 82.34 | 94.21 | 87.88 |
| BiGRU (ours) | 7.94 | 3.17 | 5.56 | 1.03 | 82.60 | 89.80 | 98.37 | 93.29 |

temporal feature extraction. The BiGRU model achieves the best overall results, with the lowest MAPE of 5.56% in Case Study I and the highest accuracy of 93.29% in Case Study II. BiLSTM follows, with an overall MAPE of 5.81% and an accuracy of 91.52%, outperforming its unidirectional counterpart. Among unidirectional models, GRU outperforms LSTM, with a MAPE of 6.31% versus 6.73% and an accuracy of 87.88% versus 85.61%. Bidirectional variants introduce only a modest increase in the number of parameters while consistently improving performance over unidirectional architectures.

We evaluate the proposed WI-BiGRU model for estimating track stiffness variations under four scenarios in Case Study I. **Table 9** summarizes the RMSE and MAPE for both k_p and k_b . The WI-BiGRU model achieves the best results in the uniform stiffness scenario, with MAPEs of 5.11 for k_p and 1.86% for k_b . The most challenging case is the single-sleeper stiffness drop, which results in the highest errors (overall MAPE of 7.81%). The three-sleeper drop scenario yields moderate performance (overall MAPE of 6.19%), while transition-zone changes are estimated more accurately (overall MAPE of 4.56%). Across all

scenarios, the model finds k_p harder to estimate than k_b , with overall MAPEs of 7.94% and 3.17%, respectively.

Fig. 10 shows examples of railpad and ballast stiffness estimations compared to their ground truth values. The blue lines represent the ground truth, while the orange lines represent the model's estimations. These examples illustrate the model's capability to estimate both parameters under various scenarios, including uniform stiffness, localized reductions, and transition zones.

Furthermore, we evaluate the WI-BiGRU performance for the transition zone identification in Case Study II. **Table 10** summarizes the model's performance in identifying sleepers on the normal track and above the viaduct at Site 4 in Case Study II. The model achieves an overall accuracy of 93.29% in both scenarios, demonstrating high and balanced classification performance. For the viaduct stiffness scenario, the model achieves 82.60% precision, 98.37% recall, and 89.80% F1-score, indicating strong detection of transition-zone sleepers with minimal false negatives. In the normal track scenario, the model achieves a precision of 99.24%, a recall of 91.12%, and an F1-score of 95.00%, reflecting reliable identification with few false positives.

Table 9
WI-BiGRU estimation accuracy performance across four track stiffness scenarios (Case Study I).

| Case Study I Scenarios | k_p | | k_b | | Overall MAPE (%) |
|------------------------------|-------------|----------|-------------|----------|------------------|
| | RMSE (MN/m) | MAPE (%) | RMSE (MN/m) | MAPE (%) | |
| (I) Uniform | 15.29 | 5.11 | 0.65 | 1.86 | 3.36 |
| (II) Drops in 1 sleeper | 21.99 | 10.76 | 1.81 | 4.66 | 7.81 |
| (III) Drops in 3 sleepers | 19.34 | 8.98 | 1.22 | 3.41 | 6.19 |
| (IV) Transition zone changes | 17.01 | 6.62 | 0.86 | 2.54 | 4.56 |
| All scenarios | 18.21 | 7.94 | 1.13 | 3.17 | 5.56 |

Table 10
WaveletInception-BiGRU model estimation accuracy performance on the track transition zone identification (Case Study II).

| Scenarios | Prec. (%) | Recall (%) | F1 (%) | Accuracy (%) |
|----------------------------|-----------|------------|--------|--------------|
| (I) Normal Track Stiffness | 99.24 | 91.12 | 95.00 | 93.29 |
| (II) Viaduct Stiffness | 82.60 | 98.37 | 89.80 | 93.29 |

6. Discussion

This paper presents two case studies demonstrating the effectiveness of the WaveletInception-BiGRU deep learning framework for on-board vibration response analysis. Our model outperforms four external baseline models and its own ablations in both track stiffness estimation and railway transition zone identification. In the first case study, the WaveletInception-BiGRU model achieves a MAPE of 7.94% for railpad stiffness and 3.17% for ballast stiffness across different scenarios of track stiffness reduction (see Table 9). In the second case study, it yields an accuracy of 93.29% for transition zone identification (see Table 10). Compared to baseline models in the literature, our framework demonstrates notable accuracy (see Table 5) and computational efficiency (see Table 6), supported by insights from our ablation studies (see Tables 3–4 and Tables 7–8).

At the core of the proposed architecture is the WaveletInception network, which integrates an LWPT stem with 1D Inception-ResNet blocks. The LWPT stem decomposes input vibration signals into multi-resolution components while preserving all original signal information due to its orthogonality. This decomposition provides physically informative features in the early stages of the network and enables downsampling without information loss.

The results in Table 3 demonstrate that the proposed LWPT-based stem module significantly outperforms conventional preprocessing techniques. Standard resizing techniques, such as those used for CWT and STFT, rely on interpolation and averaging, which introduce non-physical values and spectro-temporal decimation. This resolution loss is particularly detrimental in rail health monitoring, where accurately localized, transient signatures — such as rail squats or fastening failures — are obscured by the smoothing effects of frequency-bin averaging and reduced temporal resolution.

Furthermore, the LWPT stem addresses the critical challenge of handling variable signal lengths via zero padding. While essential for CNN-based deep learning architectures, padding often introduces artificial discontinuities and spectral leakage, leading to the degraded performance observed in STFT-based preprocessing (Accuracy: 84.78%; MAPE: 6.87%). In contrast, the proposed LWPT stem mitigates these artifacts through multi-resolution analysis, adaptively resolving frequency components without the constraints of rigid windowing. Specifically, utilizing $db4$ filters over the Haar variant proves superior; the higher-order $db4$ wavelets provide smoother boundary transitions, effectively distributing padding-induced distortions and allowing the model to isolate the low-frequency features essential for accurate defect mapping.

Empirical evidence (Table 4) also highlights the effectiveness of the employed learnable denoising function (Michau et al., 2022). Given that stochastic noise is inherent in real-world vibration data, these results suggest that integrating adaptive thresholding directly into the

architectural stem is more effective than static preprocessing (no denoising).

Since not all frequency decomposition bins of the LWPT are equally relevant for the subsequent deeper layer, our analysis showed the significance of a channel projection layer to adapt the LWPT for a supervised deep learning architecture effectively. This can boost the architecture to prioritize frequency decompositions, which are more informative for the task-specific monitoring. However, a primary limitation of the current LWPT implementation is its sequential recursive structure, which precludes parallel computation across nodes. As the decomposition depth L increases, the exponential growth of the number of nodes (2^L) can significantly increase computational overhead. Furthermore, a physical interpretation of our stem module, particularly the LWPT and its channel projection layer, remains an open question for future research.

The WaveletInception network offers both representational and computational benefits for vibration-based feature extraction. It integrates spectral information directly into the deep learning architecture. The 1D Inception-Residual network has multi-scale parallel branches that capture informative patterns across temporal scales, essential for track condition monitoring (Shen et al., 2023; Lamprea-Pineda et al., 2024; Shen et al., 2021). The Inception-ResNet blocks use filters of varying sizes to extract multi-scale, high-level features from the outputs of the stem module.

Moreover, the residual connections support stable learning and an efficient gradient flow, helping deeper models converge and avoid degradation. Bottleneck convolutions further reduce parameter count and computational cost while preserving discriminative power.

The proposed 1D convolutional network avoids the computational overhead of 2D convolutions over time-frequency representations (see Table 6). This design not only reduces processing time but also can preserve the physical interpretability of the extracted features, as they remain closely linked to the original signal. While 2D networks (Hajializadeh, 2023) can capture complex time-frequency interactions, their features are often more abstract and computationally demanding. Indeed, they convolve over frequency bins, which may not be physically interpretable. By focusing on 1D convolutions, the WaveletInception-BiGRU model can achieve an efficient balance between performance, interpretability, and computational cost.

The proposed framework leverages Bi-directional Gated Recurrent Unit (BiGRU) layers to model the long-range temporal dependencies inherent in on-board vibration responses. As shown in Table 8 and motivated by Figs. 6–8, the inclusion of bidirectional context is critical for capturing interrelated dependencies between measurement speed, excitation levels, and modal characteristics. By processing the signal in both forward and backward paths, the architecture captures the neighboring context of vibrations, which is essential for the accurate health-condition estimation of individual infrastructure components.

The proposed BiGRU-based architecture captures the relationships between onboard vibrations and operational-condition embeddings. As shown in Table 7, the late-fusion BiGRU strategy significantly outperforms CNN and LSTM counterparts, particularly in the real-world scenarios of Case Study II (93.29% accuracy). This advantage arises from the BiGRU's ability to effectively model interactions between speed embeddings and spectral features extracted by the WaveletInception network. By capturing bidirectional temporal dependencies, the

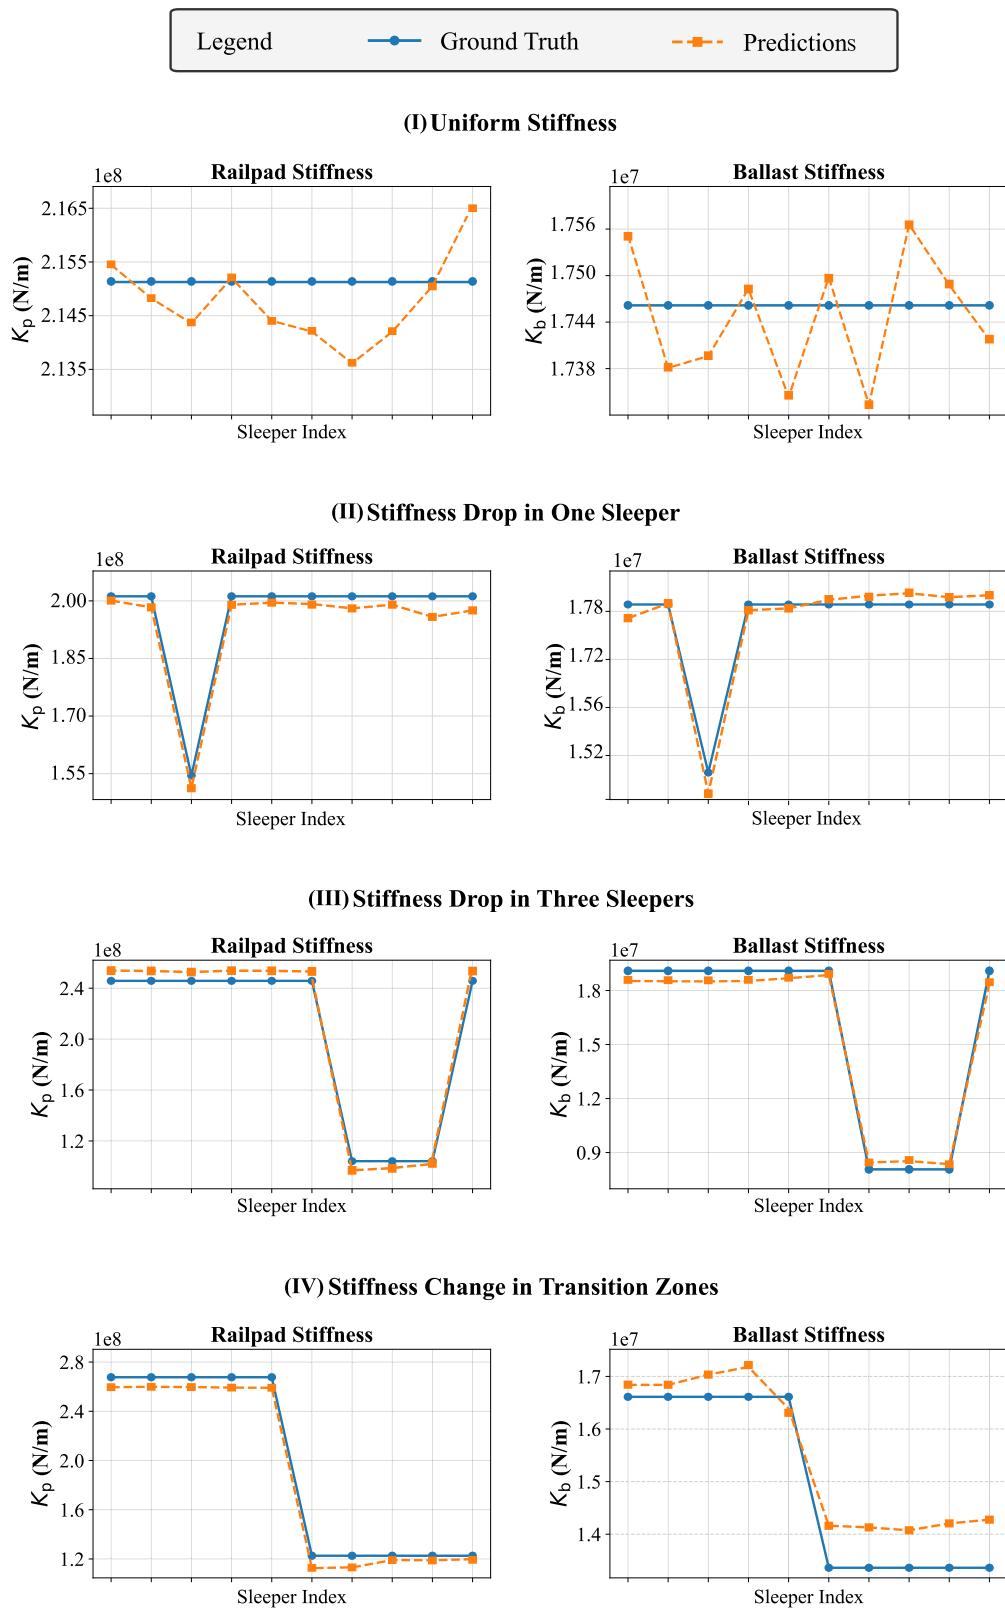


Fig. 10. WI-BiGRU model predictions on the Case Study I test set. Ground-truth values are shown in blue, and model estimates in orange.

network adaptively manages variable signal lengths, mitigates zero-padding artifacts, and adjusts the extracted vibration features under operational conditions, ensuring consistent health-condition estimation. Furthermore, the high precision observed in Case Study II, which

includes environmental noise and sensor interference typical of real-world conditions, demonstrates the framework's practical robustness. While this paper confirms the model's reliability across diverse operational conditions, future work should systematically stress-test the

model under extreme signal-to-noise ratios to further assess the noise robustness of our proposed model architecture.

A distinctive feature of our architecture is the distributed temporal-sampling strategy that generates N feature vectors for the corresponding N structural beams. Recurrent models, such as LSTM and GRU modules, often employ a many-to-one architecture, collapsing the entire sequence into a single terminal hidden state (the last time step). This can create an information bottleneck, particularly for high-resolution vibration signals, in which local transient features are lost during global compression. In contrast, our strategy performs distributed selection across the temporal dimension, maintaining a direct mapping between features and the physical, spatial distribution of the target component within the infrastructure.

This hierarchical sequence modeling approach offers significant practical advantages for on-board monitoring. By selecting N equidistant points in the temporal domain, the model effectively aligns with the physical layout of the infrastructure without requiring high-precision spatial positioning. This is particularly advantageous when measurement speed varies; while precise GPS-based positioning becomes challenging under fluctuating speeds, the number of structural targets within a segment remains constant. Consequently, the distributed sampling preserves temporally localized features and global context, enabling finer localized health assessments while mitigating the representation loss inherent in standard length-aware or packing-based sequence modules (see Table 3).

By processing vibration responses directly in the time domain, the proposed architecture eliminates the necessity for manual feature engineering. Unlike traditional diagnostic frameworks that rely on hand-crafted statistical or spectral indicators, the proposed stem module enables the network to automatically derive sensitive features directly from vibration signals collected at varying speeds. This process is further enhanced by the BiGRU-based integration of operational data, which facilitates the internal modeling of vibration signatures across varying speed profiles. Consequently, this end-to-end approach establishes a fully automated monitoring system that maintains high diagnostic precision despite fluctuations in measurement speed.

The WI-BiGRU model demonstrates high precision across all scenarios in Case Study I (Table 9), though performance varies with the localization of defects. While the overall MAPE is 5.56%, the highest misestimation occurs in Scenario (II) (single sleeper drop), where the k_p error reaches 10.76%. This indicates that highly localized, transient stiffness changes are more difficult to isolate than the sustained signatures found in 3-sleeper drops or transition zones. Furthermore, rail pad stiffness (k_p) exhibits higher error rates than ballast stiffness (k_b), likely because k_b has a more dominant influence on the global track vibration response. In Case Study II (Table 10), the model accurately identifies transition zones but reveals a lower precision for the viaduct stiffness class (82.60%). This suggests that some normal sleepers near the transition boundary are misclassified as viaduct sleepers due to mechanical coupling and boundary blurring across different track structures. These findings indicate that, although the model is highly effective for mapping structural conditions, its sensitivity can still be improved for single-point defects and boundary regions.

The dataset in Case Study I is based on a vehicle-track interaction model from Shen et al. (2023) to demonstrate our deep learning framework. Simulated drive-by responses may not capture all real-world environmental and operational factors. Phenomena like weld effects, temperature changes in rail stiffness, and rail neutral temperature are not modeled. The simulation uses a single vehicle-track setup, while real railways have varying profiles, welds, fastenings, sleepers, and structures that could affect results (Shen et al., 2023). Case Study II shows that our framework can identify transition zones in real-world conditions. Transition zones may exhibit additional degradation mechanisms, such as ballast fouling and differential settlement, suggesting that expanding the model to include additional health indicators, such as track geometry degradation and differential settlement, is essential

for future research. Additionally, the framework can also be applied to various track types, including ballast and ballastless tracks, and broader transportation infrastructure, such as railway and road bridges and tunnels.

7. Conclusions and future work

In this paper, we have proposed a novel deep learning framework, Wavelet-Inception-BiGRU, for estimating infrastructure health condition monitoring using onboard vibration signals. The core of our methodology is a new feature-extraction network, called WaveletInception, that integrates a Learnable Wavelet Packet Transform (LWPT) and 1D Inception-ResNet blocks to effectively extract multi-scale, high-level representations from vibration signals. Furthermore, the feature extraction network incorporates operational conditions, in particular measurement speed, through feature-level fusion with a bidirectional Gated Recurrent Unit (BiGRU) layer. This enables the model to learn interrelated features across different measurement conditions. For the health condition estimation head, we have proposed a sequential model using BiGRU networks. The model leverages temporal information from both the forward and backward directions of on-board measurements. Additionally, the sequential modeling approach enables more detailed estimation at the beam or component level, enabling comprehensive assessments of different segments of infrastructure.

Future research will focus on three main directions. First, this paper is one of the first to include operational conditions of on-board vibration response in a deep learning model. The BiGRU layer captures temporal dependencies and fuses vibration and operational data. In the future, other temporal encoders, like attention mechanisms, can be explored to improve feature extraction. Second, our results show that the proposed model works well for health indicators with low-frequency features. Next, the framework can be adapted for high-frequency applications, such as defect detection in superstructures and rails, using advanced sensors like laser Doppler vibrometers. Third, to support large-scale, automated, cloud-based monitoring, future research will systematically test the framework's stability when signal-to-noise ratios and operational conditions vary.

CRediT authorship contribution statement

Reza Riahi Samani: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Alfredo Núñez:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Data curation. **Bart De Schutter:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The first author thanks Hesam Araghi for insightful comments and Chen Shen for providing the simulator for Case Study I. This research was partly supported by ProRail and Europe's Rail Flagship Project IAM4RAIL - Holistic and Integrated Asset Management for Europe's Rail System [grant agreement 101101966].

Disclaimer

Funded by the European Union. Views and opinion expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or Europe's Rail Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them. The project FP3-IAM4Rail is supported by the Europe's Rail Joint Undertaking and its members.

Data availability

Data will be made available on request.

References

Aggarwal, C.C., et al., 2018. Neural Networks and Deep Learning, vol. 10, (978), Springer, <http://dx.doi.org/10.1007/978-3-031-29642-0>.

Bergstra, J., Bengio, Y., 2012. Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* 13 (1), 281–305.

Biewald, L., 2020. Experiment tracking with weights and biases. Software available from wandb.com, URL <https://www.wandb.com/>.

Bishop, C.M., Bishop, H., 2023. Deep Learning: Foundations and Concepts. Springer Nature, <http://dx.doi.org/10.1007/978-3-031-45468-4>.

Borghesani, P., Herwig, N., Antoni, J., Wang, W., 2023. A Fourier-based explanation of 1D-CNNs for machine condition monitoring applications. *Mech. Syst. Signal Process.* 205, 110865. <http://dx.doi.org/10.1016/j.ymssp.2023.110865>.

bouny, L.E., Khalil, M., Adib, A., 2020. An End-to-End Multi-Level Wavelet Convolutional Neural Networks for heart diseases diagnosis. *Neurocomputing* 417, 187–201. <http://dx.doi.org/10.1016/j.neucom.2020.07.056>.

Caprioli, A., Cigada, A., Ravaglia, D., 2007. Rail inspection in track maintenance: A benchmark between the wavelet approach and the more conventional Fourier analysis. *Mech. Syst. Signal Process.* 21, 631–652. <http://dx.doi.org/10.1016/j.ymssp.2005.12.001>.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y., 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y., 2021. Data-driven structural health monitoring using feature fusion and hybrid deep learning. *IEEE Trans. Autom. Sci. Eng.* 18 (4), 2087–2103. <http://dx.doi.org/10.1109/TASE.2020.3034401>.

Chung, J., Gulcehre, C., Cho, K., Bengio, Y., 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. <http://dx.doi.org/10.48550/arXiv.1412.3555>.

Corbally, R., Malekjafarian, A., 2024. A deep-learning framework for classifying the type, location, and severity of bridge damage using drive-by measurements. *Computer-Aided Civ. Infrastruct. Eng.* 39, 852–871. <http://dx.doi.org/10.1111/mice.13104>.

Dang, H.V., Tran-Ngoc, H., Nguyen, T.V., Bui-Tien, T., De Roeck, G., Nguyen, H.X., 2021. Data-driven structural health monitoring using feature fusion and hybrid deep learning. *IEEE Trans. Autom. Sci. Eng.* 18, 2087–2103. <http://dx.doi.org/10.1109/TASE.2020.3034401>.

European Commission, 2023. Eighth monitoring report on the development of the rail market under Article 15(4) of Directive 2012/34/EU of the European Parliament and of the Council.

Fernández-Bobadilla, H.A., Martin, U., 2023. Modern tendencies in vehicle-based condition monitoring of the railway track. *IEEE Trans. Instrum. Meas.* 72, 1–44. <http://dx.doi.org/10.1109/TIM.2023.3243673>.

Frusque, G., Fink, O., 2024. Robust time series denoising with learnable wavelet packet transform. *Adv. Eng. Informatics* 62, 102669. <http://dx.doi.org/10.1016/j.aei.2024.102669>.

González, A., Feng, K., Casero, M., 2023. Effective separation of vehicle, road and bridge information from drive-by acceleration data via the power spectral density resulting from crossings at various speeds. *Dev. Built Environ.* 14, 100162. <http://dx.doi.org/10.1016/j.dibe.2023.100162>.

Hajializadeh, D., 2023. Deep learning-based indirect bridge damage identification system. *Struct. Health Monit.* 22, 897–912. <http://dx.doi.org/10.1177/14759217221087147>.

Hoelzl, C., Dertimanis, V., Ancu, L., Kollros, A., Chatzi, E., 2023. Vold–Kalman filter order tracking of axle box accelerations for track stiffness assessment. *Mech. Syst. Signal Process.* 204, 110817. <http://dx.doi.org/10.1016/j.ymssp.2023.110817>.

Huang, J., Yin, X., Kaewunruen, S., 2022. Quantification of dynamic track stiffness using machine learning. *IEEE Access* 10, 78747–78753. <http://dx.doi.org/10.1109/ACCESS.2022.3191278>.

Khodadoost, S., Nouhi, B., Talatohari, S., Sareh, P., 2025. Intelligent vibration-based structural health monitoring systems: Methodological advances and challenges. *Sensors Actuators A: Phys.* 394, 116886. <http://dx.doi.org/10.1016/j.sna.2025.116886>.

Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization. <http://dx.doi.org/10.48550/arXiv.1412.6980>, arXiv:1412.6980.

Lahat, D., Adali, T., Jutten, C., 2015. Multimodal data fusion: an overview of methods, challenges, and prospects. *Proc. IEEE* 103 (9), 1449–1477.

Lamprea-Pineda, A.C., Connolly, D.P., Castanheira-Pinto, A., Alves-Costa, P., Hussein, M.F., Woodward, P.K., 2024. On railway track receptance. *Soil Dyn. Earthq. Eng.* 177, 10831. <http://dx.doi.org/10.1016/j.soildyn.2023.10831>.

Lederman, G., Chen, S., Garrett, J., Kovačević, J., Noh, H.Y., Bielak, J., 2017. Track-monitoring from the dynamic response of an operational train. *Mech. Syst. Signal Process.* 87, 1–16. <http://dx.doi.org/10.1016/j.ymssp.2016.06.041>.

Li, T., Zhao, Z., Sun, C., Cheng, L., Chen, X., Yan, R., Gao, R.X., 2022. WaveletKernelNet: An interpretable deep neural network for industrial intelligent diagnosis. *IEEE Trans. Syst. Man, Cybern.: Syst.* 52, 2302–2312. <http://dx.doi.org/10.1109/TSMC.2020.3048950>.

Liu, P., Zhang, H., Lian, W., Zuo, W., 2019. Multi-level wavelet convolutional neural networks. *IEEE Access* 7, 74973–74985.

Locke, W., Sybrandy, J., Redmond, L., Safro, I., Atamturktur, S., 2020. Using drive-by health monitoring to detect bridge damage considering environmental and operational effects. *J. Sound Vib.* 468, 115088. <http://dx.doi.org/10.1016/j.jsv.2019.115088>.

Malekjafarian, A., Golpayegani, F., Moloney, C., Clarke, S., 2019b. A machine learning approach to bridge-damage detection using responses measured on a passing vehicle. *Sensors* 19, 4035.

Malekjafarian, A., OBrien, E., Quirke, P., Bowe, C., 2019a. Railway track monitoring using train measurements: An experimental case study. *Appl. Sci.* 9, 4859. <http://dx.doi.org/10.3390/app9224859>.

Mallat, S., 1999. A Wavelet Tour of Signal Processing. Elsevier, <http://dx.doi.org/10.1016/B978-0-12-374370-1.X0001-8>.

Mao, X., Xia, C., Liu, J., Zhang, H., Ding, Y., Yao, Y., Liu, Z., 2025. A novel similarity measure based on dispersion-transition matrix and Jensen–Fisher divergence and its application on the detection of rail short-wave defects. *Chaos Solitons Fractals* 192, 115988. <http://dx.doi.org/10.1016/j.chaos.2025.115988>.

Michau, G., Frusque, G., Fink, O., 2022. Fully learnable deep wavelet transform for unsupervised monitoring of high frequency time series. *Proc. Natl. Acad. Sci.* 119 (8), e2106598119.

Molodova, M., Li, Z., Núñez, A., Dollevoet, R., 2014. Automatic detection of squats in railway infrastructure. *IEEE Trans. Intell. Transp. Syst.* 15 (5), 1980–1990.

Mou, L., Zhou, C., Zhao, P., Nakisa, B., Rastgoo, M.N., Jain, R., Gao, W., 2021. Driver stress detection via multimodal fusion using attention-based CNN-LSTM. *Expert Syst. Appl.* 173, 114693. <http://dx.doi.org/10.1016/j.eswa.2021.114693>.

Peng, Z., Li, J., Zhong, Y., Hao, H., 2025. Interpretable damage sensitive feature extraction for drive-by structural health monitoring using knowledge distillation-based deep learning. *J. Sound Vib.* 618, 119303. <http://dx.doi.org/10.1016/j.jsv.2025.119303>.

Phusakulkajorn, W., Núñez, A., Wang, H., Jamshidi, A., Zoeteman, A., Ripke, B., Dollevoet, R., De Schutter, B., Li, Z., 2023. Artificial intelligence in railway infrastructure: Current research, challenges, and future opportunities. *Intell. Transp. Infrastruct.* 2, liad016. <http://dx.doi.org/10.1093/iti/liad016>.

Phusakulkajorn, W., Unsiwilai, S., Chang, L., Li, Z., Núñez, A., 2025a. A hybrid neural model approach for health assessment of railway transition zones with multiple data sources. *IEEE Trans. Instrum. Meas.* 74, 1–17. <http://dx.doi.org/10.1109/TIM.2025.3575986>.

Phusakulkajorn, W., Unsiwilai, S., Chang, L., Li, Z., Núñez, A., 2025b. A hybrid neural model approach for health assessment of railway transition zones with multiple data sources. *IEEE Trans. Instrum. Meas.* 74, 1–17. <http://dx.doi.org/10.1109/TIM.2025.3575986>.

Phusakulkajorn, W., Unsiwilai, S., Chang, L., Li, Z., Núñez, A., 2025d. A hybrid neural model approach for health assessment of railway transition zones with multiple data sources. *IEEE Trans. Instrum. Meas.*

Phusakulkajorn, W., Zeng, Y., Li, Z., Núñez, A., 2025c. Unsupervised representation learning for monitoring rail infrastructures with high-frequency moving vibration sensors. *IEEE Trans. Intell. Transp. Syst.* 1–15. <http://dx.doi.org/10.1109/TITS.2025.3557712>.

Quirke, P., Cantero, D., OBrien, E.J., Bowe, C., 2017. Drive-by detection of railway track stiffness variation using in-service vehicles. *Proc. Inst. Mech. Eng. Part F: J. Rail Rapid Transit* 231, 498–514. <http://dx.doi.org/10.1177/0954409716634752>.

Riahi Samani, R., Nunez Vicencio, A., De Schutter, B., 2025. Bidirectional long short-term memory approach for infrastructure health monitoring using on-board vibration response. *Transp. Res.* 2679 (10), 344–356. <http://dx.doi.org/10.1177/03611981251342236>.

Riahi Samani, R., Nunez Vicencio, A., De Schutter, B., 2025. WaveletInception-BiGRU: Deep learning framework for on-board infrastructure health monitoring. GitHub repository, URL <https://github.com/rezars9701/WaveletInception-BiGRU>.

Samani, R.R., 2025. Data underlying the publication: WaveletInception networks for drive-by-vibration-based infrastructure health monitoring, 4TU.ResearchData. <http://dx.doi.org/10.4121/8d245b5-2d81-436f-896c-83a78e22d213>.

Samani, R.R., Nunez, A., Schutter, B.D., 2024. A bidirectional long short term memory approach for infrastructure health monitoring using on-board vibration response. <http://dx.doi.org/10.48550/arXiv.2412.02643>, arXiv:2412.02643.

Sansinena, A., Rodriguez-Arana, B., Arrizabalaga, S., 2025. A systematic review of acceleration-based estimation of railway track quality. *Veh. Syst. Dyn.* 1–28. <http://dx.doi.org/10.1080/00423114.2025.2483972>.

Shen, C., Dollevoet, R., Li, Z., 2021. Fast and robust identification of railway track stiffness from simple field measurement. *Mech. Syst. Signal Process.* 152, 107431. <http://dx.doi.org/10.1016/j.ymssp.2020.107431>.

Shen, C., Zhang, P., Dollevoet, R., Zoeteman, A., Li, Z., 2023. Evaluating railway track stiffness using axle box accelerations: A digital twin approach. *Mech. Syst. Signal Process.* 204, 110730. <http://dx.doi.org/10.1016/j.ymssp.2023.110730>.

Shi, C., Zhou, Y., Xu, L., Zhang, X., Guo, Y., 2023. A critical review on the vertical stiffness irregularity of railway ballasted track. *Constr. Build. Mater.* 400, 132715. <http://dx.doi.org/10.1016/j.conbuildmat.2023.132715>.

Unsiwilai, S., Phusakulkajorn, W., Shen, C., Zoeteman, A., Dollevoet, R., Núñez, A., Li, Z., 2024. Enhanced vertical railway track quality index with dynamic responses from moving trains. *Heliyon* 10, e38670. <http://dx.doi.org/10.1016/j.heliyon.2024.e38670>.

Unsiwilai, S., Wang, L., Núñez, A., Li, Z., 2023. Multiple-axle box acceleration measurements at railway transition zones. *Measurement* 213, 112688. <http://dx.doi.org/10.1016/j.measurement.2023.112688>.

Wang, H., Markine, V.L., 2018. Methodology for the comprehensive analysis of railway transition zones. *Comput. Geotech.* 99, 64–79. <http://dx.doi.org/10.1016/j.compgeo.2018.03.001>.

Wang, H., Markine, V., Liu, X., 2018b. Experimental analysis of railway track settlement in transition zones. *Proc. Inst. Mech. Eng. Part F: J. Rail Rapid Transit* 232 (6), 1774–1789. <http://dx.doi.org/10.1177/0954409717748789>.

Wang, J., Wang, Z., Li, J., Wu, J., 2018a. Multilevel wavelet decomposition network for interpretable time series analysis. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. In: KDD '18, pp. 2437–2446. <http://dx.doi.org/10.1145/3219819.3220060>.

Yang, Y.B., Yang, J.P., 2018. State-of-the-art review on modal identification and damage detection of bridges by moving test vehicles. *Int. J. Struct. Stab. Dyn.* 18, 1850025. <http://dx.doi.org/10.1142/S0219455418500256>.

Yao, T., Pan, Y., Li, Y., Ngo, C.-W., Mei, T., 2022. Wave-ViT: Unifying wavelet and transformers for visual representation learning. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (Eds.), *Computer Vision – ECCV 2022*. pp. 328–345. http://dx.doi.org/10.1007/978-3-031-19806-9_19.

Zhang, Y., Li, H., Wang, H., 2025. Data-driven wind-induced response prediction for slender civil infrastructure: Progress, challenges and opportunities. *Structures* 74, 108650. <http://dx.doi.org/10.1016/j.istruc.2025.108650>.

Zhang, Y., Marie d'Avigneau, A., Hadjidemetriou, G.M., de Silva, L., Girolami, M., Brilakis, I., 2024. Bayesian dynamic modelling for probabilistic prediction of pavement condition. *Eng. Appl. Artif. Intell.* 133, 108637. <http://dx.doi.org/10.1016/j.engappai.2024.108637>.

Zhang, R., Nie, F., Li, X., Wei, X., 2019. Feature selection with multi-view data: A survey. *Inf. Fusion* 50, 158–167. <http://dx.doi.org/10.1016/j.inffus.2018.11.019>.

Zhang, Y.-M., Wang, H., Bai, Y., Mao, J.-X., Xu, Y.-C., 2022a. Bayesian dynamic regression for reconstructing missing data in structural health monitoring. *Struct. Health Monit.* 21 (5), 2097–2115. <http://dx.doi.org/10.1177/14759217211053779>.

Zhang, Y.-M., Wang, H., Bai, Y., Mao, J.-X., Xu, Y.-C., 2022b. Bayesian dynamic regression for reconstructing missing data in structural health monitoring. *Struct. Health Monit.* 21 (5), 2097–2115. <http://dx.doi.org/10.1177/14759217211053779>.

Zhu, X.Q., Law, S.S., Huang, L., 2018. Identification of railway ballasted track systems from dynamic responses of in-service trains. *J. Aerosp. Eng.* 31, 04018060. [http://dx.doi.org/10.1061/\(ASCE\)AS.1943-5525.0000898](http://dx.doi.org/10.1061/(ASCE)AS.1943-5525.0000898).