Track-monitoring from the dynamic response of an operational train

George Lederman^{a*}, Siheng Chen^b, James Garrett^a, Jelena Kovačević^{b,c}, Hae Young Noh^a, Jacobo Bielak^a

^a Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

^b Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

^c Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

*Corresponding author: Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

E-mail Addresses: george.lederman@gmail.com (G. Lederman), sihengc@andrew.cmu.edu (S. Chen), garrett@cmu.edu (J. Garrett), jelenak@cmu.edu (J. Kovacevic), noh@cmu.edu (H. Noh), jbielak@cmu.edu (J. Bielak)

ABSTRACT

We explore a data-driven approach for monitoring rail infrastructure from the dynamic response of a train in revenue-service. Presently, track inspection is performed either visually or with dedicated track geometry cars. In this study, we examine a more economical approach where track inspection is performed by analyzing vibration data collected from an operational passenger train. The high frequency with which passenger trains travel each section of track means that faults can be detected sooner than with dedicated inspection vehicles, and the large number of passes over each section of track makes a data-driven approach statistically feasible. We have deployed a test-system on a light-rail vehicle and have been collecting data for the past two years. The collected data underscores two of the main challenges that arise in train-based track monitoring: the speed of the train at a given location varies from pass to pass and the position of the train is not known precisely. In this study, we explore which feature representations of the data best characterize the state of the tracks despite these sources of uncertainty (i.e., in the spatial domain or frequency domain), and we examine how consistently change detection approaches can identify track changes from the data. We show the accuracy of these different representations, or features, and different change detection approaches on two types of track changes, track replacement and tamping (a maintenance procedure to improve track geometry), and two types of data, simulated data and operational data from our

test-system. The sensing, signal processing, and data analysis we propose in the study could facilitate safer trains and more cost-efficient maintenance in the future. Moreover, the proposed approach is quite general and could be extended to other parts of the infrastructure, including bridges.

KEYWORDS: rail maintenance; signal processing; change detection; vehicle-based inspection; data acquisition; position uncertainty

1. Introduction

While continuous monitoring of engineered assets has become cost effective in fields like manufacturing and aerospace, the majority of civil assets are still monitored by visual inspection [1,2]. This paper focuses on railroad infrastructure, which is monitored by a combination of visual inspections and specialty track inspection vehicles, known as track geometry cars [3]. Inspection cars are expensive to operate, and as such, they are used to inspect the tracks infrequently [4]. A number of researchers have proposed using vibration-based monitoring to assess the track profile as a lower-cost alternative to the type of optical sensors currently employed on track geometry cars [5–10]. All of these studies assume that the sensors are placed on the axle of the train so that they have a direct mechanical connection to the track. In addition, they assume the position of the train is known, typically achieved by requiring that the inspection car operates at a constant speed.

In this paper, we explore a more economical approach: using operational trains and allowing the sensors to be placed inside the cabin [11]. By using operational trains, no dedicated monitoring vehicles are required, track closures are avoided, and tracks can be interrogated more frequently. The high frequency of interrogation not only means that issues can be flagged more rapidly, but also more passes over areas of interest can be collected, meaning greater statistical confidence about the

state of the track. Lastly, by allowing the sensors to be placed in the cabin, the sensors can be installed more easily and can be protected from the external environment.

There are numerous challenges that must be overcome for this method to become practical, two of which we explore in detail in this paper. The first is that a train's speed over a section of track differs with each pass, so methods robust to train speed must be found for comparing data between passes. This is particularly challenging when the sensors are in the train's cabin, because the train's suspension filters the roughness from the track differently depending on speed. The second challenge is that we do not know the precise location of the train due to GPS noise, so position uncertainty must be considered in attempting to detect track changes from the vibration signal.

We learned about these challenges through our long-term monitoring project in partnership with the Port Authority of Allegheny County. We have placed accelerometers and a GPS antenna on an operational light-rail vehicle and have been collecting data over the last two years. We have focused on determining which analysis techniques are best suited for train-based monitoring and understanding why certain techniques work better than others. To gain insight into the vehicle-based monitoring problem, we also modeled a simplified version of the train-track interaction using a single degree-of-freedom oscillator travelling over a rough track. This parametric simulation allowed us to explore the effects of variable speed and position uncertainty in a controlled setting before validating on the light-rail data.

Simulation has three main benefits. First, we can generate much more data than we could collect from instrumented vehicles. Although we have been collecting data from the light-rail system for two years, there are relatively few recorded maintenance events in the rail-network each year that we can use to test our approach. With simulation, we can rapidly generate hundreds of track changes. Second, we can simulate a wider range of parameters than the narrow band we have

observed in our operational system. For example, while the fundamental natural frequency of the instrumented light-rail vehicle has a single value, we can simulate numerous oscillators with different fundamental frequencies to ensure that the analysis techniques we find are general. Third, in the data we have collected, we do not have the ground-truth of the train's position. Through simulation, we can study the effect of position error, and propose techniques that work well for the level of error we expect in operational systems.

While our ultimate objective is to build a complete track-monitoring system, we focus in this paper on two main components required for automated track monitoring: meaningful feature extraction from the raw vibration signals and detection of track-changes from these features. For the first component, we examined four different features, and then used supervised classification to determine which one provides the most reliable information about the state of the track. The objective is to determine which features are most robust to the sources of uncertainty inherent in train-based monitoring. The fact that the train's speed varies between each pass (or the speed of the travelling damped oscillator in simulation) made comparing the data from multiple passes challenging. The signals are recorded in the time domain, but given that track changes are localized in space, the spatial domain is often more useful for detecting track changes. We examined features based on the spatial-domain representation of the signal, spatial frequency domain and temporal frequency domain as well as features based on the energy in the signal. We show the discriminate power of these features on data from our parametric simulation as well as on the light-rail dataset.

The second component in our automated detection system, change detection, is a challenging task because we do not know a-priori the type of change we are trying to detect [12]. The most closely related study of automated track anomaly detection from the vibration signal of a train is that done by Molodova et al. [13] where a system for detecting track squats (a type of rail-surface

indentation) was proposed. In the study, a detection event was triggered by a vibration signal above a pre-defined threshold. We aim to build a broader detection system where anomalies are defined as changes relative to historical behavior. This ensures that areas with consistently high vibrations (like track switch gear or joints) are not labeled as anomalies, while changes, even in areas with low vibration-amplitude, are detected. For example, we will show that our method detects changes in track geometry due to tamping that can have safety implications despite their small amplitude. This historical detection technique allows for the monitoring of an entire network without manually tagging problematic areas as the method in Molodova et al. requires.

To perform change detection, we experimented with common methods like cumulative sum chart control (CUSUM) [14] and generalized likelihood ratio (GLR) [15], as well as a simplistic Haar filter [16,17] borrowed from the field of signal processing and computer vision. We report the performance of these approaches both on our simulated data as well as on the light-rail dataset.

As a validation, we apply these feature extraction and change detection approaches to data we collected from the instrumented light-rail vehicle. We examine our ability to detect a broken rail (a localized spike in the track profile) which is then replaced, as well as tamping, a maintenance procedure which subtly corrects track geometry. In both cases, information about when and where the maintenance took place can be found in the track maintenance logs, thus detection accuracy for our system can be evaluated.

2. Rail-Monitoring Method

Our proposed rail-monitoring method uses data collected from accelerometers inside a train to identify changes over time in rail infrastructure. In thinking about how to analyze the data, our first goal was to understand how track roughness, filtered by the train's suspension, produces vibrations

in the train's cabin. To do this we used the simplest possible model, a travelling damped oscillator as shown in Fig. 1 and conducted a simulation study as will be described in the following section. We sought to determine how best to detect changes in a section of track as the oscillator traveled repeatedly over it. This same type of model has been used previously in the track-monitoring literature [8,18], but we extended that work to include variable train speed and differing levels of position uncertainty [19].

2.1 Feature Selection Simulation

We conducted a simulation of an oscillator travelling over a rough track to understand which features, when extracted from its dynamic response, are sensitive to track changes, but robust to speed and position uncertainty. The framework of the simulation is shown in Fig. 2. First we simulated a track profile and a change in that profile, for various types of track repairs. We then simulated the oscillator passing over the roughness 100 times before and after the track-repair, where each pass over the track had a unique speed profile, and extracted relevant features from the dynamic response of the oscillator. Finally, we quantified how well we could differentiate between the data before and after the repair. We repeated the steps of the flow chart in Fig. 2 for oscillators of different natural frequencies and damping ratios. The goal was not only to understand the behavior of the oscillator for parameters most similar to those of the light-rail system, but also to gain insight into how the results would vary for different rail systems so as to make the results more general.

Fig. 1. Traveling oscillator moving over a rough track.

Fig. 2. Flow chart of the simulation. Note that this process was repeated for the three types of track changes, for a variety of damping ratios, natural frequencies, and position uncertainty levels.

Fig. 3. Three types of roughness changes (left plots), with detail of each (right plots). (a) A toy-model change (used later to provide visual intuition about the simulation) and (b) the detail of the toy-model change. (c) A spike change, characteristic of a broken rail before and after replacement, with a realistic track roughness and (d) the detail of the spike

change at 150m. (e) The tamping change, simulated using a filter with the same smoothing effect of the tamping machine and (f) the detail of tamping which occurs between 150m and 250m.

When simulating the roughness, we generated a 1km section of track for three change types, each of which is shown in Fig. 3. The first track change, in Fig. 3a, is a toy-model of roughness used to visualize the simulation later in the paper (greater detail is shown in Fig. 3b). For the second and third track changes (Fig. 3c and Fig. 3e), we simulated two types of track changes we had observed in practice, a large localized spike which is removed, simulating the replacement of a broken rail (detail in Fig. 3d), and a smoothing of a track profile, simulating tamping (detail in Fig. 3f). For each of these realistic changes, we also simulated a realistic track roughness using the parameters found in the literature [20]. For the tamping change, we filtered this track roughness as per the smoothing effect of the tamping machine in [21] as shown in Fig. 4, reducing the standard deviation of the track profile from 2mm to 1.5mm over a 200m section of track. For each simulated profile, we generated 25 samples per meter for 25,000 samples per profile; the wavelengths which most influence the fundamental frequency of train cabin (and thus our oscillator) are between 2m and 50m, so this sampling rate provided adequate resolution.

Fig. 4. Simulated change due to tamping. The wavelengths which are most effected by tamping are between 2m and 50m which are the wavelengths which most affect the dynamic response of the light-rail. The tamping change is based on data from Figure 3 of [21].

Once we generated the track roughness, the next step was to generate the response of the oscillator; two realizations of this process over the same roughness with different speed profiles are shown in Fig. 5. Fig. 5a shows the speed profiles in the time domain, while Fig. 5b shows them in the spatial domain (i.e. plotted against position). Note when the speed approached zero, no distance is covered, producing the scalloping effect in Fig. 5b, a phenomenon which is common when the train stops at a station. When generating the speed profile, we limited it to be between 0 and 15m/s (35 MPH / 55 KPH) which is the same as that of the light-rail vehicle in our deployment. The toy-model roughness profile is shown in Fig. 5d as a function of position, but the train experiences this profile in the time domain Fig. 5c. We generated the response of the oscillator in Fig. 5e by solving the following differential equation [22]:

$$\ddot{u}(t) + 2\zeta\omega_n(\dot{u}(t) - \dot{r}(t)) + \omega_n^2(u(t) - r(t)) = 0, \quad (1)$$

where ζ is the damping ratio of the oscillator, ω_n is the natural frequency of the oscillator, u is the displacement of the oscillator, and r the track roughness as shown in Fig. 1. For each pass, we simulated 200 seconds of the dynamic response train, with a time-step of 0.01s; thus each pass had a length of 20,000 samples.

Although the bumps in the track occur at the same location in space, due to the difference in the speed profiles, the oscillator is excited in the two passes at different points in time, as highlighted by the box in Fig. 5e. This posed a challenge when comparing multiple passes in the time domain, so we interpolated the signal into the spatial domain as seen in Fig. 5f. Given the response of the oscillator is smoother than the track roughness, when performing linear interpolation, we sample from the dynamic response at a rate of 10 samples/m.

The bumps experienced by the oscillator line up more closely in the spatial domain, but the effects of varying speed are still visible. The wavelengths of the oscillators (the distance over which they occur) varies considerably in Fig. 5f due to the speed (as highlighted in the box), even though the periods of the oscillations (their duration in time) are invariant, as can be seen in Fig. 5e.

Fig. 5. Two passes over the toy-roughness shown for illustration: (a) the speed profiles in the time domain; (b) speed profiles in the spatial domain; (c) the roughness interpolated in time; (d) the roughness in space; (e) the acceleration of the oscillator in time, the box highlights two oscillations from the same bump which occur at different points in time due to the oscillator's speed profile; and (f) the acceleration of the oscillator in space, the box highlights a location where the wavelengths of the oscillations differ.

The variation in Fig. 5 highlights the challenges of dealing with variable speed; a second challenge lies in not knowing the exact position of the oscillator, a phenomenon that occurs in practice due to GPS error. In order to plot the signal in the spatial domain (Fig. 5f), we used the

position of the oscillator, $x(t)$. In practice, we only know some approximation of the position, $x_\epsilon(t) = x(t) + \epsilon(t)$, where ϵ is the error. Let us assume that this error is normally distributed with zero-mean and standard deviation σ . We show the effect of this error with different standard deviations in Fig. 6.

As the train's position and the associated uncertainty are a central part of this study, the method in which the uncertainty is applied requires further comment. The position vector, x , is generated based on the speed profiles as shown in Fig. 5. Because the speeds are positive, the position either increases or remains the same with each step. In order to simulate position uncertainty, we add Gaussian noise, ϵ_0 . From a practical standpoint, in order to interpolate the data spatially, the position vector must be monotonically increasing, so this noise poses a challenge. Physically, if the position vector does not increase monotonically, it means the train is moving backwards at some time-steps, something we know does not happen. To avoid this problem, we sort the values,

$$x_\epsilon = \text{sort}(x + \epsilon_0).$$

This new position vector differs from the original position vector according to ϵ ,

$$\epsilon = x - x_\epsilon.$$

This resulting error follows the normal distribution according to $\epsilon \sim N(0, \sigma)$. The standard deviation of the resulting error, σ , is used to characterize the extent of the position uncertainty throughout this text.

While in Fig. 5f we showed the vibration signal for two passes in the spatial domain, in Fig. 6a we show 200 passes, where each horizontal line is a vibration signal from a single pass, and the color is indicative of the instantaneous acceleration. Note that we show 100 passes before the track change, and 100 passes after the track change from Fig. 3a. In Fig. 6b and Fig. 6c we show higher

levels of position uncertainty by varying the standard deviation of the error, σ . As the position uncertainty grows, detecting changes in the tracks becomes harder, although detection is still possible (even visually) for this type of a trivial track change.

Fig. 6. Effect of position uncertainty. Each row of the above plots shows one pass of the oscillator over the toy-model roughness shown in Fig. 3. Passes 1-100 correspond to the “before” roughness, while passes 101-200 correspond to the “after” roughness. The color along each line represents the acceleration of the oscillator in m/s^2 but has been truncated at current bounds to show greater clarity. Note that each pass has a unique speed profile. (a) Shows the response of the oscillator in space with no position uncertainty, (b) shows the response with added uncertainty (zero mean and standard deviation of $\sigma = 10\text{m}$).

In Fig. 6 we show the vibration signal plotted spatially as a more intuitive signal representation given our interest in particular locations of track. To quantify how well this representation portrays track condition, we consider it as a feature, and compare it with three other features, temporal-frequency, spatial-frequency and signal-energy, to see which provides the greatest discrimination of track condition. Each is described in the equations below where $\mathcal{F}[\cdot]$ denotes the Fourier Transform, $\ddot{u}_n$ is the vector of length p of collected data from the n^{th} pass, f_n is the feature vector (length p), and x_ϵ is the position vector with added noise ϵ (also length p).

Temporal-frequency	$f_n = \mathcal{F}[\ddot{u}_n]$
Spatial-frequency	$f_n = \mathcal{F}[\ddot{u}_n _{x_\epsilon}]$
Spatial-amplitude	$f_n = \ddot{u}_n _{x_\epsilon}$
Signal-energy	$f_n = \ddot{u}_n^2 _{x_\epsilon}$

The motivation for having two types of frequency-based features is the prevalence of frequency-based features in the literature [23–26] or features related to frequency, such as wavelets [27]. “Temporal” here means the signal in the time domain as it was originally acquired, whereas “spatial” denotes the signal has been linearly interpolated spatially given its speed profile. The signal-energy feature is similar to the spatial-amplitude, but it is squared prior to interpolation. This squaring makes a difference because it makes the feature have a non-zero mean. Prior to classification, the spatial-amplitude and signal-energy features were averaged over a moving 25m window of track. This step increased robustness to position uncertainty for both features, but can filter out some zero-mean oscillations.

To evaluate the features, we used supervised classification; the goal was to see which features allow for the clearest indication of track change. Supervised classification means that the “label” of the pass is known during a “supervised” training phase. In this case we use a binary label as to whether the change in the track has occurred. (By contrast, “unsupervised” approaches are a broad family of algorithms which require no training, like the change detection approaches explained later in the text.) Here we used a support vector machine classifier with a linear kernel (we choose a simple model to avoid overfitting given the relatively few passes we use in our experiments with operational data). We selected 150 passes out of the original 200 for training (the remaining 50 are for testing), and we repeated for 50 fold cross validation. In each case, we focus on a 500m section of track around the change of interest, a procedure we also use later in this study for

examining changes in the operational data. (This means the feature vector has a length of 5000 samples for the spatial data, and between 2000 and 10000 samples for the time data, depending on the train's speed).

High classification accuracy means that the data is more easily separable, i.e. the feature is useful. The features were long, so we select only the 50 most discriminative indices of the feature. To accomplish this, we use the technique described in [25], where we find the mean signal for both classes from the training data, then define the most discriminative indices as those with the greatest difference between the two mean signals. In total, to explore feature selection, we simulate 180,000 passes of the oscillator over 1km of simulated track (3 types of track changes $\times$ 6 damping ratios $\times$ 5 natural frequencies $\times$ 10 levels of position uncertainty $\times$ 200 passes each). The 6 damping ratios are 0, 0.1, 0.2, 0.3, 0.4 and 0.5. The five natural frequencies are $3\pi/2$, 2π , $5\pi/2$, 3π , and $7\pi/2$ rad/s.

2.2 Feature Selection Simulation Results

The classification results from the simulation are shown in Fig. 7 and Fig. 8. This was a binary classification, so 50% accuracy would mean the labels are random, while higher accuracy means more consistency in determining the state of the track given the selected features. We report accuracy while varying the level of position uncertainty (Fig. 7a), the oscillator's damping ratio (Fig. 7b) and the oscillator's natural frequency (Fig. 7c). In Fig. 7a, we kept the damping and frequency constant at values we observe in the light-rail system ($\zeta = 0.2$; $\omega_n = 5\pi/2$) [28]. We found that temporal-frequency features provide low accuracy at all levels of position uncertainty, spatial-amplitude and spatial-frequency fall in accuracy as the position uncertainty increases, and signal-energy offer relatively high accuracy at all levels of position uncertainty. It is not surprising that temporal-frequency performed badly because it is sensitive to speed changes, which we varied

between each run. Spatial-amplitude and spatial-frequency require interpolation by x_ϵ , so they are sensitive to increases in the uncertainty ϵ . Signal-energy also relies on interpolation by x_ϵ , but is less sensitive to position error as it represents the level of roughness or excitation caused by the track, rather than the specific state of the oscillator (it is always positive independent of the oscillator).

We also investigated the effect of varying the damping ratio (Fig. 7b) and the natural frequency (Fig. 7c) if the location uncertainty is kept constant at $\sigma = 7\text{m}$. For this level of position uncertainty, signal-energy is the only feature with strong discriminative power. When the damping ratio was zero, the oscillations kept growing over the course of a pass, so the localized track change are undetectable even from the signal-energy feature. If the damping ratio is large, i.e. $\zeta = 0.5$, the changes are so localized that they become difficult to detect given the position uncertainty. Thus, a moderate level of damping ($\zeta \approx 0.2$) appears preferable; this conveniently corresponds with the observed damping ratio of the train. For variation of the natural frequency (Fig. 8c), oscillator stiffness has relatively little effect on the accuracy over the range of values considered. In the case of tamping in Fig. 8c, a stiffer oscillator leads to more localized changes and because tamping is effectively many small smoothing changes over a section of track, these localized changes are more difficult to detect.

Overall, the signal-energy feature outperforms the other features as it is robust to position uncertainty for both types of track changes. This is particularly convenient as the signal-energy feature could potentially be universally applied in this type of monitoring system. In industry, the most common parameter for measuring track geometry is the standard deviation of the track profile, and signal-energy of the vibration signal is related to this parameter, as both involve squaring their respective data points. The relationship between the vibration signal and the track profile, given in

Eqn. (1), is linear and time invariant, which means that the frequency components of the response are simply scaled versions of the frequency components of the roughness. The strong performance of signal-energy shows that the traditional parameters for track monitoring can inspire new features that are useful for statistical track monitoring.

Fig. 7. Classification accuracy for spike change. (a) Effect of position uncertainty for oscillator with $\zeta = 0.2$ and $\omega_n = 2.5\pi$ rad/s. (b) Effect of varying damping ratio while uncertainty $\sigma = 7m$ and $\omega_n = 2.5\pi$ rad/s. (c) Effect of varying natural frequency while $\zeta = 0.2$ and uncertainty $\sigma = 7m$.

2.3 Change Detection Simulation

Having established that among the different criteria examined, signal-energy represents the most robust feature, in this section we explore how we can achieve our second goal: to determine when a change occurs in the tracks.

Fig. 8. Classification accuracy for tamping change. (a) Effect of position uncertainty for oscillator with $\zeta = 0.2$ and $\omega_n = 2.5\pi$ rad/s. (b) Effect of varying damping ratio while uncertainty $\sigma = 7m$ and $\omega_n = 2.5\pi$ rad/s. (c) Effect of varying natural frequency while $\zeta = 0.2$ and uncertainty $\sigma = 7m$.

Fig. 9. Flow chart of change detection simulation. Note that this procedure is repeated for different track changes and different levels of position uncertainty.

We conducted a simulation to study change detection approaches by following the flow chart shown in Fig. 9. First we generated many passes of the train before and after a track change (varying the speed profile as in the previous simulation study). We computed the signal-energy feature from the vibration signal, then extracted the data at a particular track location across all N passes. Although a number of change-detection approaches exist, many of them are designed for finding statistical change in a scalar quantity. In this application, we compared data from different passes of a moving vehicle, so each pass was a vector. To format the data in a way that we could use these general change detection approaches, we build a matrix, F , where each feature vector f_n^T is a row in the matrix (the superscript “T” indicates the transpose of the vector):

$$F = \begin{bmatrix} f_1^T \\ f_2^T \\ \vdots \\ f_n^T \\ \vdots \\ f_N^T \end{bmatrix}.$$

Each element within the matrix, $F(n, i)$, describes the signal-energy feature at a specific location on the track, i , and a particular pass, n . Using this formulation, we experimented with three different change detection filters: cumulative summation control chart (CUSUM), generalized likelihood ratio (GLR) and a Haar filter.

All three change-detection methods effectively detect when there is a change in the mean-value of the feature at a particular location of track compared to that value for a window (set number) of previous passes. Details of each will be presented briefly to explain how they are applied within the two dimensional feature matrix.

1. Cumulative summation control chart (CUSUM) estimates the mean value over a window of previous passes (known as a sliding window), then compares the current (or most recent value) to the historic mean [15]. If there is a succession of passes all deviating from the mean in one direction, it will trigger a detection event because it is likely that the mean has changed. Mathematically, the mean, $\hat{\Theta}$, is estimated from historical data over a window of $w + 1$ previous points,

$$\hat{\Theta}(n, i) = \frac{1}{w + 1} \sum_{k=n-w}^n F(k, i).$$

We then find how the data from the current pass differs from this mean,

$$\delta(n, i) = F(k, i) - \hat{\Theta}(n, i).$$

Finally we add this difference, δ , to a running sum of the differences,

$$g(n, i) = g(n - 1, i) + \delta(n, i).$$

If the data varies from the mean consistently in one direction (positive or negative) intuitively the data has changed because the mean no longer represents the data. We say a change has occurred when the magnitude of g exceeds a threshold, h , $|g(n, i)| > h$.

2. The generalized likelihood ratio (GLR) looks for a change in the underlying distribution, and quantifies the log-likelihood that the recent data is derived from the historical distribution [15]. If the recent data appears dissimilar from the historical data, then the

approach will trigger a detection event. Mathematically, let us call the first window of data y_0 , and the second window of data y_1 :

$$y_0(n, i) = F(n - w - 1 : n - 1, i)$$

$$y_1(n, i) = F(n : n + w, i).$$

The distribution of y_0 is defined by H_0 and the distribution of y_1 by H_1 . We want to see the likelihood ($L[\cdot]$) that y_1 comes from the same distribution as y_0 , H_0 , versus the likelihood it comes from H_1 .

$$g(n, i) = 2 \log \frac{L[y_1(n, i) | H_1(n, i)]}{L[y_1(n, i) | H_0(n, i)]}$$

We say a change has occurred when g exceeds a threshold, h , i.e $g(n, i) > h$. In other words, beyond some threshold, the data is so much more likely to have come from a new distribution than from the historical distribution that we say a change in the data has occurred.

3. The Haar filter finds the difference between the sum of recent data and the sum of historical data, triggering a detection event if the difference is too large. Mathematically, the difference between the two windows of data of width $w + 1$ can be written as

$$g(n, i) = \left| \sum_{k=n}^{n+w} F(k, i) - \sum_{k=n-w-1}^{n-1} F(k, i) \right|.$$

As in previous approaches, we say a change has occurred when $g(n, i) > h$. This approach tends to find locations where a step change has occurred between the two windows of data.

Each of the approaches relies on a sliding window of data: shorter-sized windows have the potential for detecting changes more rapidly after they occur, while longer windows allow for greater statistical significance; here we chose a window size of 20 passes, which would allow our

system to detect a change within a week (assuming a few passes over the track per day). The results of applying each of these filters to the simulated changes can be seen in Fig. 10 as will be discussed in more detail in the next section.

2.4 Change Detection Simulation Results

In this section, we discuss the performance of the three change-detection approaches described previously. Results of each approach applied to one example of each type of simulated track change are shown in Fig. 10. Successful change-detection filters should produce a high value after the change occurs at the same location in the track, and successful detection has been indicated with a red-box. Note that the change is not necessarily detected immediately, as it takes several passes after the change to build up statistical significance that a change has indeed occurred. All three approaches detected the track change on the left in Fig. 10b-d (with an ideal threshold), while only the Haar filter was able to detect the tamping change in Fig. 10h.

Fig. 10. The change detection filters applied to the data. (a) one example of a track replacement change showing the energy feature where the spike (due to a broken rail) is removed at 150m, (b-d) show three change detection approaches applied to this track change, (e) shows a tamping change between 150 and 250m using the signal-energy feature, while (f-h) show the three change detection approaches applied to this data. Values in each figure (shown in color) have been normalized on $[0,1]$. The red-boxes indicate the true-positive events that could be detected with an ideal threshold. There are no boxes in (f) and (g) as the methods failed to detect the change. The data shown in this figure has no position uncertainty.

CUSUM is perhaps the most versatile of the methods, as it could detect a gradual change in the mean over time, while GLR and Haar work best for abrupt changes in the data. In our case, we assume the track change occurs over a short period of time, so the change is complete before the train passes over that section of track again. Given this assumption, GLR and Haar might be

expected to perform better than CUSUM. GLR is more complex than the Haar filter as it involves not only the mean, but also the standard deviation in calculating the distribution of the data before and after the change. A change in variance could trigger a detection event, but it is unlikely in our application that high variance would be due to an infrastructure change, so GLR is less robust for this application. It is the simplicity of the Haar filter that enables it to consistently detect tamping along the section of track where it was simulated in Fig. 10h.

Fig. 10 shows just one example of each type of change, but we quantified the success of each approach over the 100 changes we simulated. We defined a true positive as a change detected within 20 passes and within 25m of a true change. A false positive was a change detected outside of this window. We report false positives and false negatives (true positives that were not detected) for a range of thresholds in Fig. 11a, because selecting the appropriate threshold itself can be difficult. Methods for selecting the best threshold are beyond the scope of this paper, so we report the error level assuming the ideal threshold was selected in order to see which change detection technique has the most potential. Assuming false positives, FP , and false negatives, FN , are equally bad, we select the threshold level that will minimize the larger of the two errors, which occurs where the false negative and false positive curves intersect as show in Fig. 11a. In other words, we want to select a threshold, h , where

$$h = \underset{h}{\operatorname{argmin}}(\max[FN, FP]).$$

In total, for understanding the trade-off in the change detection approaches, we simulated 200,000 passes (2 types of track changes $\times$ 100 track changes $\times$ 100 passes $\times$ 10 levels of uncertainty). While in the feature selection study (Section 2.2), we used only 3 track profiles (track change, tamping change and toy-roughness) in this case we simulate 100 track profiles for track changes and 100 track profiles for tamping changes. Also, in feature selection study (Section 2.2), we considered a

range of natural frequencies and damping ratios. In this simulation study we use only the damping ratio and natural frequency values we have observed in our train system ($\omega_n = 2.5\pi, \zeta = 0.2$).

Fig. 11. Change detection results from the simulation. (a) A typical plot of false negatives (FN) and false positives (FP) as the threshold is varied, shown with data for the 100 examples of this type. In this case, the plot is shown for CUSUM with no position uncertainty, and where the two lines cross, there is 43% error of both types. The data has been normalized on $[0,1]$ so that the threshold spans the whole range. (b) The minimum error for all three approaches and all position uncertainty levels for the spike change. (c) The minimum error for the tamping change.

In Fig. 11b the lowest error (i.e. ideal threshold) is reported for all change detection approaches on the track change simulation. While the error increased for all methods as position uncertainty increases, the Haar filter consistently performed the best. Fig. 11c shows the same error quantification but for the tamping change simulation, and again the Haar filters outperforms the other methods. Because tamping occurs over a larger section of track, position uncertainty matters less as we saw in the classification results shown in Fig. 8a, so it is consistent that the Haar filter remains relatively unaffected by position uncertainty.

3. Validation on the Light-Rail Vehicle Dataset

The previous sections provide a general understanding of which features and detection filters work well to detect changes in a rough track from the dynamic response of a travelling oscillator. In

this section, we investigate whether the same findings hold true for our test-system deployed on a light-rail vehicle.

3.1 The Light-Rail Vehicle Dataset

We have collected an extensive dataset over the last two years; this is the first study to our knowledge on vibration-based track monitoring in which data was collected over such a long period while the trains were in operational service and all of the raw data was kept. We provide details about our instrumentation, about two infrastructure changes of interest, and, finally, we assess the validity of the approaches on this dataset.

The sensors on the train were installed through our partnership with the Port Authority of Allegheny County; the location of the sensors can be seen in Fig. 12. We placed two uni-axial accelerometers inside the cabin of the train to measure vertical vibrations, one tri-axial accelerometer on the central wheel-truck (bogie), and a BU-353 GPS antenna mounted just under the roof. In this study, we only use data from the accelerometer mounted inside the cabin above the front wheel-truck, a Vibra-Metrics model 5102 sensor. The train car is a CAF light-rail vehicle that has an empty mass of 40 metric tons, and a length of 27m.

We have collected data from hundreds of complete passes through the 30km rail network, sampling continuously from the GPS at 1Hz and from the accelerometer at 1.6kHz. Details of this deployment can also be found in a previous conference proceeding [11]. One of the challenges in analyzing the data has been handling the varying speed and the position uncertainty, issues which inspired the above simulations.

Fig. 12. Instrumentation of light rail vehicle. A GPS antenna is just below the roof the train, the data acquisition system is in a cabinet near in the conductor, and the accelerometers are inside the main cabin and on the central wheel truck. In this study, all data is from the accelerometer highlighted with the red box, which measures vertical vibrations inside the cabin of the train.

Our simulations aimed to model the natural frequency of the train; a comparison between the simulated data and the collected data is shown in Fig. 13. While there is more noise at high frequencies in the measured data (Fig. 13c), both the oscillator and train have a fundamental natural frequency of around 1.25 Hz. Our simple oscillator does not attempt to model all the complexities of the train system. For example, the air-conditioning system is one of the biggest sources of noise within the train cabin. However, the train's response to the track roughness below 3 Hz does appear similar to the single-degree-of-freedom oscillator.

Fig. 13. Frequency response of the simulated oscillator and the operational train. (a-b) Response of the simulated oscillator at two different scales. (c-d) Response of the train at two different scales. In (c), the measured response has a strong narrow peak at 30Hz, which is a resonance frequency of the 60 Hz electricity used inside the cabin. In general, the dominant response of the train is around 1.25 Hz, which is the same as the natural frequency of the oscillator we have simulated. Note that the amplitude of the lowest frequencies of the measured response (0 - 0.5 Hz) are reduced by the sensor; it is likely that the values below 0.5 Hz should be higher, but the accelerometer, which is a shear piezo-electric, has low sensitivity at such low frequencies.

3.2 Track Change in the Light-Rail Dataset

In September 2014, the owner of the light-rail system replaced the track in an old road-crossing. We use this known maintenance activity, where faulty track was replaced with good track, to test our signal analysis pipeline. Fig. 14a and b show two characteristic passes of the train over the 1km section of track. In the second pass the train stops at two stations, the first at 200m and the second at 520m, whereas in the first pass it only stops at one of the stations (at 200m). The train stops when there are passengers to pick up or let off, which poses a challenge when comparing data

between passes. One hundred passes over this section of track are shown in Fig. 14c, where the color indicates the value of accelerations recorded. Although difficult to see in Fig. 14, there is a high frequency vibration event (up to 100 Hz) when crossing a road at 220m, an event which no longer occurs after the 50th pass, when the repair is done. Although the simulations presented earlier only model the fundamental natural frequency of the train, in this particular case, the broken rail excited higher modes of the train. The spatial-amplitude feature does not show this high-frequency event clearly, but the signal-energy feature, shown in Fig. 15, does show this clearly. As the train crosses the road at 220m, high signal-energy can be seen in Fig. 15a. This spike is absent in Fig. 15b or after pass #50 in Fig. 15c due to the repair.

Fig. 14. Spatial signal. (a) A pass before repair showing both the train speed and vertical vibrations from a sensor inside the train. (b) A pass after repair. (c) 50 passes before and 50 passes after the repair, where each pass is a horizontal line and the color indicates the instantaneous acceleration. With the spatial-amplitude feature, the track change is nearly impossible to see.

Note as well that the first station stop in Fig. 14(a) and Fig. 15(a) occurs just after 200m, whereas in (b) it occurs just before 200m. This is likely due to the orientation of the train. The GPS is at one end of the train, and the train always stops at the same location within the station. If the orientation of the car changes, the position can differ by 27m, the length of the car. Due to the

sensor's location on the train, it will interrogate the tracks near the stations at different speeds. A sensor at the front of the train will travel quickly over tracks at the beginning of the station, and slowly over tracks towards the end of the station; the reverse is true for a sensor at the back of the train, adding to the difficulty of comparing data directly between different runs.

Fig. 15. Signal-energy. (a) A pass before repair showing both the train speed and vertical vibrations from a sensor inside the train. (b) A pass after repair. (c) 50 passes before and 50 passes after the repair, where each pass is a horizontal line and the color indicates the magnitude of the signal-energy feature. With the signal-energy feature, the track change is clearly visible.

As in the simulation, we extracted different features from the data and used classification to determine which features allow the infrastructure change to be detected most readily. In a binary classification between the two states, we achieve 91% accuracy in Fig. 16a, drawing data from a 500m section of track shown in Fig. 16c. The second best feature appears to be the temporal-frequency feature. However, as we saw in the simulation, this feature would not be expected to detect track changes given the train's variable speed. It is far more likely that the feature is detecting temperature differences between the two classes as the data was collected over a year, and Pittsburgh is a temperate region with significant temperature variability. Temperature affects the

data in a number of ways: for example, we have observed higher noise levels inside the train cabin during the summer months due to the vibrations of the air conditioning system. More subtle effects include changes in the elasticity of rubber, which affects both the train's damping and the properties of the track [29]. That temperature impacts the signal is supported by the results shown in Fig. 16b where classification results are presented for an adjacent 500m section of track where no work was done. Classification using the temporal-frequency feature allows for 78% accuracy despite the fact that the infrastructure has not changed, whereas classification using the signal-energy feature allows for 67% accuracy (close to random) meaning it is more robust to environmental conditions. Changing external temperature, which could affect the frequency of the signal, is one possible explanation for the apparent difference in the signal despite the fact that the track itself did not change. Again we use 75% of the data for training, 25% for testing and 50-fold cross validation as in the classification in the simulation study.

Fig. 16. Classification accuracy of a 500m section (a) of a track where rails were replaced, and (b) of a track nearby where no work was done. (c) Shows where the data for the classification was drawn from and the two classes used in the binary classification. High classification accuracy means the classes are separable, and 50% accuracy means the classes are not separable leading to random classification. Signal-energy is sensitive to track changes because it achieves 91% accuracy when there is a track change, and is close to random when there are no track changes, meaning it is not classifying based on environmental factors.

3.3 Tamping Change in the Light-Rail Dataset

Tamping is an important maintenance procedure used to improve track geometry. The tamping machine measures the profile of the track, then adjusts the ballast below the track to produce a smoother, safer ride. One future goal could be to use our data to optimize tamping schedules. At this early stage though, we are most interested in identifying which features are sensitive to tamping, and which detection approach works best for this feature. In the summer of 2014, three 500m sections of track were tamped as shown in Fig. 17, starting at 1500m, 2000m and 2500m. Unlike the track change in the previous section, the change due to tamping is subtle, but occurs over a much larger section of track. Notice how in the regions with no tamping, the signal is relatively consistent over all 100 passes, while in the areas with tamping, there is more signal-energy (i.e. more bumpy ride) before tamping, and less signal-energy (i.e. smoother ride) after tamping.

We tested all four features in a binary classification to discriminate between before and after a tamping event; we found that the signal-energy feature is the best for detecting this change. The classes in the classification and the region from which the data was drawn are shown in Fig. 18c. Signal-energy achieves 90% classification accuracy (shown in Fig. 18a) with a change, and <60% accuracy (effectively random) in Fig. 18b when there is no change. This means that the signal-energy feature is classifying due to the change in the track and not an environmental variable. Frequency and spatial-frequency features on the other hand report 70% accuracy if there is a change (as in Fig. 18a), and 70% with no change (as in Fig. 18b). We can assume the discriminative power stems from temperature changes and not from infrastructure changes. Note that for consistency, we only classified data from a 500m section of track as in the previous section, but the tamping occurs

over a larger region. Classification accuracy can be slightly increased by considering a larger section of track.

Fig. 17. Region of track where tamping occurs with (a) showing the spatial-amplitude feature and (b) showing the signal-energy feature. The tamping maintenance was done three times on three separate days due to the limitations on how much work the tamping machine can do per day. Note that the peaks between 2400 and 2500m (and between 400 and 500m) are due to switchgear in the tracks.

Fig. 18. Classification accuracy of a 500m section (a) of a track where tamping work was done, and (b) of a track nearby where no work was done. As in Fig. 16, the high accuracy of the signal-energy feature where there is track work shows it is sensitive to infrastructure changes, and the low accuracy (almost random) where no work has been done shows it is robust to environmental variables. (c) We show the data used for the classification in relation to Fig. 17b, both in terms of which 500m sections of track were used, and how the two classes in the binary classification were defined.

Fig. 19. Change detection on the light-rail dataset for the track replacement change (a)-(d) and tamping change (e)-(h). Panels (a) and (e) show the signal-energy feature while (b)-(d) and (f)-(h) show the respective change detection techniques. The red-boxes indicate true-positive changes that could be detected with an appropriate threshold.

3.4 Detecting Change in the Light-Rail Dataset

The ultimate goal of the project is to automatically detect changes in the infrastructure. As such, change detection is a vital step. Although we have collected hundreds of passes through the

rail-network, we still have relatively few known infrastructure changes. Thus, we do not have sufficient data to rigorously test different threshold levels as we did in the simulation. Instead, we applied the change detection methods on the operational data and present the results with a basic threshold applied in Fig. 19. This figure mirrors the results in Fig. 10, where the raw data and all three approaches are shown. Figure 19a shows the signal-energy feature for the track replacement change, and Figure 19b for the tamping changes. Fig. 19 c-h show the results of applying the three change detection approaches to each of these two changes, in which the red-boxes indicate the change-detection approach was successful.

For both types of track changes on the light-rail dataset, the Haar filter performed the best, reliably detecting both track and tamping changes (Fig. 19d and h) as predicted by our simulation results. For the track-type change, CUSUM and Haar filter (Fig. 19b and Fig. 19d respectively) could detect a change with zero error given the correct threshold. For the tamping change, CUSUM and GLR fail to detect the change at all. CUSUM fails because the variability in the sections with high energy (like the switchgear at 2500m) are larger than the change from tamping itself, which occurs on low-energy sections of track. It is unclear why the GLR method has false-positives, although it could be due to the GLR's sensitivity to changes in variance as discussed earlier. It is important to note that the Haar filter succeeds in our original goal: detecting changes relative to historical behavior, rather than simply detecting areas of the track that produce large vibrations. The high values around 2500m in Fig. 19a due to switchgear do not affect the detection in Fig. 19h.

4. Concluding Remarks

In this paper, we propose a new way to monitor rail infrastructure by recording data from accelerometers in the cabin of a train while it operates in revenue service. To realize such a system, new features must be extracted from the vibration signal that are able to characterize the rails

despite the filtering from the train's suspension, the train's changing speed, and the noise in the GPS location. We tested four features, temporal-frequency, spatial-frequency, spatial-domain and signal-energy on simulated data and found signal-energy to be the best feature for detecting both track changes and tamping changes. We then tested the same features on the operational data from the light-rail system, and showed that signal-energy was the most sensitive to infrastructure changes and the least sensitive to environmental variability. Both in our simulations and in the operational data, we found that frequency-based features do not work well, despite their widespread use in many structural health-monitoring studies.

In addition to feature extraction, we studied a variety of unsupervised change detection approaches, including CUSUM, GLR, and the Haar filter. We found that the Haar filter outperformed the other approaches on both the simulated data and operational data, as it was particularly robust to position uncertainty.

From the results of this study, it appears that vehicle-based monitoring could be a low-cost approach to monitor our infrastructure; we were able to detect changes on an operational system using just a single sensor on a train in revenue service. Future work should look at data for additional types of changes in rail infrastructure, and more instances of each type, as well as how thresholds can be selected when performing change detection. Although only sudden changes have been studied in this paper, detecting gradual change should be studied as many changes of interest in a rail network occur gradually. Lastly, this paper presents statistical approaches which do not consider the underlying physics of the problem; hybrid approaches combining statistical approaches with physics-based approaches could provide additional insight.

Acknowledgements

This material is based on work supported by the National Science Foundation through a Graduate Research Fellowship for Lederman under Grant No. 0946825, by National Science Foundation awards 1130616 and 1017278, and a University Transportation Center grant (DTRT12-G-UTC11) from the US Department of Transportation. The authors also gratefully acknowledge the reviewers for their helpful comments and the Port Authority of Allegheny County for their partnership, particularly David Kramer, and the many helpful workers of Amalgamated Transit Union Local 85.

References

- [1] C.R. Farrar, K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*, John Wiley & Sons, 2012.
- [2] J.M.W. Brownjohn, Structural health monitoring of civil infrastructure, *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.* 365 (2007) 589–622. doi:10.1098/rsta.2006.1925.
- [3] D. Barke, W.K. Chiu, Structural Health Monitoring in the Railway Industry: A Review, *Struct. Health Monit.* 4 (2005) 81–93. doi:10.1177/1475921705049764.
- [4] C. Esveld, *Modern Railway Track: Digital Edition*, 3.4 edition, MRT-Productions, 2015.
- [5] M. Molodova, Z. Li, R. Dollevoet, Axle box acceleration: Measurement and simulation for detection of short track defects, *Wear.* 271 (2011) 349–356. doi:10.1016/j.wear.2010.10.003.
- [6] P.F. Westeon, C.S. Ling, C. Roberts, C.J. Goodman, P. Li, R.M. Goodall, Monitoring vertical track irregularity from in-service railway vehicles, *Proc. Inst. Mech. Eng. Part F J. Rail Rapid Transit.* 221 (2007) 75–88. doi:10.1243/0954409JRRT65.
- [7] Y. Oshima, K. Yamamoto, K. Sugiura, A. Tanaka, M. Hori, Simultaneous monitoring of the coupled vibration between a bridge and moving trains, in: *Proc. Fifth Int. IABMAS Conf. Phila. USA 11-15 July 2010*, n.d. <http://www.crcnetbase.com/doi/abs/10.1201/b10430-114> (accessed November 5, 2012).
- [8] J. Real, P. Salvador, L. Montalbán, M. Bueno, Determination of Rail Vertical Profile through Inertial Methods, *Proc. Inst. Mech. Eng. Part F J. Rail Rapid Transit.* 225 (2011) 14–23. doi:10.1243/0954409JRRT353.
- [9] M. Bocciolone, A. Caprioli, A. Cigada, A. Collina, A measurement system for quick rail inspection and effective track maintenance strategy, *Mech. Syst. Signal Process.* 21 (2007) 1242–1254. doi:10.1016/j.ymssp.2006.02.007.
- [10] Y. Naganuma, M. Kobayashi, M. Nakagawa, T. Okumura, Condition monitoring of shinkansen tracks using commercial trains, in: *2008 4th IET Int. Conf. Railw. Cond. Monit.*, 2008: pp. 1–6.
- [11] G. Lederman, H. Young Noh, J. Bielak, Rail-infrastructure monitoring through the dynamic response of a passing train, in: *IWSHM*, Stanford, CA, 2015. <http://dpi-proceedings.com/index.php/SHM2015/article/view/844/844> (accessed November 9, 2015).
- [12] H. Noh, R. Rajagopal, A.S. Kiremidjian, Sequential structural damage diagnosis algorithm using a change point detection method, *J. Sound Vib.* 332 (2013) 6419–6433. doi:10.1016/j.jsv.2013.07.005.
- [13] M. Molodova, Z. Li, A. Nunez, R. Dollevoet, Automatic Detection of Squats in Railway Infrastructure, *IEEE Trans. Intell. Transp. Syst.* 15 (2014) 1980–1990. doi:10.1109/TITS.2014.2307955.
- [14] E.S. Page, Continuous Inspection Schemes, *Biometrika.* 41 (1954) 100–115. doi:10.2307/2333009.
- [15] F. Gustafsson, *Adaptive Filtering and Change Detection*, 1 edition, Wiley, Chichester New York, 2000.
- [16] P. Viola, M. Jones, Rapid object detection using a boosted cascade of simple features, in: *Proc. 2001 IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. 2001 CVPR 2001*, 2001: p. I-511-I-518 vol.1. doi:10.1109/CVPR.2001.990517.
- [17] M. Vetterli, J. Kovacevic, *Wavelets and Subband Coding*, CreateSpace Independent Publishing Platform, 2013.
- [18] E.G. Berggren, M.X.D. Li, J. Spännar, A new approach to the analysis and presentation of vertical track geometry quality and rail roughness, *Wear.* 265 (2008) 1488–1496. doi:10.1016/j.wear.2008.01.029.
- [19] T. Matarazzo, S. Pakzad, STRIDE for Structural Identification Using Expectation Maximization: Iterative Output-Only Method for Modal Identification, *J. Eng. Mech.* (2016) 4015109. doi:10.1061/(ASCE)EM.1943-7889.0000951.
- [20] Statistical Representations of Track Geometry, Federal Railway Administration, 1980. <http://ntl.bts.gov/lib/45000/45000/45025/DOT-TSC-FRA-80-04-II.pdf>.
- [21] C. Esveld, A. Jourdain, G. Kaess, M.J. Shenton, Historic data on track geometry in relation to maintenance, *Rail Eng. Int. Ed.* (1988). http://www.esveld.com/Download/TUD/Esveld_Historic_data_on_track_geometry.pdf.

- [22] R.R. Craig, A.J. Kurdila, *Fundamentals of Structural Dynamics*, 2nd ed., Wiley, 2006.
- [23] G. Lederman, Z. Wang, J. Bielak, H.Y. Noh, J.H. Garrett, S. Chen, J. Kovacevic, F. Cerda, P. Rizzo, Damage quantification and localization algorithms for indirect SHM of bridges, in: IABMAS, Shanghai, CN, 2014.
- [24] Y. Oshima, K. Yamamoto, K. Sugiura, Damage assessment of a bridge based on mode shapes estimated by responses of passing vehicles, *Smart Struct. Syst.* Volume 13, Number 5 (2014).
- [25] F. Cerda, S. Chen, J. Bielak, J.H. Garrett, P. Rizzo, J. Kovacevic, Indirect structural health monitoring of a simplified laboratory-scale bridge model, *Smart Struct. Syst.* 13 (2014) 859–868.
- [26] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J.H. Garrett, J. Kovacevic, Semi-Supervised Multiresolution Classification Using Adaptive Graph Filtering With Application to Indirect Bridge Structural Health Monitoring, *IEEE Trans. Signal Process.* 62 (2014) 2879–2893. doi:10.1109/TSP.2014.2313528.
- [27] A. Caprioli, A. Cigada, D. Raveglia, Rail inspection in track maintenance: A benchmark between the wavelet approach and the more conventional Fourier analysis, *Mech. Syst. Signal Process.* 21 (2007) 631–652. doi:10.1016/j.ymssp.2005.12.001.
- [28] Dumitriu, Madalina, Influence of the Suspension Damping on the Ride Comfort of Passenger Railway Vehicles, *Sci. Bull. Univ. Politehnica Buchar. Ser. D.* 74 (2012).
- [29] G. Squicciarini, D.J. Thompson, M.G. Toward, R.A. Cottrell, The effect of temperature on railway rolling noise, *Proc. Inst. Mech. Eng. Part F J. Rail Rapid Transit.* (2015) 954409715614337. doi:10.1177/0954409715614337.