Validation of vibration-based structural health monitoring on retrofitted railway bridge KW51[⋆,⋆⋆]

K. Maes[a,*], L. Van Meerbeeck[a], E. P. B. Reynders[a], G. Lombaert[a]

[a] *KU Leuven, Department of Civil Engineering, Leuven, Belgium*

Abstract

In this paper, the feasibility of structural health monitoring based on natural frequencies is investigated for a steel bowstring railway bridge in Leuven, Belgium. The data used in the study are obtained from an ongoing long-term monitoring campaign on the railway bridge and include acceleration measurements on the bridge deck and the arches. During the monitoring period, the railway bridge has been retrofitted, resulting in data for two distinct states of the structure. Particular attention is paid to removing the effects of environmental conditions, such as temperature, which affect the modal characteristics of the structure and therefore may lead to false-positive or false-negative damage detection. A comparison is made between standard linear regression and robust principal component analysis (PCA), two black-box modeling techniques which are adopted to remove natural frequency variations resulting from changes in the environmental conditions. In order to assess the success rate of these techniques, a receiver operating characteristic (ROC) curve analysis is performed, considering the actual retrofit as well as a number of more subtle structural changes, which are modeled using a detailed finite element model of the structure. The state transition can be observed for the actual retrofit as well as for smaller structural modifications that result in relatively small natural frequency shifts.

Keywords: structural health monitoring, system identification, damage detection, environmental variability, railway bridge KW51

1. Introduction

Structural health monitoring (SHM) relies on the observation that structural changes, for example due to damage, affect the behavior of a structure, i.e. its deflection, vibrations, etc. Monitoring the response of a structure during its lifetime has therefore been proposed as a tool to detect damage at the earliest possible stage. According to Rytter [1], SHM is composed of the following four identification stages:

[⋆] *Postprint submitted to Mechanical Systems and Signal Processing*

[⋆⋆] *Published version:* K. Maes, L. Van Meerbeeck, E. P. B. Reynders, and G. Lombaert. Validation of vibration-based structural health monitoring on retrofitted railway bridge KW51. *Mechanical Systems and Signal Processing*, 165:108380, 2022.
`https://doi.org/10.1016/j.ymssp.2021.108380`

[*] Corresponding author. Tel.: +32 (0) 16 32 25 73.
Email address: `kristof.maes@kuleuven.be` (K. Maes)

1. Detection – Is there any damage present in the structure?
2. Localization – What is the location of the damage?
3. Assessment – What is the type and severity of the damage?
4. Prediction – What is the remaining service life of the structure?

Damage detection (stage 1) very often makes use of dynamic response data, for example accelerations, to extract damage-sensitive features, for example natural frequencies. Those features are used to detect novelties that could possibly be attributed to damage. Once damage has been detected, the data are adopted to localize and assess the damage (stages 2 and 3). This step often involves the use of a model, for example when model updating is applied [2]. In the last and most challenging stage, damage progression models are adopted to predict the remaining lifetime of the structure. This paper focuses on damage detection (stage 1).

Many damage-sensitive features proposed in the literature, including natural frequencies, are influenced by environmental and operational conditions. Environmental conditions include for example temperature, relative humidity, and wind speed. Operational conditions involve altering of loading conditions, for example the traffic crossing a bridge, or control operations for the case of a wind turbine. The changes in damage-sensitive features due to variations in environmental and operational conditions are often of the same order of magnitude as those resulting from damage, which may lead to false-positive or false-negative damage detection. It is the differentiation between damage and the conditions determined by the environment and operation of the structure which poses a big challenge for damage detection [3, 4].

The use of physical models to predict and remove the influence of environmental and operational conditions on the dynamic response of the structure is in most cases very difficult. This is mainly due to the complex and often non-linear relation between the environmental and operational parameters and the mechanical properties of the building materials and the support conditions [5, 6]. As such, black-box models are commonly used in damage detection as a valuable alternative for physical models. Black-box models are purely data-driven and are obtained by means of system identification, also referred to as machine learning. The structure of the model and its parameters are not based on physical laws. Two types of data-driven models exist, which are referred to as input-output models and output-only models. Inputs consist of the environmental and operational variables that drive the so-called natural variation in the structural behavior (temperature, relative humidity, ...). Outputs consist of the damage-sensitive features (e.g. natural frequencies). In this contribution, it is assumed that only data obtained in the original (undamaged) state of the structure can be used in the training of the model. This situation is in particular relevant for civil engineering structures which are mostly characterized by a unique design and a large variety of possible damage scenarios.

The simplest approach to determine an input-output model that allows removing the environmental and operational conditions is linear regression [3, 7]. Alternative approaches that do not assume a linear relation between the inputs and the outputs but allow identifying a global non-linear input-output mapping are based on neural networks [8] and support vector machines [4, 9]. Input-output modeling requires that all relevant environmental and operational variables are measured. This is rarely the case. As indicated by Rainieri et al. in [6], the identification of an input-output model becomes very challenging when multiple environmental and operational variables act simultaneously and are potentially correlated. This is why many approaches in the literature aim at compensating the damage-sensitive features by means of so-called output-only modeling, where measurements of the environmental and operational variables are not required. The most popular

approach for output-only modeling relies on linear principal component analysis (PCA) [10, 11], where a linear relationship between the measured features is identified. When the variation of the features does no longer match the identified relationship, this may point towards damage. As an alternative to PCA, Rainieri et al. have recently proposed the use of second-order blind identification (SOBI) [6]. This technique has the advantage over linear PCA that it can provide fundamental insight in the causes of the natural variability. The relationship between damage-sensitive features is often non-linear [4], which poses another challenge for damage detection. In this case, linear PCA and SOBI can no longer be applied. A very promising approach that allows identifying a non-linear output-only model is kernel PCA, that was proposed for application in SHM by Reynders et al. in [4]. Kernel PCA consists of a computationally efficient non-linear version of linear PCA, for which the structure of the non-linear model does not need to be explicitly defined. A recent application of kernel PCA can be found in [12].

The previously discussed and commonly adopted linear PCA technique is characterized by a large sensitivity to outliers in the training data [24]. Robust PCA, a recently developed technique which is widely adopted in the field of image processing, serves as an extension of linear PCA and aims at removing outliers in the training data prior to the actual determination of the principal components describing the variation in the data [23, 24]. The use of robust PCA was introduced into the field of structural health monitoring by Gharibnezhad et al. in [25], where differente types of robust PCA algorithms are compared and discussed. This paper explores one specific type of robust PCA, i.e. robust PCA via sparse plus low-rank (S+LR) [23].

The presence of actual damage in a structure that is being monitored is rare. This complicates the validation of damage detection approaches. A common approach to investigate the detection capabilities of SHM methods in those cases where no actual damage is present is to generate frequency shifts that correspond to realistic damage scenarios, based on a numerical model of the structure [13, 12]. Many of the aforementioned black-box modeling approaches for damage detection have also been applied to experimental data from the Z24 benchmark problem, where long-term monitoring was performed on a highway bridge prior to progressive damage tests [4, 3, 14]. Several approaches, including kernel PCA, succeed in identifying the damage present in this case. The damage induced in some of the progressive damage tests is rather large, however. The detection of more subtle structural changes (in an early stage) remains a challenge. In addition, it is desirable to develop validation cases which consider different types of structures and other types of damage.

This paper investigates the use of natural frequencies to monitor the health of steel bowstring railway bridge KW51 in Leuven, Belgium, shown in Fig. 1. The data used in this study are obtained from an ongoing long-term monitoring campaign on the railway bridge, which includes acceleration measurements on the bridge deck and the arches. The acceleration data have been used to extract the evolution of the natural frequencies over a period of 15 months. In this period, the connections of the diagonals to the bridge deck and the arches were strengthened after the observation of damage. As such, the data set includes data for two different states of the structure, i.e. before and after retrofitting.

Using the aforementioned natural frequency data, a comparison is made between standard linear regression and robust PCA via sparse plus low-rank (S+LR). It is investigated to what extent these techniques enable removing variations in the natural frequency data that are due to changes in environmental conditions, and, as such detecting the actual changes in the structural behavior that are introduced by damage or, in this case, by the retrofit. Furthermore, a calibrated

detailed finite element model of the structure is used to simulate a number of more subtle structural changes. This enables a more extensive comparison of damage detection methods.

The paper is organized as follows. First, Section 2 describes the vibration monitoring of railway bridge KW51 and presents the identified modal characteristics obtained from the monitoring campaign. Next, Section 3 presents the two damage detection methods and their application to the monitoring data. In Section 4, a detailed finite element model of the railway bridge is adopted to simulate structural modifications other than the actual retrofit. It is investigated to what extent the two damage detection methods succeed in detecting the state transition would these structural modifications occur. Finally, Section 5 concludes the work.

A detailed description of the data used in this paper can be found in [15]. The data are available for download at [16].

2. Vibration monitoring of railway bridge KW51

2.1. Introduction

Railway bridge KW51, shown in Fig. 1, is located on railway line L36N between Leuven and Brussels, Belgium, where it enables the crossing of the canal Leuven-Mechelen. The bridge, of the type bowstring, has a length of 115 m and a width of 12.4 m.

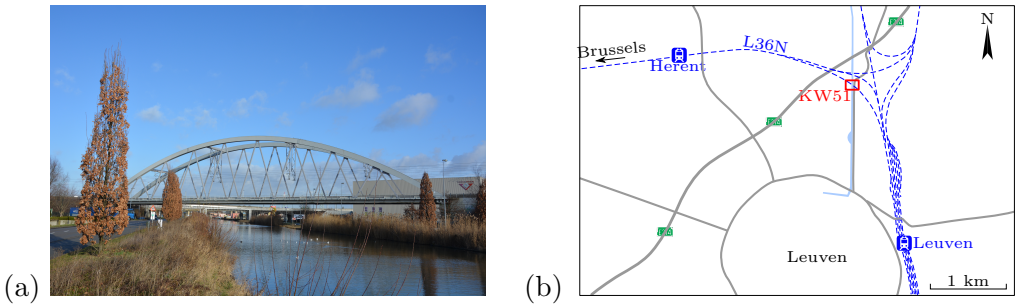

Fig. 1: (a) Railway bridge KW51 in Leuven, Belgium and (b) situation sketch.

The railway bridge consists of two ballasted tracks, referred to as track A (at the north-side) and track B (at the south-side). A maximum speed of 160 km/h is imposed. Both tracks are curved, with a curve radius of 1125 m for track A and 1121 m for track B. The bridge is used by passenger trains and was opened for traffic in 2003.

2.2. Measurement setup

A monitoring system was installed on the bridge in September 2018. The following quantities are measured (see also Fig. 2):

- Six accelerations on the bridge deck in the lateral and vertical direction.
- Six accelerations on the arches in the lateral direction.
- Eight uniaxial strains on the bridge deck along the longitudinal direction.
- Four uniaxial strains on two braces connecting the arches and the bridge deck (2 times 2).
- Two biaxial strains on one rail of tracks A and B (directions $+45°$ and $-45°$ with respect to the neutral axis of the rail).

– Two uniaxial laser-based displacements at the abutments of the bridge side Leuven, along the longitudinal direction.
– The surface temperature and the relative humidity at one point below the bridge deck.

The measurements are performed by means of a National Instruments (NI) data acquisition system, consisting of a NI cDAQ 9178 chassis equipped with two NI 9234 ICP modules for the acceleration measurements, four NI 9237 analog input bridge modules for the strain measurements, and one NI 9219 analog input module for the displacement, temperature, and relative humidity measurements. The data acquisition hardware is installed in an IP66 steel cabinet that is mounted on one of the transverse girders below the bridge deck. In the remainder of this paper, only the acceleration and temperature measurements are considered. The reader is referred to [15] for a detailed description of the monitoring system.

In addition to the measurements at the bridge site, additional environmental data are collected at the Vliet Building from the KU Leuven Building Physics Section. This building is located on the KU Leuven campus, at a distance of about 5 km from the railway bridge. The environmental data are provided on an hourly basis and include air temperature, relative humidity, vapor pressure, solar radiation, rain, wind speed, and wind direction.

Finally, a second independent vibration monitoring system consisting of 80 fiber Bragg grating (FBG) strain measurements is present on the railway bridge [17]. The monitoring covers a shorter time period. These data have not been considered here.

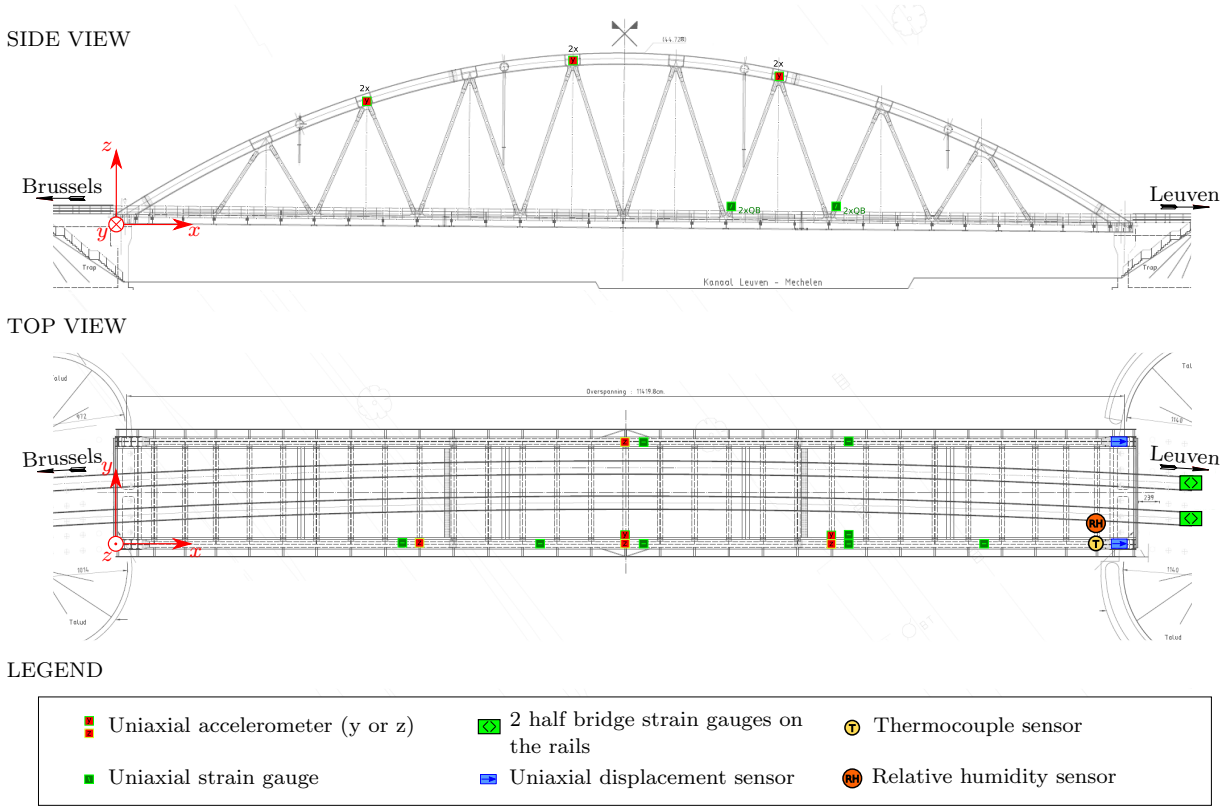

Fig. 2: Overview of the measurement setup. Sensors installed on the bridge deck are only indicated in the top view. Sensors installed on the diagonals and the arches are only indicated in the side view.

2.3. Structural retrofitting

In the period from 15 May to 27 September 2019, the bridge was retrofitted to resolve a construction error that was noticed during inspection. The retrofit consisted of strengthening the connections of the diagonals to the arches and the bridge deck. Fig. 3 shows a picture of these connections before and after the retrofit. For every diagonal, a steel box was welded around the original bolted connection at the intersection with the arches and the bridge deck.

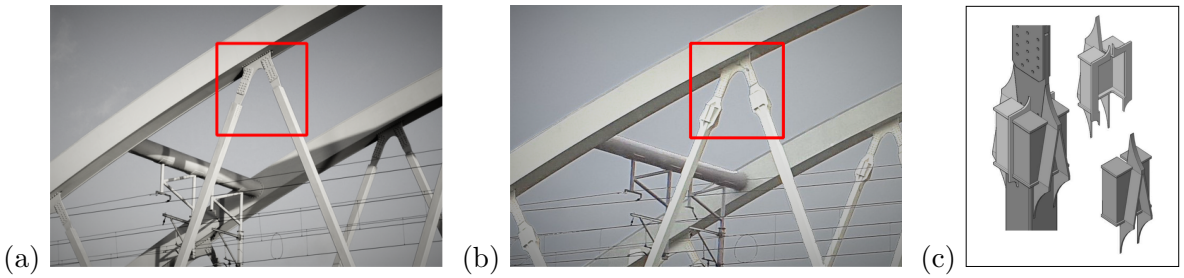

Fig. 3: (a) Original connections between the diagonals and the arches and the bridge deck and (b) connections after retrofitting, and (c) detail of the steel boxes welded around the original bolted connection.

The measurements adopted in this paper have been collected during a period of 7.5 months before the retrofit (02/10/2018 – 15/05/2019), 4.5 months during the retrofit (15/05/2019 – 27/09/2019), and 3.5 months after the retrofit (27/09/2019 – 15/01/2019).

2.4. Identified modal characteristics

The modal characteristics of the structure have been experimentally determined on an hourly basis using automated operational modal analysis (OMA) [15]. In the OMA, the acceleration measurements corresponding to a period of five minutes of ambient vibration (no train passages on the bridge) are processed using the reference-based covariance-driven stochastic subspace identification (SSI-cov/ref) algorithm [18]. An automated interpretation of the stabilization diagram representing the results of the SSI-cov/ref algorithm is performed using the clustering approach proposed in [19]. Finally, an additional clustering step is performed in order to group modes that occur in multiple date sets. The adopted procedure is described in detail in [15].

In total, 14 modes are tracked over time. Note that none of the modes is identified with a success ratio of 100%, i.e. every hour of the monitoring period. This is due to the fact that the excitation of the modes and therefore their practical identifiability changes over time. Table 2.4 provides an overview of the identified modes, including the time-average natural frequency and the success ratio (SR) of the identification before and after the retrofit, as well as a description of the mode shapes. This description is based on the mode shapes obtained from a detailed multi-setup operational modal analysis (OMA) that was performed prior to the monitoring on 16/02/2018, see also [15]. The success ratio SR of the identification changes after the retrofitting. This is especially the case for the lateral bending modes of the arches, which is explained by the fact that the sensor network after retrofitting is extended by six additional accelerometers that were installed on the arches. The acceleration measurements obtained from those sensors are more sensitive to the lateral bending modes of the arches compared to the (earlier installed) measurements on the bridge deck.

	Before retrofit		After retrofit		
No.	f_{id} [Hz]	SR [%]	f_{id} [Hz]	SR [%]	Description
1	0.51	34.3	0.51	94.4	1^{st} lateral mode of the arches (arches in phase)
2	1.24	25.3	1.23	91.7	2^{nd} laterale mode of the arches (arches in phase)
3	1.89	99.0	1.88	92.4	1^{st} lateral mode of the bridge deck
4	2.45	57.6	2.42	90.1	3^{rd} lateral mode of the arches (arches in phase)
5	2.58	95.6	2.57	81.0	1^{st} global vertical mode
6	2.92	99.5	2.98	99.7	2^{nd} global vertical mode
7	3.55	19.7	3.55	88.8	4^{th} lateral mode of the arches (arches in phase)
8	3.88	23.3	3.92	70.6	4^{th} lateral mode of the arches (arches in phase)
9	4.11	99.2	4.08	99.5	1^{st} global torsional mode
10	4.30	66.7	4.39	33.5	3^{rd} global vertical mode
11	4.84	63.1	4.94	68.8	2^{nd} lateral mode of the bridge deck (+ torsion)
12	5.33	79.5	5.44	43.3	4^{th} global vertical mode
13	6.33	98.2	6.42	78.4	5^{th} global vertical mode
14	6.88	60.9	7.06	28.9	3^{rd} lateral mode of the bridge deck (+ torsion)

Table 1: Overview of the tracked modes before and after the retrofitting (f_{id}: time-average identified natural frequency, SR: success ratio of the identification). Time windows where the data acquisition was interrupted have not been included in the calculation of the SR.

Fig. 4 displays the evolution of the identified natural frequencies over time, together with the air temperature at the Vliet Building and the surface temperature below the bridge deck. The influence of the temperature on the identified natural frequencies is clearly visible in the period before the retrofitting. This is especially the case for the frost period in January 2019. During the retrofitting, the natural frequencies change significantly as well. This is mainly due to the presence of scaffolding at both sides of the bridge. For some modes, it was not possible to continue the mode tracking during the retrofitting. In what follows, only the modal characteristics from the period before the retrofit (original state) and the period after the retrofit (modified state) are considered.

It is also noticed that the fluctuations of the bridge surface temperature (Fig. 4b, black) are significantly smaller than the fluctuations of the air temperature (Fig. 4b, blue). This is most likely due to the large thermal inertia of the structure. Finally, note that the bridge surface temperature is not available before 22/11/2018, due to a data acquisition problem.

Fig. 5 shows the identified natural frequencies as a function of the measured bridge surface temperature. A distinction is made between data corresponding to the period before the retrofit (black) and after the retrofit (gray). The following observations are made:

- Before the retrofit, a bilinear relation between the identified natural frequencies and the bridge temperature is observed for most modes, with a knee point around 0°C. Both for positive and negative temperatures, the natural frequencies tend to decrease with increasing temperature. A similar bilinear relation has been previously observed for the Z24 bridge in [3], where it was found that this bilinearity is due to a significant increase of the stiffness of the asphalt layer on the bridge deck during frost periods. For the KW51 railway bridge, it was verified by means of finite element model updating that the significant increase of the natural frequencies for temperatures below 0°C is likely due to frost of the porphyry ballast layer supporting the tracks.
- For modes 6, 7, 8, 10, 11, 12, 13, and 14, the retrofit results in an increase of the natural frequency. Taking into account the fact that the retrofit results in relatively small mass and

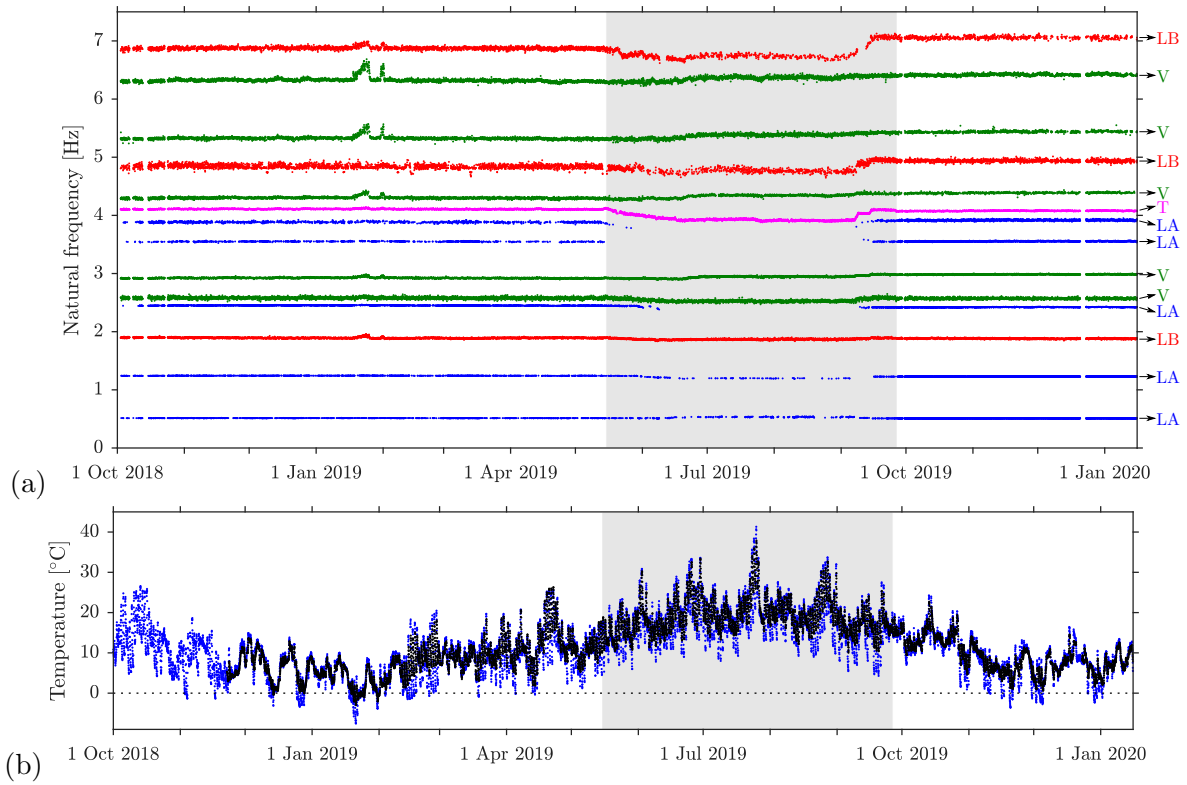

Fig. 4: (a) Evolution of the identified natural frequencies over time (blue: lateral modes of the arches (LA), red: lateral modes of the bridge deck (LB), green: global vertical modes (V), and magenta: global torsional modes (T)) and (b) time history of the surface temperature below the bridge deck (black) and the air temperature at the Vliet building (blue). The period during which the retrofit is performed is indicated with a gray shading.

stiffness changes, it can be reasonably assumed that the increase in those natural frequencies is attributed to the increased stiffness of the connections between the diagonals, the bridge deck, and the arches [20]. This is also verified by introducing the stiffness increments corresponding to the retrofit in the detailed finite element model of the railway bridge described in Section 4.1.

– For modes 1, 2, 3, 4, 5, and 9, the retrofit results in a decrease of the natural frequency. Similarly, the decrease in those natural frequencies is attributed to the additional mass of the steel boxes, which is found to play an important role for the three lowest lateral bending modes of the arches (modes 1, 2, and 4). This is also verified by introducing the local mass increments corresponding to the retrofit in the detailed finite element model of the railway bridge described in Section 4.1.

– For the majority of the identified modes, two separate clusters are obtained, where one is corresponding to the original state, before retrofitting, and the other is corresponding to the modified state, after retrofitting. For some modes, the shift in natural frequency due to the retrofit is larger than the observed variation of the natural frequencies due to changes in the environment. This is for example the case for mode 4. One could argue that a removal of environmental influences on the natural frequencies for those modes is not necessary to enable damage/novelty detection. However, when aiming at the detection of damage in an early stage, removing these influences is essential, as demonstrated in the following sections.

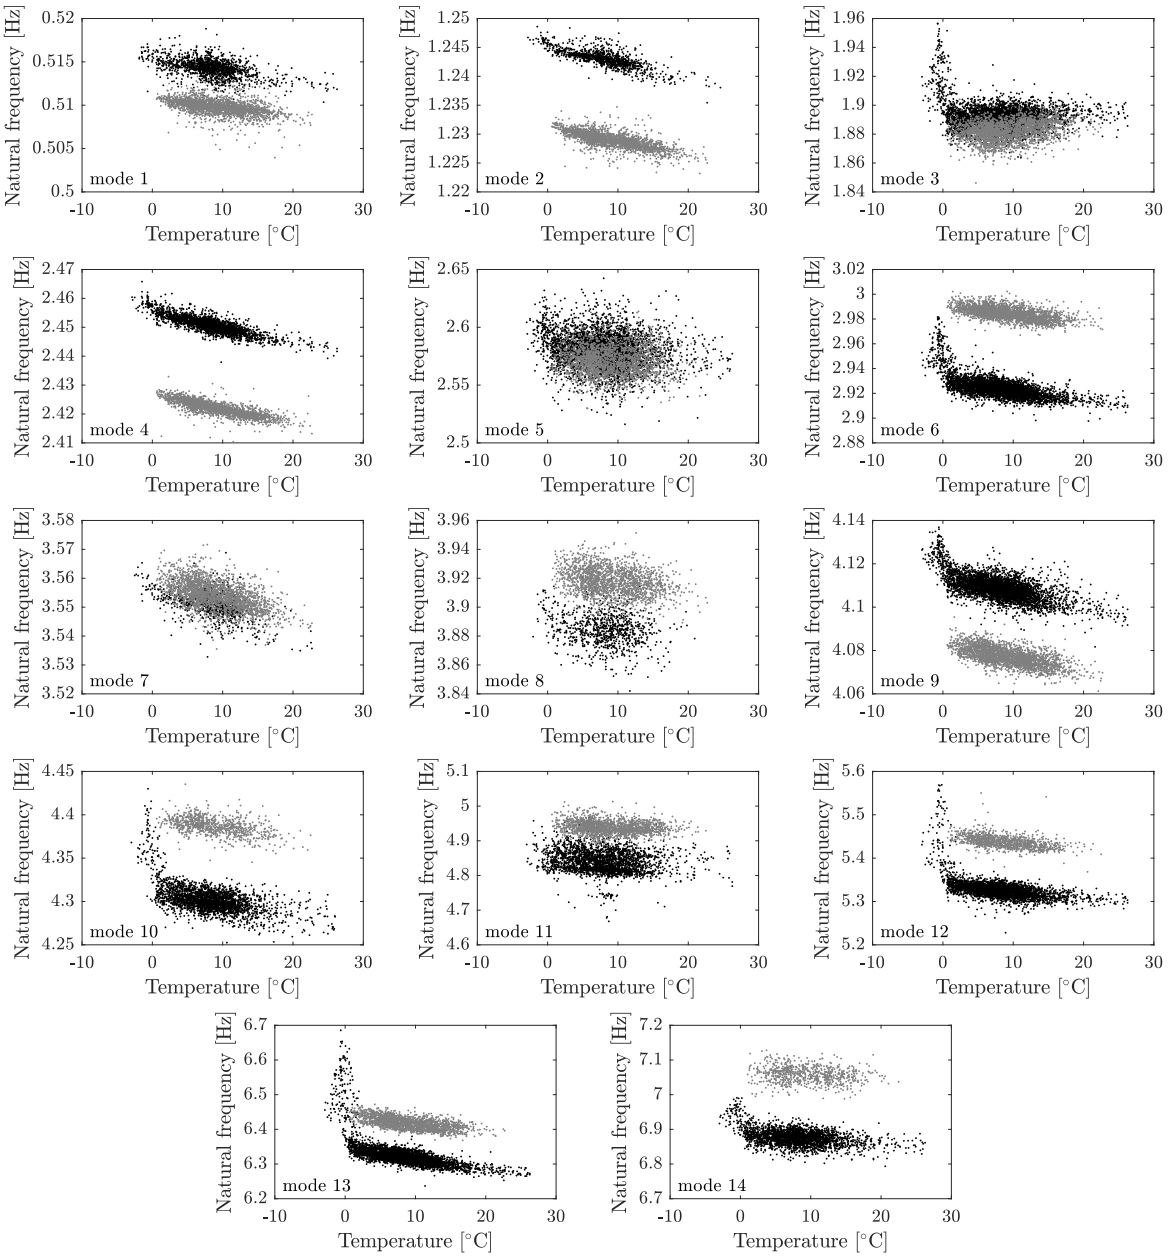

Fig. 5: Evolution of the identified natural frequencies with the bridge temperature before the retrofit (black) and after the retrofit (gray).

Apart from the clear variation of the natural frequencies with temperature, it is expected that other environmental variables play a role. This is further addressed in Section 3.2. Nonetheless, it cannot be excluded that part of the variation in the natural frequencies is attributed to uncertainty in the system identification. This uncertainty mainly stems from the fact that some modes are not properly excited and/or do not contribute significantly to the measured bridge acceleration response. As mentioned previously, additional sensors have been installed on the arches during the retrofit. This results in a clearly lower variability of the natural frequency estimates for some modes, including for example mode 11 (Fig. 5), indicating that those modes were not significantly contributing to the response when only measuring accelerations on the bridge deck.

Finally, it is also important to bear in mind that the excitation of the railway bridge is generally twofold [15]. If the period considered in the system identification directly follows a train passage, which is often the case during the day, the measured accelerations mainly consist of a free decay, where the majority of the modes under consideration has a clear contribution to the response. If, on the other hand, no train passages precede the selected time period, for example during the night, the bridge vibration is mainly caused by wind loading and vehicle passages underneath the bridge. In this case, the practical identifiability largely depends on level and the properties (e.g. wind direction) of the ambient excitation. Some modes may not be properly excited. If this is the case, these modes are not identified or the modal characteristic estimates are prone to larger uncertainty.

3. Damage detection

This section compares two techniques to remove the environmental influences on the identified natural frequencies, aiming at improved damage detection; (1) linear regression and (2) robust PCA via sparse plus low-rank (S+LR), a robust version of classical linear PCA. For both techniques, a brief recapitulation of the method is provided, before they are applied to the data described in Section 2.4.

3.1. Linear regression

Linear regression consists of the identification of a linear input-output mapping. For a given mode j, the linear regression model is described by the following equation:

$$m_{j,k} = \mathbf{W}_j \mathbf{u}_k + b_j + e_{j,k} \tag{1}$$

The output of the model is $m_{j,k} \in \mathbb{R}$, consisting in this case of the identified natural frequency corresponding to mode j for a time window k, shown in Fig. 4a. The input of the model is a vector $\mathbf{u}_k \in \mathbb{R}^{n_\mathrm{u}}$ of n_u measured environmental variables, that are typically averaged over the corresponding time window k. In this case, only a single input is considered ($n_\mathrm{u} = 1$), consisting of the steel surface temperature measured below the bridge deck, shown in Fig. 4b (black). The scalar $e_{j,k} \in \mathbb{R}$ is the time-dependent error term that accounts for the misfit between the model predictions and the observed behavior of the structure.

The model training consists of estimating the coefficient matrix $\mathbf{W}_j \in \mathbb{R}^{1 \times n_\mathrm{u}}$ and the bias term $b_j \in \mathbb{R}$. The estimation is based on a set of measured and hence known outputs $m_{j,k}$ and inputs $\mathbf{u}_k$ for n_t (not necessarily consecutive) training points $k = 1, 2, \ldots, n_\mathrm{t}$ obtained from the structure in its original (undamaged) state. The estimation is performed for every mode separately in a least-squares sense:

$$\begin{bmatrix} \hat{\mathbf{W}}_j & \hat{b}_j \end{bmatrix} = \begin{bmatrix} m_{j,1} & m_{j,2} & \ldots & m_{j,n_\mathrm{t}} \end{bmatrix} \begin{bmatrix} \bar{\mathbf{u}}_1 & \bar{\mathbf{u}}_2 & \ldots & \bar{\mathbf{u}}_{n_\mathrm{t}} \end{bmatrix}^\dagger \tag{2}$$

where $\bar{\mathbf{u}}_k = \begin{bmatrix} \mathbf{u}_k^\mathrm{T} & 1 \end{bmatrix}^\mathrm{T}$ and $\dagger$ denotes the Moore-Penrose pseudo-inverse.

Once $\hat{\mathbf{W}}_j$ and $\hat{b}_j$ are estimated, the model prediction $\hat{m}_{j,k}$, which consists of an estimate of natural frequency j given the bridge surface temperature at time window k, can be predicted as follows:

$$\hat{m}_{j,k} = \hat{\mathbf{W}}_j \mathbf{u}_k + \hat{b}_j \tag{3}$$

As previously discussed in Section 2.4, a bilinear relation between the bridge surface temperature and the natural frequencies is observed, with a knee point located near 0°C. The assumed linear regression in Eq. (1) cannot capture such bilinear relation. In order to overcome this issue, one could train two independent models, one for negative temperatures and one for positive temperatures. This is illustrated for mode 13 in Fig. 6. The curve fitting for negative temperatures is complicated, however, by (1) the large variation of the identified natural frequencies corresponding to negative temperatures, potentially due to other environmental factors that play an important role, (2) the very limited range of negative temperatures present in the data, and (3) the small number of data points corresponding to negative temperatures. As such, it was decided to consider only positive temperatures. In addition, the model is limited to temperatures above 2°C (instead of 0°C). This margin of 2°C is imposed to avoid the large variation in natural frequencies observed below 2°C, which leads to a bias in the linear regression, also illustrated in Fig. 6. The restriction to temperatures above 2°C comes with the drawback that the damage detection is inevitably interrupted during frost periods. For the structure under consideration, however, the number of days per year where the bridge deck temperature drops below the cutoff limit of 2°C is typically small (see also Fig. 5b).

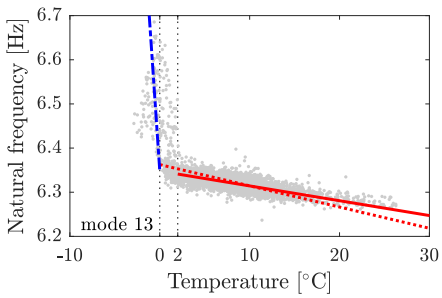

Fig. 6: Illustration of the linear regression for mode 13, for temperatures below 0°C (dash-dotted blue line), temperatures above 0°C (dotted red line), and temperatures above 2°C (solid red line). The curve fit for negative temperatures (solid blue line) is manually tuned to obtain a proper visual agreement with the data.

For every mode, the linear regression model is trained using data obtained before the retrofit. The time series of natural frequency data available for this period is subdivided in periods of one week (seven consecutive days), where data obtained for odd weeks (week 1, week 3, ...) are used to train the model and data obtained for even weeks (week 2, week 4, ...) are used for validation. By alternating short blocks of training and validation data, a large variation in environmental conditions is obtained for both subsets.

After the model training, the model prediction $\hat{m}_{j,k}$ in Eq. (3) can be used to calculate the model misfit $\hat{e}_{j,k}$ for any time window k, quantifying to what extent the identified natural frequency meets the assumed linear input-output mapping:

$$\hat{e}_{j,k} = m_{j,k} - \hat{m}_{j,k} = m_{j,k} - \hat{\mathbf{W}}_j \mathbf{u}_k - \hat{b}_j \tag{4}$$

Table 2 presents the coefficients of determination R_j^2 for all modes j, which are calculated as follows:

$$R_j^2 = 1 - \frac{\sum_{k=1}^{n_t} \hat{e}_{j,k}}{\sum_{k=1}^{n_t} \left(m_{j,k} - \frac{1}{n_t} \sum_{l=1}^{n_t} m_{j,l} \right)} \tag{5}$$

These coefficients explain to what extent the observed variation in the natural frequencies can be attributed to variations in the bridge deck temperature. Although a clear relation between the natural frequencies and the bridge deck temperature was observed in Fig. 5, very low values of R_j^2 are obtained, indicating that a major part of the natural frequency variations is resulting from other factors, including other environmental influences (e.g. solar radiation), the excitation level, and identification errors.

Mode No. j	1	2	3	4	5	6	7	8	9	10	11	12	13	14
R_j^2	0.17	0.49	0.00	0.62	0.01	0.31	0.14	0.01	0.31	0.20	0.05	0.24	0.50	0.06

Table 2: Coefficient of determination R_j^2 for the linear regression model composed for each of the different modes j and considering only temperatures above 2°C.

Fig. 7 shows the evolution of the model misfit $\hat{e}_{j,k}$ over time, for each of the 14 modes separately. The results shown for temperatures below 2°C (in gray) are obtained by extending the linear regression model towards negative temperatures and are only shown for the sake of completeness. They are not considered in the actual damage detection discussed hereafter.

Before the retrofit, the misfit $\hat{e}_{j,k}$ is mostly contained in the $[-3\sigma_j, +3\sigma_j]$ interval (also indicated in Fig. 7), both for the training and validation data, where σ_j is the standard deviation of the misfit $\hat{e}_{j,k}$ in the training phase. After the retrofit, a clear shift in the misfit is observed for most modes, pointing towards a change in the dynamic behavior of the structure. As previously mentioned in Section 2.4, the retrofitting results in an increased natural frequency and consequently a positive misfit for some modes, while leading to a decreased natural frequency and a negative misfit for other modes. For some modes, the misfit $\hat{e}_{j,k}$ clearly exceeds the $[-3\sigma_j, 3\sigma_j]$ interval after the retrofit, resulting in a pronounced gap between the misfit for the original state and the modified state. This is for example the case for modes 2, 4, and 6. These modes are characterized by a large sensitivity to the structural modification and/or a small uncertainty on the identified natural frequencies.

In practice, the influence of potential damage on each of the modes is not known in advance. It can therefore not be predicted which modes will be most sensitive and therefore most suited to detect actual damage in an early stage. Rather than monitoring the model misfit for each mode individually, it is more convenient to consider the following overall normalized misfit $\hat{\epsilon}_k$, which takes into account the contribution of the different (n_m) modes:

$$\hat{\epsilon}_k = \sqrt{\sum_{j=1}^{n_\mathrm{m}} \left(\frac{\hat{e}_{j,k}}{\sigma_j} \right)^2} \tag{6}$$

Under the assumption that the misfits $\hat{e}_{j,k}$ for the different modes are uncorrelated, this normalized misfit $\hat{\epsilon}_k$ corresponds to the Mahalanobis distance [21]. In addition, assuming that the misfits $\hat{e}_{j,k}$ for the original state of the structure follow a normal distribution with zero mean and standard deviation σ_j, the squared normalized misfit $\hat{\epsilon}_k^2$ follows a chi-square distribution. This allows calculating the critical value of $\hat{\epsilon}_k$ that is only exceeded with a probability of 0.003, similar to the $[-3\sigma_j, 3\sigma_j]$ interval for the individual misfits $\hat{e}_{j,k}$. This value is referred to hereafter as $c_{\chi^2}^{0.997}$, with $\mathrm{P}(\hat{\epsilon}_k \leq c_{\chi^2}^{0.997}) = 0.997$.

Fig. 8 shows the evolution of the normalized misfit $\hat{\epsilon}_k$ over time. Only the contribution of modes 3, 4, 5, 6, 9, 11, and 13 is accounted for in the calculation of this misfit ($n_\mathrm{m} = 7$), where an additional

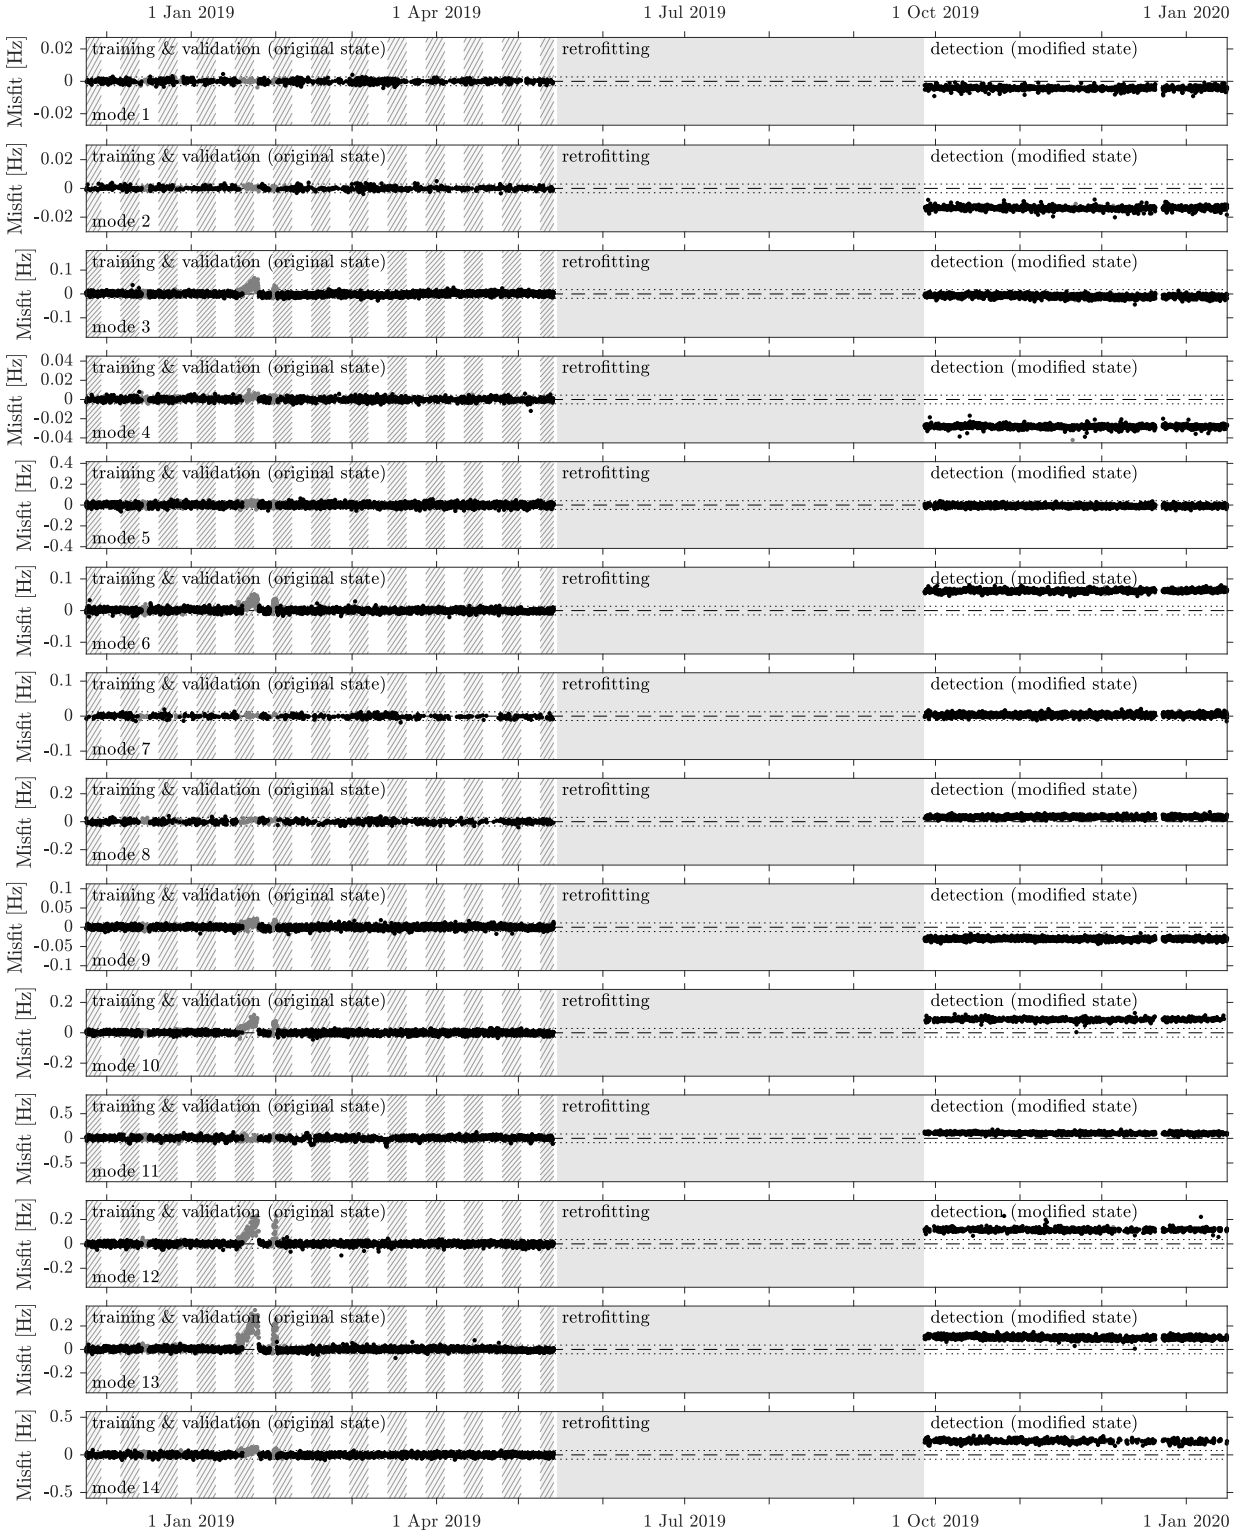

Fig. 7: Evolution of the linear regression model misfit $\hat{e}_{j,k}$ for modes 1 to 14 over time. Results corresponding to a steel surface temperature above 2°C are indicated in black, those corresponding to a temperature below 2°C in gray. The horizontal black dashed line indicates the zero misfit level. The data blocks composing the training period are hatched. The horizontal black dotted lines indicate the $+3\sigma_j$ and $-3\sigma_j$ misfit levels.

linear interpolation of the identified natural frequencies is adopted to fill small gaps when some data points are missing (see also Section 2.4). For the other modes (1, 2, 7, 8, 10, 12, and 14), the time evolution of the identified natural frequencies before and/or after the retrofit is characterized by clear gaps corresponding to time periods where the modes have not been successfully identified (Fig. 7). These modes have been excluded.

During the training and validation phase, the normalized misfit is contained within the [0, $c_{\chi^2}^{0.997}$] interval for the majority of the data points. After the retrofit, a clear shift of the misfit is observed, which again indicates that the structure has changed. In order to further improve the detection capacity, the misfits $\hat{e}_{j,k}$ are averaged over a period of 24 hours before calculation of the normalized misfit $\hat{\epsilon}_k$ in Eq. (6). The standard deviation σ_j occurring in Eq. (6) is re-computed from the 24 hour averaged misfits $\hat{e}_{j,k}$ in the training phase. The normalized misfit after 24 hour averaging is shown in Fig. 9.

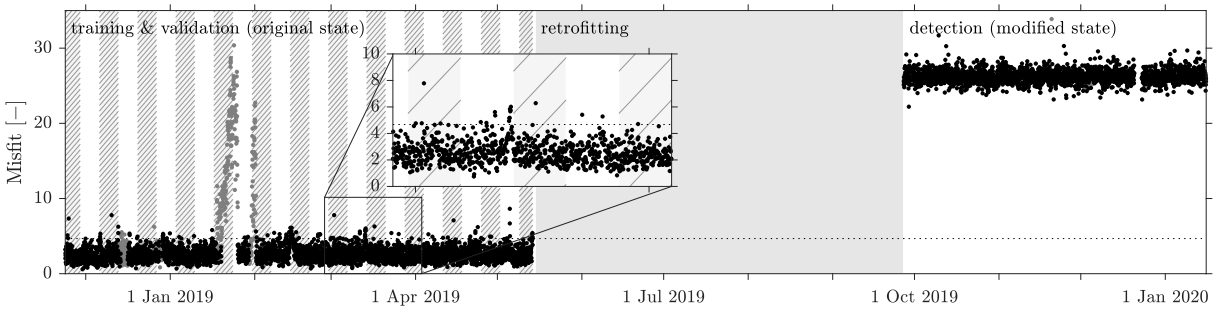

Fig. 8: Evolution of the normalized linear regression model misfit $\hat{\epsilon}_k$ over time. Results corresponding to a steel surface temperature above 2°C are indicated in black, those corresponding to a temperature below 2°C in gray. The data blocks composing the training period are hatched. The horizontal black dotted line indicates the $c_{\chi^2}^{0.997}$ misfit level ($c_{\chi^2}^{0.997} = 4.67$).

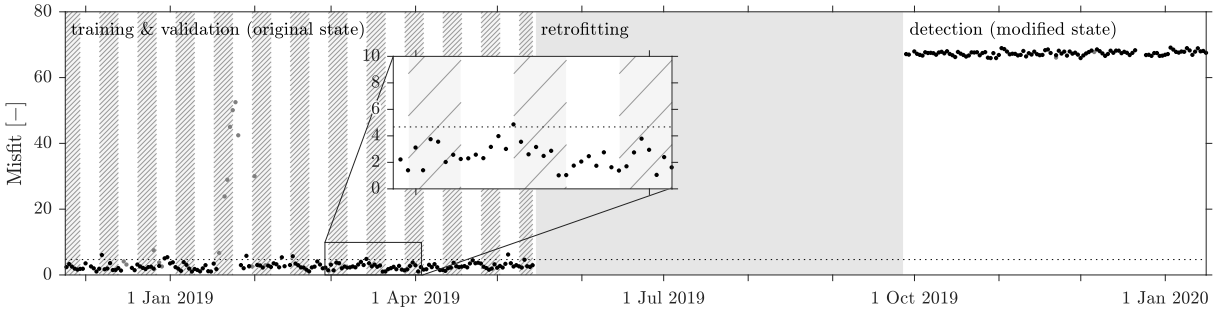

Fig. 9: Evolution of the normalized 24 hour averaged linear regression model misfit $\hat{\epsilon}_k$ over time. Results corresponding to a steel surface temperature above 2°C are indicated in black, those corresponding to a temperature below 2°C in gray. The data blocks composing the training period are hatched. The horizontal black dotted line indicates the $c_{\chi^2}^{0.997}$ misfit level ($c_{\chi^2}^{0.997} = 4.67$).

Due to the significant reduction in noise level achieved by the additional averaging, the shift in the misfit introduced by the retrofitting is even more pronounced with respect to the case where no averaging is applied (Fig. 8). Note, however, that the averaging introduces a delay of 24 hours in the damage detection. If a fast detection is required, it is advised to monitor both the original hourly-based misfit (Fig. 8) and the 24 hours averaged misfit (Fig. 9).

In the considered case, the model misfit $\hat{\epsilon}_k$ enables a clear distinction between the original and the modified state of the structure. Due to the relatively large influence of the structural retrofit

on the considered natural frequencies, the choice of a suitable threshold to distinguish between both states is rather straightforward. Indeed, any threshold value larger than 8.6 would result in a perfect distinction between the original and the modified state, even without additional averaging of the data (Figs. 8 and 9). When aiming at the identification of actual damage in an early stage, however, the selection of a proper threshold value that limits the number of false-positives and false-negatives is crucial. This is further addressed in Section 4.4.

To conclude this subsection, the main advantages and disadvantages of a linear regression model are summarized. A linear regression model has a limited complexity. This comes with the advantage that it is very straightforward to use and easy to implement. On the other hand, it comes with the (often restricting) assumptions that (1) the output is linearly related to the inputs, and (2) all inputs that drive the variation of the output are measured. Although the bridge surface temperature has a clear influence on the natural frequencies of railway bridge KW51, other environmental variables seem to have a significant influence as well. This issue is further addressed in Section 3.2.

3.2. Robust PCA

A principal component analysis (PCA) model describes the variation in the data in terms of its so-called principal components, without requiring knowledge on the inputs that drive this variation. The resulting model is generally described by:

$$\mathbf{m}_k = \mathbf{\Psi}\mathbf{x}_k + \mathbf{b} + \mathbf{e}_k \tag{7}$$

where the vector $\mathbf{m}_k = \begin{bmatrix} m_{1,k} & \dots & m_{n_\mathrm{m},k} \end{bmatrix}^\mathrm{T}$ contains the output for a given time window k, which in this case is the set of n_m identified natural frequencies. The matrix $\mathbf{\Psi} \in \mathbb{R}^{n_\mathrm{m} \times n_\mathrm{c}}$ contains n_c principal components that describe the variation in the data, $\mathbf{x}_k \in \mathbb{R}^{n_\mathrm{c} \times 1}$ contains the contribution of each principal component, also called score, for a given time window k, $\mathbf{b} \in \mathbb{R}^{n_\mathrm{m} \times 1}$ is the bias vector, and $\mathbf{e}_k = \begin{bmatrix} e_{1,k} & \dots & e_{n_\mathrm{m},k} \end{bmatrix}^\mathrm{T}$ is the time-dependent error term that accounts for the misfit between the model predictions and the observed behavior of the structure.

The principal component matrix $\mathbf{\Psi}$ and the bias term $\mathbf{b}$ are determined from a set of measured and hence known outputs $\mathbf{m}_k$ for n_t (not necessarily consecutive) training points $k = 1, 2, \dots, n_\mathrm{t}$ obtained from the structure is in its original (generally undamaged) state. The bias vector $\mathbf{b}$ is estimated as follows:

$$\hat{\mathbf{b}} = \frac{1}{n_\mathrm{t}} \sum_{k=1}^{n_\mathrm{t}} \mathbf{m}_k \tag{8}$$

In classical linear PCA, the principal component matrix $\mathbf{\Psi}$ is typically derived from an eigenvalue decomposition of the sample covariance matrix of the centered training data ($\bar{\mathbf{m}}_k = \mathbf{m}_k - \hat{\mathbf{b}}$). This process is known to be sensitive to outliers in the training data, leading to a misalignment of the principal directions [25]. Robust PCA aims at removing or diminishing the effect of outliers in the extraction of the model. Prior to the determination of the matrix $\mathbf{\Psi}$, the training data are "cleaned" in a separate preprocessing step, as explained next.

In the particular case of robust PCA via sparse plus low-rank (S+LR), presented in [23, 24], the training data matrix $\mathbf{M} = [\bar{\mathbf{m}}_1\ \bar{\mathbf{m}}_2\ \cdots\ \bar{\mathbf{m}}_{n_\mathrm{t}}] \in \mathbb{R}^{n_\mathrm{m} \times n_\mathrm{t}}$, in this case containing the centered

natural frequencies for n_m modes on n_t time steps, is decomposed into a structured low-rank matrix $\mathbf{L} \in \mathbb{R}^{n_\mathrm{m} \times n_\mathrm{t}}$ and a sparse matrix $\mathbf{S} \in \mathbb{R}^{n_\mathrm{m} \times n_\mathrm{t}}$ containing outliers and corrupt data:

$$\mathbf{M} = \mathbf{L} + \mathbf{S} \tag{9}$$

It is assumed that the principal components $\mathbf{\Psi}$ of the resulting matrix $\mathbf{L}$ are robust to the outliers and corrupt data in $\mathbf{S}$. In general, the goal is to determine $\mathbf{L}$ and $\mathbf{S}$ as the solution of the following optimization problem:

$$\min_{\mathbf{L},\mathbf{S}} \ \mathrm{rank}(\mathbf{L}) + ||\mathbf{S}||_0 \ \text{ subject to } \ \mathbf{L} + \mathbf{S} = \mathbf{M} \tag{10}$$

Since optimization problem (10) is non-convex, matrices $\mathbf{L}$ and $\mathbf{S}$ are in practice typically determined using the following convex relaxation of (10):

$$\min_{\mathbf{L},\mathbf{S}} \ ||\mathbf{L}||_* + \lambda||\mathbf{S}||_1 \ \text{ subject to } \ \mathbf{L} + \mathbf{S} = \mathbf{M} \tag{11}$$

where $||\cdot||_*$ denotes the nuclear norm, i.e. the sum of singular values, serving as a proxy for the matrix rank [24]. The scalar λ is a regularization parameter that determines the weight assigned to the rank of the matrix $\mathbf{L}$ and the sparsity of the outlier matrix $\mathbf{S}$. The convex optimization problem (11) is solved using the iterative augmented Lagrange multiplier (ALM) algorithm described in [24], where the Matlab code for this algorithm is also provided.

Upon deriving the low-rank matrix $\mathbf{L}$, its principal components $\mathbf{\Psi}$ are calculated by means of an eigenvalue decomposition of the sample covariance matrix of the centered and cleaned training data:

$$\frac{1}{n_\mathrm{t}}\mathbf{L}\mathbf{L}^\mathrm{T} = \begin{bmatrix}\mathbf{U}_1 & \mathbf{U}_2\end{bmatrix} \begin{bmatrix}\mathbf{\Sigma}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_2\end{bmatrix} \begin{bmatrix}\mathbf{U}_1^\mathrm{T} \\ \mathbf{U}_2^\mathrm{T}\end{bmatrix} \tag{12}$$

where $\mathbf{\Sigma}_1 \in \mathbb{R}^{n_\mathrm{c} \times n_\mathrm{c}}$ contains the n_c largest eigenvalues on its diagonal, with $\mathbf{U}_1$ containing as columns the corresponding eigenvectors, which correspond in this case to the so-called principal directions of the data. Under the assumption that only the first n_c out of the n_m principal components contribute significantly to the data ($n_\mathrm{c} \leq n_\mathrm{m}$), the eigenvector matrix $\mathbf{U}_1$ serves as an estimate of the principal component matrix $\mathbf{\Psi}$, i.e. $\hat{\mathbf{\Psi}} = \mathbf{U}_1$. The selection of the number of retained principal components n_c is essential in the development of a model aiming at damage detection and is further addressed below.

Once $\hat{\mathbf{b}}$ and $\hat{\mathbf{\Psi}}$ are obtained from the training data, the model prediction error $\hat{\mathbf{e}}_k$ for a given time window k can be estimated as follows:

$$\hat{\mathbf{e}}_k = \left(\mathbf{I}_{n_\mathrm{m}} - \hat{\mathbf{\Psi}}\hat{\mathbf{\Psi}}^\dagger\right)\left(\mathbf{m}_k - \hat{\mathbf{b}}\right) = \mathbf{U}_2\mathbf{U}_2^\mathrm{T}\left(\mathbf{m}_k - \hat{\mathbf{b}}\right) \tag{13}$$

Note that the model prediction error $\hat{\mathbf{e}}_k$ is calculated from the original data $\mathbf{m}_k$ without outlier removal by means of robust PCA. The principal directions, contained in the matrix $\hat{\mathbf{\Psi}}$, however, are obtained from the cleaned data, as explain previously. Although the data contained in $\mathbf{m}_k$ may contain outliers, the preprocessing step by means of solving the optimization problem in Eq. (11) is not repeated in the actual detection stage, since the outlier removal may reduce detection capacity. Indeed, when the response of the structure changes due to damage, the resulting data points may be classified as outliers when applying robust PCA, which should therefore be avoided.

Prior to the discussion of the results of the robust PCA method, the relation between the variation in the different natural frequencies is studied. As an example, Fig. 10 shows the relation between the identified natural frequencies for modes 6 and 13, before the retrofitting. As in Section 3.1, a distinction is here made between temperatures below 2°C (gray) and temperatures above 2°C (black). At temperatures above 2°C, a roughly linear trend is observed between the natural frequencies of of both modes. This is in line with the near to linear dependence on temperature that both natural frequencies showed (see Section 3.1). Although the range spanned by the natural frequencies below 2°C is considerably smaller, the trend here also seems roughly linear. In line with what was considered for the linear regression model in Section 3.1, two independent PCA models could be trained, for example using the methodology proposed in [22]. For the same reasons as those mentioned in Section 3.1, however, the model is here limited to temperatures above 2°C. In addition, only modes 3, 4, 5, 6, 9, 11, and 13 are accounted for in the derivation of the robust PCA model ($n_{\mathrm{m}} = 7$), where linear interpolation of the identified natural frequencies is used to generate missing data points (see also Section 3.1). The remaining modes have been omitted, due to the presence of clear gaps in the time history of the identified natural frequencies.

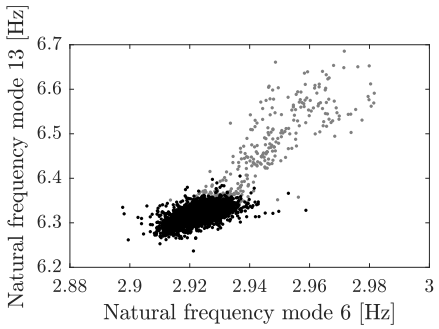

Fig. 10: Observed relation between the identified natural frequency of modes 6 and 13 before the retrofit. Data points corresponding to temperatures below 2°C are indicated in gray. Data points corresponding to temperatures above 2°C are indicated in black.

At this point, one could consider the estimation of a simple linear regression model in Eq. (1) to describe the relation between the identified natural frequencies in Fig. 10 for temperatures above 2°C, rather than developing a PCA model. In fact, such a linear regression model would be similar to a PCA model with only one principal component ($n_{\mathrm{c}} = 1$). Although a PCA model with more than one principal component ($n_{\mathrm{c}} \geq 1$) still assumes a linear relation between the observed natural frequencies, it accounts for multiple linear trends represented in the data, which are described by the principal components. This may be important when the dependence on temperature (or other environmental or operational factors) stems from different physical phenomena. For example, one of the components could represent the variation of the natural frequencies with the bridge deck temperature, while other components could for example represent the variation with the temperature of the arches, which may affect the natural frequencies in a different way. Under the assumption of a sufficient amount of training data, a PCA model (linear PCA or robust PCA) with multiple components will therefore be more effective in modeling the natural variation in the output data and therefore enable more accurate damage detection, as illustrated next.

In order to study the outlier removal by means of robust PCA, Fig. 11 presents a comparison of the original natural frequency training data ($\mathbf{m}_k$) and the low-rank representation of the data resulting from robust PCA ($\mathbf{l}_k + \hat{\mathbf{b}}_k$, with $\mathbf{l}_k$ the columns of the low-rank matrix $\mathbf{L}$), when plotted as a function of the bridge deck temperature. The regularization parameter λ occurring in Eq. (11)

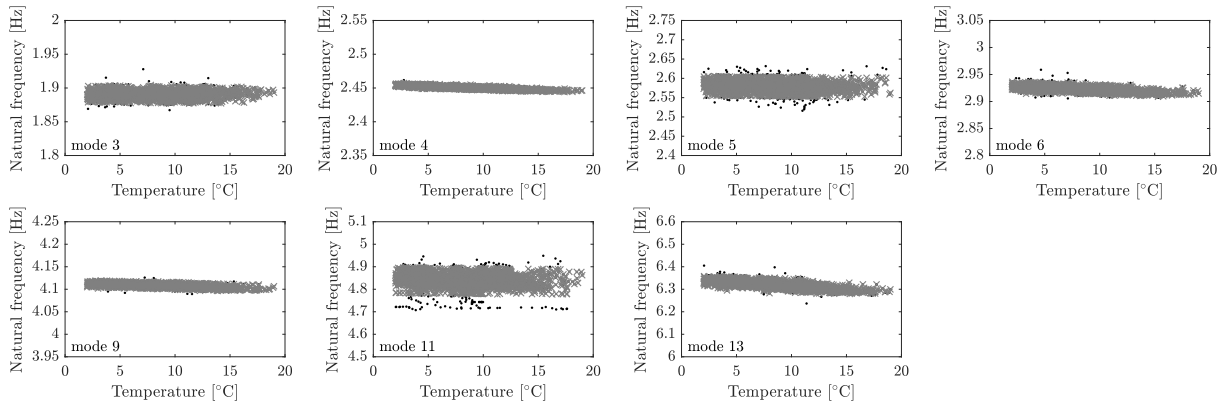

Fig. 11: Comparison of the original natural frequency training data (matrix $\mathbf{M}$, black dots) and the low-rank representation of the data resulting from robust PCA (matrix $\mathbf{L}$, gray $\times$-markers), plotted as a function of the bridge deck temperature. Both for matrices $\mathbf{M}$ and $\mathbf{L}$, the offset contained in the vector $\hat{\mathbf{b}}$ is added before plotting, in order to facilitate interpretation.

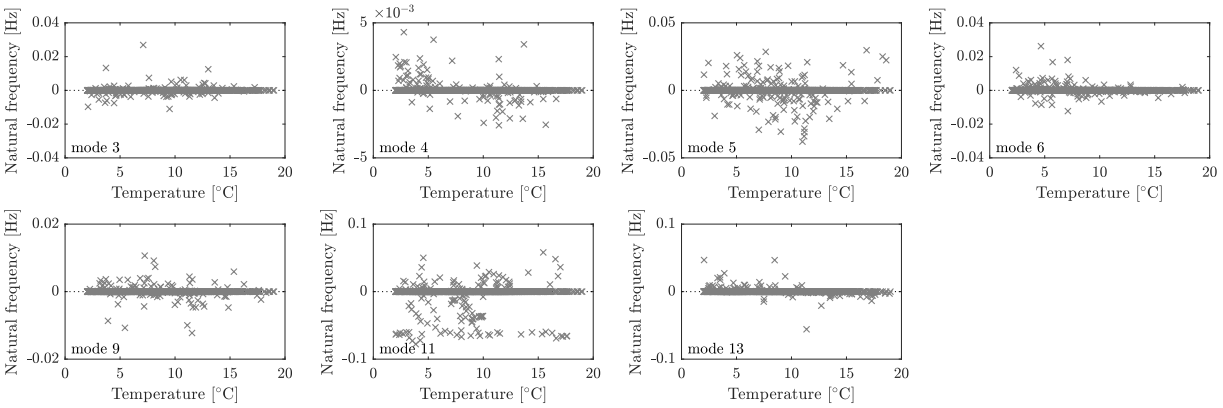

Fig. 12: Outlier matrix $\mathbf{S}$ plotted as a function of the bridge deck temperature.

is assigned the value $2/\sqrt{\max(n_\mathrm{m}, n_\mathrm{t})} = 2/\sqrt{n_\mathrm{t}}$. This value is chosen two times larger than the value proposed in [24] (i.e. $\lambda = 1/\sqrt{\max(n_\mathrm{m}, n_\mathrm{t})}$), since it was found that a lower value resulted in an outlier matrix $\mathbf{S}$ which was insufficiently sparse. It is apparent from the comparison of the original and the cleaned training data in Fig. 11 that some of the data points in $\mathbf{M}$ are classified as outliers and have therefore been corrected by the robust PCA algorithm, see for example modes 5 and 11. This also becomes clear from the outlier matrix $\mathbf{S}$, which is shown in Fig. 12.

Fig. 13 shows the identified principal components $\hat{\mathbf{\Psi}}$ obtained from robust PCA (in black). For the sake of comparison, the principal components of the original data matrix $\mathbf{M}$ are also shown (in gray), corresponding to the case of regular linear PCA. It is observed that the principal components obtained from robust PCA and linear PCA are nearly identical, allowing to conclude that outliers in this case do not play an important role. Since the principal components obtained from both methods are very similar, so is their detection capacity. In the remainder of this work, only the results of robust PCA are included.

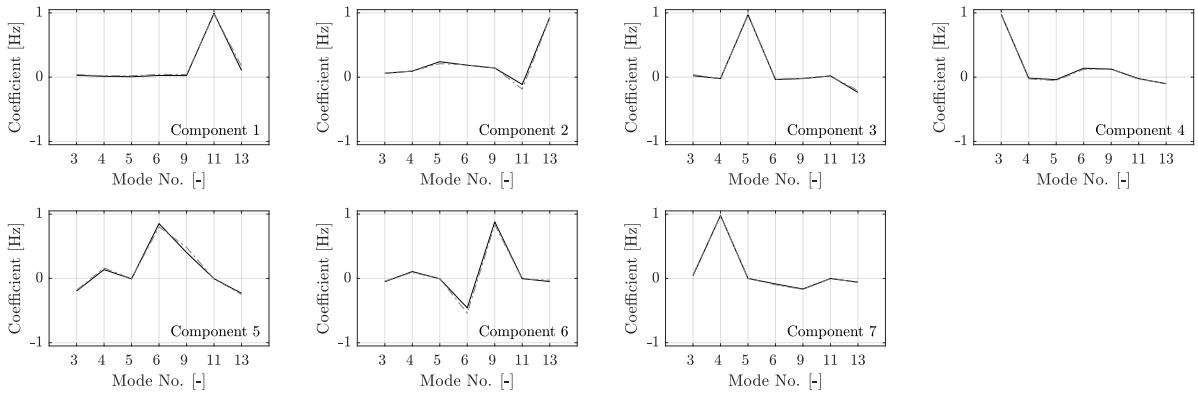

Fig. 13: Comparison of the principal components $\hat{\mathbf{\Psi}}$ obtained from robust PCA (solid black line) and linear PCA (dash-dotted gray line).

In order to investigate the contribution of the identified principal components to the data and select the optimal number of principal components n_c, the misfit during the training phase and the validation phase (before the retrofit) is studied. Fig. 14 presents the overall training misfit $\hat{\epsilon}_\mathrm{t}$ and the validation misfit $\hat{\epsilon}_\mathrm{v}$ (both in %), defined in Eq. (14), for all possible values of n_c, i.e. from 1 to 7. As a reference, the training misfit for the previously considered case of linear regression is also indicated ($\hat{\epsilon}_\mathrm{t} = 53.8\%$). In the evaluation of Eq. (14), it is assumed that the data points are rearranged such that the first n_t points correspond to training and the next n_v points correspond to validation.

$$\hat{\epsilon}_\mathrm{t} = \sqrt{\frac{\sum_{k=1}^{n_\mathrm{t}} ||\hat{\mathbf{e}}_k||_2^2}{\sum_{k=1}^{n_\mathrm{t}} ||\mathbf{m}_k - \hat{\mathbf{b}}||_2^2}}, \quad \hat{\epsilon}_\mathrm{v} = \sqrt{\frac{\sum_{k=n_\mathrm{t}+1}^{n_\mathrm{t}+n_\mathrm{v}} ||\hat{\mathbf{e}}_k||_2^2}{\sum_{k=n_\mathrm{t}+1}^{n_\mathrm{t}+n_\mathrm{v}} ||\mathbf{m}_k - \hat{\mathbf{b}}||_2^2}} \tag{14}$$

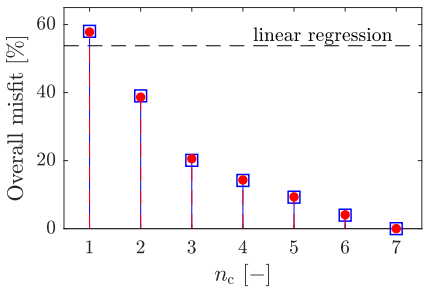

Fig. 14: Evolution of the overall training misfit $\hat{\epsilon}_\mathrm{t}$ (blue hollow squares) and the overall validation misfit $\hat{\epsilon}_\mathrm{v}$ (red solid circles) with the number of principal components n_c.

For the case of a single principal component ($n_\mathrm{c} = 1$), the training misfit ($\hat{\epsilon}_\mathrm{t} = 58.0\%$) is similar to the misfit obtained for the case of linear regression ($\hat{\epsilon}_\mathrm{t} = 53.8\%$). Further increasing the number of principal components leads to a clear reduction of the misfit, both in the training and the validation phase, showing the importance of the higher principal components in the description of variations present in the data. Indeed, while the first two components clearly have the largest contribution to the observed variation in the data, components 3 to 6 still represent about 5% of the observed variation each. It should be noticed that adding those higher principal components

does not lead to an overfit. This can be concluded from the fact that the validation misfit decreases along with the training misfit when increasing the number of principal components n_c. For $n_c = 7$, the misfit (by definition) equals zero. This is due to the fact that the principal components in this case constitute an orthogonal basis which enables describing any trend in the data. As such, selecting 7 principal components will not enable detecting any anomalies in the structural behavior and the number of principal components n_c should always be smaller than the number of outputs n_m to enable damage detection ($n_c < n_m$).

It is reminded that the actual number of components represented in the data, referred to as $n_c^\star$, equals the number of independent inputs driving the variation in the output, i.e. in the identified natural frequencies. This number is typically identified from a drop in the training misfit spectrum (Fig. 14) between $n_c = n_c^\star$ and $n_c = n_c^\star + 1$, where the misfit for $n_c > n_c^\star$ approximates zero. The first $n_c^\star$ principal components are then assumed to represent the variation in data driven by the independent inputs, while higher principal components represent noise in the data. The optimal choice of n_c equals $n_c^\star$. In the presented case, however, the training misfit spectrum in Fig. 14 is not characterized by such a clear drop. This may be due to the fact that the actual number of driving inputs $n_c^\star$ is equal to or larger than the number of outputs accounted for in the analysis, i.e. $n_c^\star \geq n_m = 7$. Since the first six principal components contribute significantly to the variation in the data, they are all retained in the construction of the PCA model for damage detection ($n_c = 6$).

From the previous analysis, it can be assumed that at least six independent environmental factors simultaneously affect the natural frequency estimates. One of them most likely relates to the bridge deck temperature considered in the linear regression. Other factors could for example relate to air humidity, the ballast temperature and the temperature of the bearings. The latter two are most probably different from the steel temperature below the bridge deck, due to thermal inertia of the structure and solar radiation. Some of the influencing factors are measured while others are not. The explicit study of the factors influencing the dynamic bridge response is outside the scope of this paper.

Fig. 15 shows the evolution of the model misfit $\hat{\mathbf{e}}_k$ for each of the 7 retained modes and $n_c = 6$. Similar to the case of linear regression, the model trained for temperatures above 2°C is extended to negative temperatures (in gray), for the sake of completeness. The results obtained for temperatures below 2°C are disregarded in the actual damage detection.

In line with the results obtained for the linear regression model in Section 3.1, the misfit $\hat{e}_{j,k}$ before the retrofit is mostly contained in the $[-3\sigma_j, +3\sigma_j]$ interval, both for the training and validation phase, where σ_j is the standard deviation of the misfit $\hat{e}_{j,k}$ during the training phase. Please note that the the standard deviation σ_j of the training misfit for the robust PCA model for all modes except for mode 4 is significantly smaller than the corresponding value obtained for the linear regression model, see also Table 3. This indicates that the selected robust PCA model better captures the trends in the natural frequency data. After the retrofit, a very large shift in the misfit of the robust PCA model is observed, exceeding for each of the seven modes the $[-3\sigma_j, 3\sigma_j]$ interval for all data points.

Fig. 16 shows the time history of the overall normalized misfit $\hat{\epsilon}_k$ of the robust PCA model, calculated in accordance to Eq. (6). Similar to the case of linear regression, this overall misfit accounts for the contribution of all modes and is therefore the preferred indicator in the process of damage detection. A distinction is made between the hourly basis misfit and the misfit obtained after averaging over 24 hours. Due to the normalization (w.r.t. the misfit standard deviation σ_j),

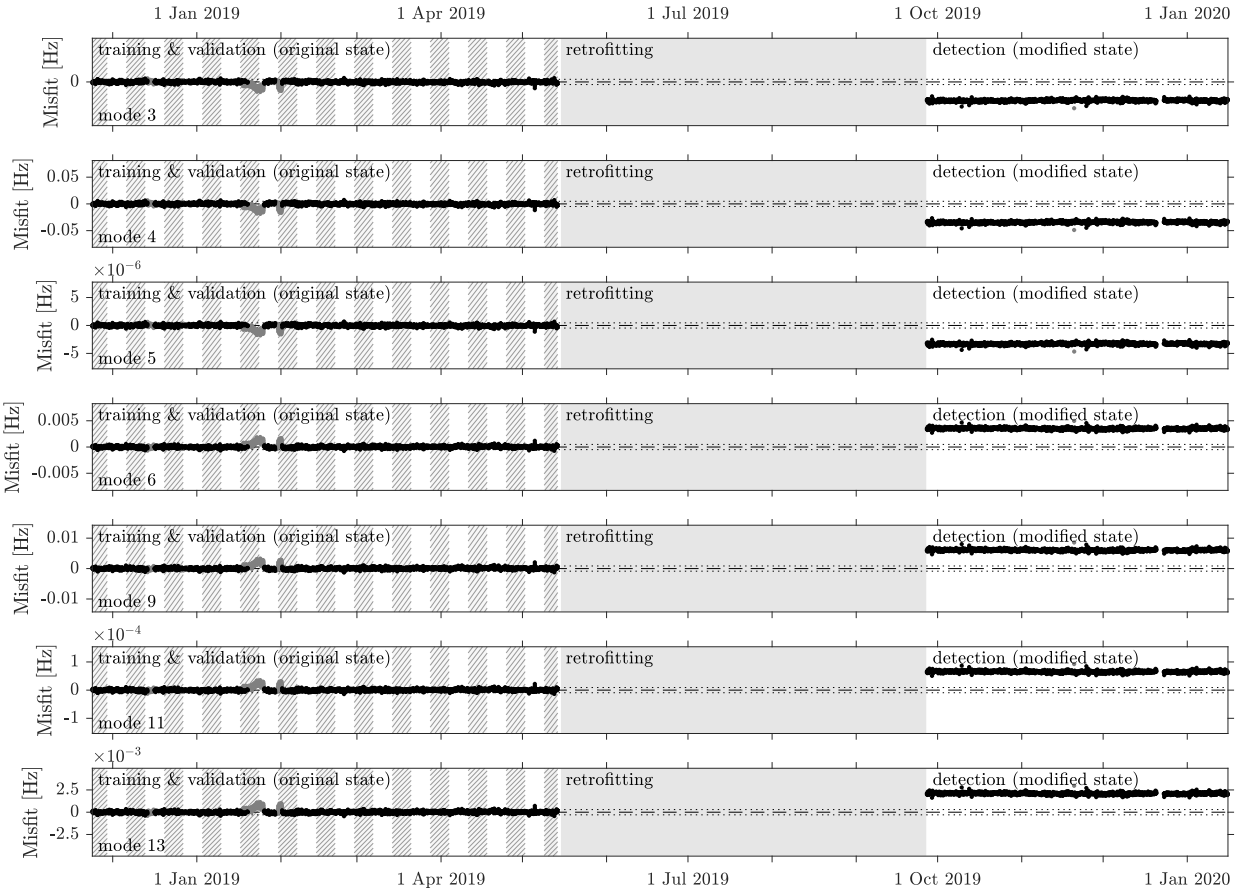

Fig. 15: Evolution of the robust PCA model misfit $\hat{e}_{j,k}$ for modes 3, 4, 5, 6, 9, 11, and 13 over time. Results corresponding to a steel surface temperature above 2°C are indicated in black, those corresponding to a temperature below 2°C in gray. The horizontal black dashed line indicates the zero misfit level. The data blocks composing the training period are hatched. The horizontal black dotted lines indicate the $+3\sigma_j$ and $-3\sigma_j$ misfit levels.

Mode No. j	3	4	5	6	9	11	13
Linear regression	3.0×10^{-3}	5.7×10^{-4}	4.4×10^{-3}	1.8×10^{-3}	1.4×10^{-3}	2.3×10^{-2}	5.1×10^{-3}
Robust PCA	9.2×10^{-5}	1.6×10^{-3}	1.6×10^{-7}	2.9×10^{-4}	2.7×10^{-4}	3.1×10^{-6}	9.9×10^{-5}

Table 3: Comparison of the standard deviation σ_j of the misfit $\hat{e}_{j,k}$ (in Hz) during the training phase between linear regression and robust PCA ($n_c = 6$).

the results can be directly compared to those obtained for the linear regression model in Figs. 8 and 9. The shift in the robust PCA model misfit is nearly twice as large as the shift in the linear regression model misfit. This conclusion holds for both the hourly-based misfit and the 24 hour averaged misfit. As such, it can be concluded that the robust PCA model outperforms the linear regression model in the identification of the present structural modification.

In comparison to the linear regression model, the robust PCA model has a slightly larger complexity. Nonetheless, it is still straightforward to implement and easy to interpret the results. The main advantage of a PCA model over a linear regression model is that it does not require knowledge on the inputs that drive the variation in the natural frequencies. Since it is practically infeasible to identify and measure all inputs, the robust PCA model better captures the trends in the natural frequency data, although still assuming a linear relation between the different natural frequencies.

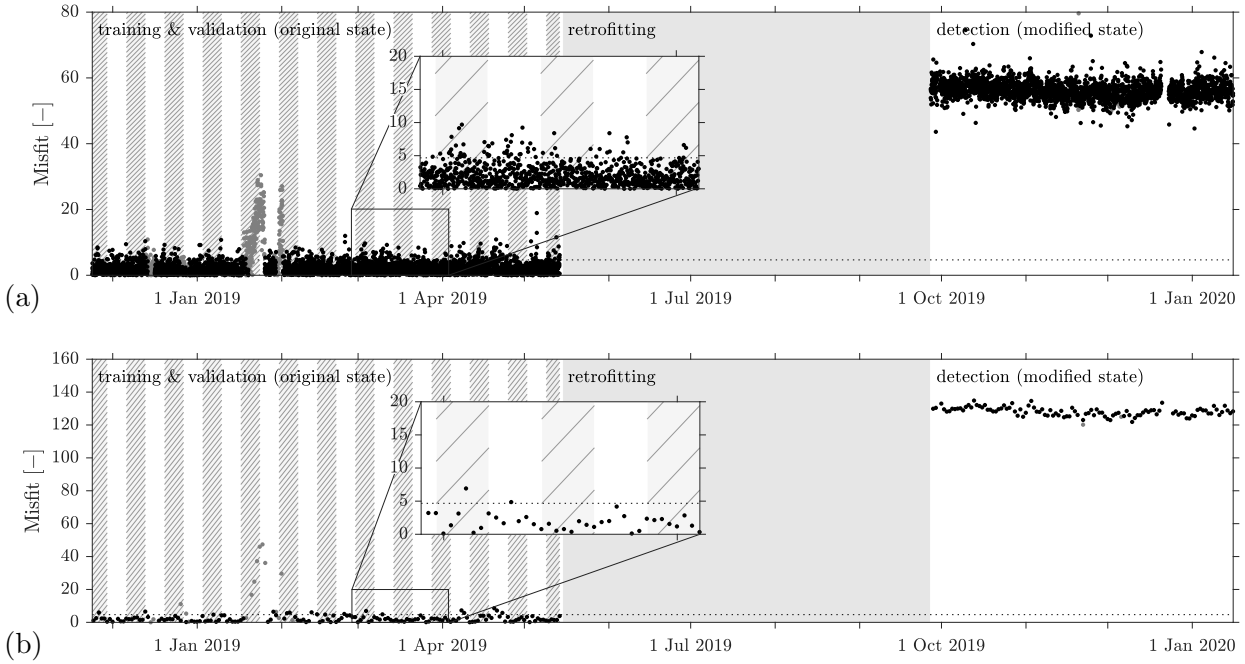

Fig. 16: Evolution of the normalized robust PCA model misfit $\hat{\epsilon}_k$ over time, (a) on an hourly basis and (b) averaged over 24 hours. Results corresponding to a steel surface temperature above 2°C are indicated in black, those corresponding to a temperature below 2°C in gray. The data blocks composing the training period are hatched. The horizontal black dotted line indicates the $c_{\chi^2}^{0.997}$ misfit level ($c_{\chi^2}^{0.997} = 4.67$).

In a final step, the performance of a robust PCA model trained using data for both positive and negative temperatures is evaluated. Fig. 17 shows the time history of the overall normalized misfit $\hat{\epsilon}_k$ obtained in this case for $n_c = 6$. In the outlier removal, the regularization parameter λ is further increased to $6/\sqrt{\max(n_m, n_t)} = 6/\sqrt{n_t}$ (instead of $2/\sqrt{n_t}$ in the previously considered case). This reduces the outlier detection capacity of robust PCA, since it puts more focus on the sparsity of the outlier matrix $\mathbf{S}$. However, when the value of λ is not increased, the data points corresponding to negative temperatures are wrongly classified as outliers and corrected. In contrast to the previously considered case where only data corresponding to temperatures above 2°C are used in the training phase, the misfit does not show a clear increase during frost periods. Meanwhile, it is also observed that the misfits during the training and validation phases (before the retrofit) have a very similar amplitude, indicating that the PCA model succeeds in modeling the overall variation in the natural frequencies, including both positive and negative temperatures. Nevertheless, extending the model towards temperatures below 2°C reduces its capacity to detect novelty. Indeed, the shift in the model misfit observed after the retrofit is smaller than in the original case where the model is limited to temperatures above 2°C (see also Table 4). The observed shift, however, is still slightly larger than the shift observed for the linear regression model in Fig. 8a. Similar conclusions can be drawn from the results obtained after 24 hour averaging, which are here omitted. As previously mentioned, the number of days per year where the temperature drops below 2°C for the considered structure is low. Limiting the model to temperatures above 2°C in this case only comes with an interruption of the detection during a few (potentially consecutive) days, meanwhile resulting in a higher potential to detect damage. If continued monitoring is needed in periods of frost, one might however choose to extend the model to both positive and negative temperatures.

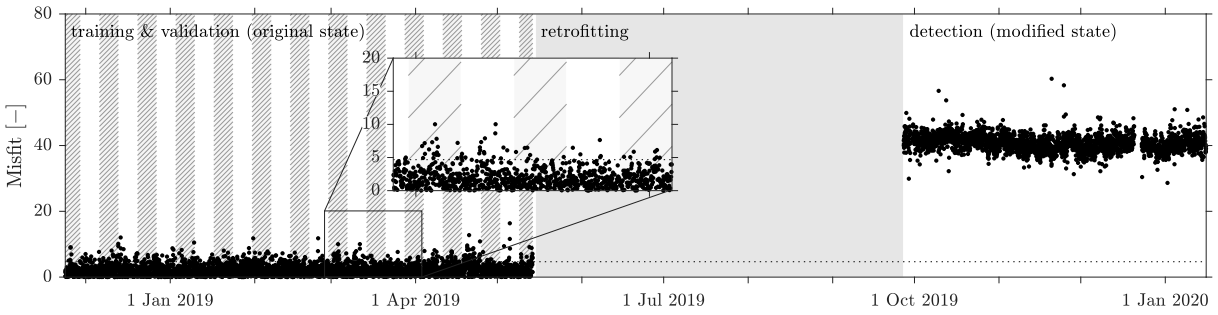

Fig. 17: Evolution of the normalized robust PCA model misfit $\hat{\epsilon}_k$ over time on an hourly basis, for the case where all data (corresponding to both positive and negative temperatures) are adopted in the training. The data blocks composing the training period are hatched. The horizontal black dotted line indicates the $c_{\chi^2}^{0.997}$ misfit level ($c_{\chi^2}^{0.997} = 4.67$).

	hourly basis	24 hour basis
Linear regression, $T > 2°C$	26.21	67.43
Robust PCA, $T > 2°C$	56.14	128.23
Robust PCA, no T limitations	40.72	76.28

Table 4: Comparison of the mean misfit level $\hat{\epsilon}_k$ observed after the retrofit (modified state) between different model types (linear regression and robust PCA) and temperature (T) application ranges.

4. Verification for simulated structural changes

In the previous section, it was observed that both linear regression and robust PCA succeed in identifying the changes in dynamic behavior that are attributed to the retrofit of railway bridge KW51. In this section, a number of more subtle structural changes is modeled using a detailed finite element model of the structure. This enables a more extensive comparison of damage detection methods.

4.1. Finite element model of railway bridge KW51

A detailed finite element (FE) model of the railway bridge has been developed using the open source MATLAB toolbox Stabil, developed at the Structural Mechanics Section of KU Leuven [26, 27]. The model is based on blueprints of the structure. The main and transverse girders of the bridge deck, the diagonals, the arches, and the braces connecting the arches are modeled using Timoshenko beam elements. The orthotropic bridge deck is modeled using four-noded shell elements in combination with Timoshenko beam elements to model the u-shaped stiffeners. The ballast is also modeled using four-noded shell elements. A Young's modulus of 200 GPa and 550 MPa are assigned to the construction steel and the ballast layer, respectively. The offset between different elements is accounted for by means of so-called constraint equations. The bridge is supported by four pot bearings. The restrictions imposed by the pot bearings are indicated in Figs 18a and 18b. The longitudinal stiffness of the pot bearings at side Leuven is accounted for in the model by means of a linear spring with stiffness k. The resulting FE model has a total of 6079 nodes and 4857 elements. Fig. 18c provides a three-dimensional view of the FE model.

4.2. Model calibration

The FE model is calibrated using a set of experimental modal characteristics obtained by means of a detailed output-only operational modal analysis (OMA), carried out before the retrofit on 16/02/2018, see also [15, Appendix 3]. In the OMA, the acceleration response of the arches and the bridge deck is processed using the reference-based covariance-driven stochastic subspace identification (SSI-cov/ref) algorithm [18]. The calibration parameters considered in the model updating consist of the thickness h of the ballast layer, the longitudinal stiffness k of the pot bearings (Fig. 18b), the out-of-plane bending stiffness EI_{yy} of the connections between the diagonals, the bridge deck, and the arches (indicated in blue in Fig. 18c), and the axial stiffness EA of these connections. In the original model (i.e. before calibration), the connections between the diagonals and the bridge deck and the arches are assigned the same section properties as the central part of the diagonals, which corresponds to a tubular section with large bending stiffness (Fig. 19). At the connections of the diagonals with the bridge deck and the arches, however, the cross section is reduced to a single plate with a much lower out-of-plane bending stiffness. As shown in Fig. 19, the out-of-plane bending stiffness EI_{yy} at the connection between the tubular profile and the plate equals 0.5×10^6 Nm^2, which is about 150 times lower than the bending stiffness of the tubular profile. Notice, however, that the calibration parameter EI_{yy} accounts for the entire connection between the tubular profile of the diagonal and the main girder, which also includes sections with a larger bending stiffness. The out-of-plane bending stiffness of the plate cross section indicated in Fig. 19 therefore consists a lower boundary for the model calibration parameter EI_{yy}. The same applies to the axial stiffness of the connections EA.

The model calibration is performed by means of an extensive grid search, where the relative difference between the measured and predicted natural frequencies is minimized and the correlation between the measured and predicted mode shapes, measured in terms of the MAC-values [28], is maximized. The calibration is performed in two steps. In the first step, a coarse grid is adopted for each of the four parameters. The search points for the ballast thickness h, the out-of-plane bending stiffness of the diagonal connections EI_{yy} and the axial stiffness of the diagonal connections EA are spaced linearly. Meanwhile, the search points for the bearing stiffness k are spaced logarithmically. From this first analysis, it became clear that the fit with the experimental data was improved as the stiffness parameter k was increased. A very high stiffness value, corresponding to hinged supports, results in the overall best match between the measured and the predicted modal characteristics. This suggests that friction in the bearings is preventing horizontal displacements during low-amplitude ambient vibrations. In a second updating step, the supports at side Leuven

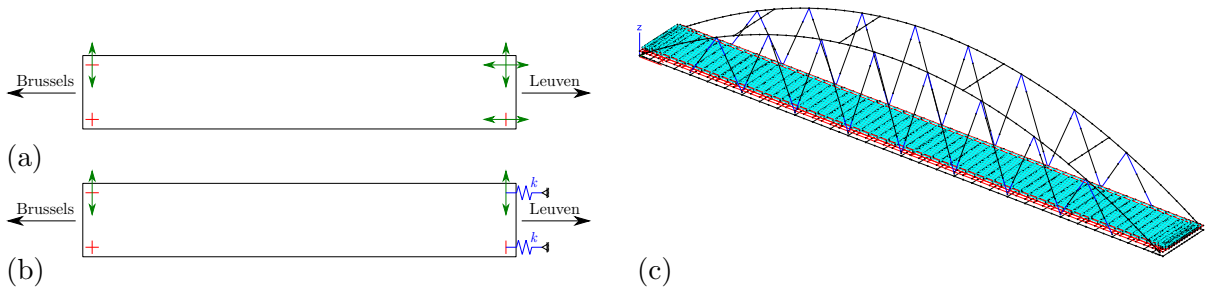

Fig. 18: Top view indicating the boundary conditions (a) according to the blueprints and (b) as modeled, and (c) 3D view of the FE model.

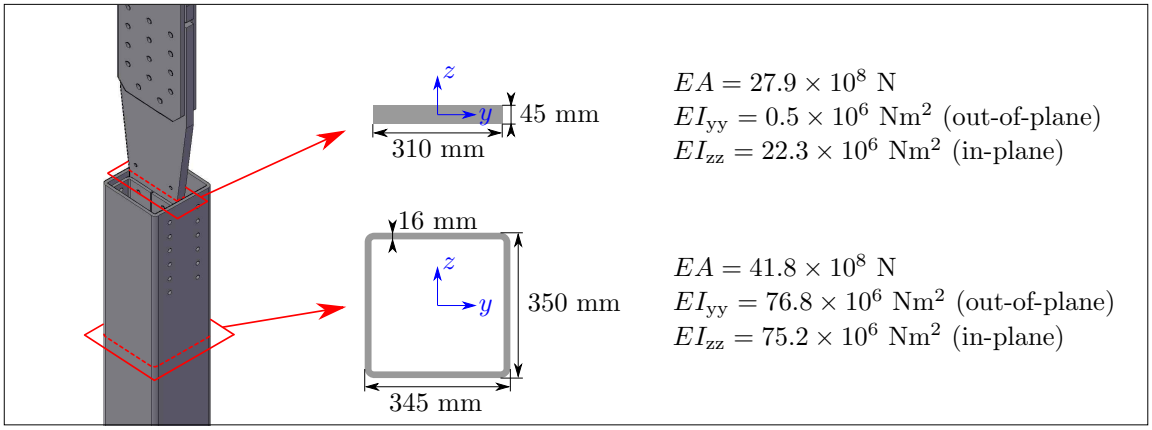

Fig. 19: Section properties of the tubular profile constituting the diagonals (bottom) and the plate section at the connection between the diagonals and the bridge deck (top).

are fixed ($k = \infty$) and a fine grid is adopted for the three remaining parameters, considering in total 1287 parameter combinations. The ballast thickness h is varied from 0.6 to 0.72 m in steps of 0.01 m (8.33% of the range). The out-of-plane bending stiffness of the diagonal connections EI_{yy} is varied from 3.46×10^6 to 6.53×10^6 Nm^2 in steps of 0.38×10^6 Nm^2 (12.5% of the range). Finally, the axial stiffness of the diagonal connections EA is varied from 29.2×10^8 to 37.6×10^8 N in steps of 0.8×10^8 N (10% of the range). The search range of the parameters is based on the results of the first coarse grid search. The best match between the measured and the predicted modal characteristics is obtained for the parameter values listed in Table 5.

	Initial model	Calibrated model
Ballast thickness h	0.60 m	0.66 m
Bearing stiffness k	0 N/m	∞
Bending stiffness EI_{yy}	76.8×10^6 Nm^2	4.6×10^6 Nm^2
Axial stiffness EA	41.8×10^8 N	32.6×10^8 N

Table 5: Value of the calibration parameters before and after model calibration (before retrofitting).

Table 6 compares the identified modal characteristics to those obtained from both the initial and the calibrated model. It is clear that the model calibration results in a significant improvement of the agreement between the measurements and the model predictions. This can be seen from (1) the increased number of modes that can be matched (9 modes for the initial model and 14 (all) modes for the calibrated model), (2) the significant overall reduction in the prediction errors on the natural frequencies, and (3) the much closer overall agreement between the identified and calculated mode shapes, as is evidenced by the very high MAC values close to 1 [28]. In general, all modes are very well predicted by the calibrated model, except for mode 11, which is the second lateral bending mode of the bridge deck. For this mode, the natural frequency is accurately predicted but a rather large discrepancy of the mode shape remains.

OMA		Initial model				Calibrated model			
No.	f_{id} [Hz]	No.	f_{fem} [Hz]	ϵ [%]	MAC	No.	f_{fem} [Hz]	ϵ [%]	MAC
1	0.51	1	0.70	36.98	0.99	1	0.53	2.94	1.00
2	1.24	2	1.41	13.87	0.98	2	1.26	1.57	1.00
3	1.90	3	1.61	-15.44	0.95	3	1.83	-3.91	0.97
4	2.45	5	2.51	2.47	0.98	4	2.38	-2.93	0.97
5	2.60	-	-	-	-	5	2.69	3.51	0.97
6	2.92	6	3.00	2.90	0.95	6	2.84	-2.48	0.99
7	3.54	-	-	-	-	7	3.51	-0.75	0.89
8	3.90	-	-	-	-	15	3.77	-3.36	0.90
9	4.11	-	-	-	-	20	4.14	0.61	0.90
10	4.28	10	4.42	3.33	0.71	25	4.29	0.38	0.84
11	4.85	11	4.92	1.53	0.63	27	4.74	-2.35	0.56
12	5.34	-	-	-	-	30	5.29	-0.86	0.87
13	6.33	18	6.73	6.32	0.89	39	6.41	1.30	0.89
14	6.80	20	6.84	0.62	0.72	41	6.78	-0.30	0.78

Table 6: Comparison of the identified modal characteristics (before retrofitting) to those obtained from the initial and calibrated FE model (f_{id}: identified natural frequency, f_{fem}: calculated natural frequency, ϵ: frequency error ($= (f_{\text{fem}} - f_{\text{id}})/f_{\text{id}}$), and MAC: Modal Assurance Criterion [28]).

4.3. Retrofit modeling

In a next step of the model validation, it is investigated to what extent the FE model is able to capture the natural frequency shifts induced by the actual retrofit, which consisted of welding steel boxes around the original bolted connection at every intersection of the diagonals with the arches and the bridge deck (see Section 2.3 and Fig. 3c).

For the sake of simplicity, it was decided to account for the stiffening effect introduced by the retrofitting by means of two properly selected model parameters, i.e. the out-of-plane bending stiffness EI_{yy} and the axial stiffness EA of the connections between the diagonals, the bridge deck, and the arches, rather than modeling the modification of these connections in full detail. The stiffness parameters EI_{yy} and EA corresponding to the situation after retrofitting are determined by means of an additional model calibration step. In addition to the adjustment of the stiffness of the connections in the model, the mass of the steel boxes is accounted for by means of a local point mass of 333 kg at the location of each steel box.

Rather than tuning the calibration parameters to match the identified natural frequencies after the retrofit, the sum of least square differences between the observed and predicted shifts in natural frequencies (in %) between the original and retrofitted state is now minimized. It was decided to consider the shifts in natural frequencies between the original and retrofitted state as opposed to actual values, as it was believed that these would be less affected by modeling errors. The frequency shifts are determined from the temperature-compensated misfit shown in Fig. 7, and are calculated for every mode separately as the difference between the average misfit after the retrofit and the average misfit during the training phase before the retrofit. The resulting frequency shifts are reported in Table 8. In the model-based prediction of the frequency shifts, the previously discussed calibrated FE model of the railway bridge, corresponding to the situation before the retrofit (Section 4.2), is used as a reference.

Table 7 compares the values of the stiffness parameters of the FE models representing respectively the structure before retrofitting (model 1) and after retrofitting (model 2). As expected,

the installation of the boxes on the diagonals results in a significant increase of both the bending stiffness EI_{yy} (increase of 42%), and the axial stiffness EA (increase of 21%).

	Calibrated model 1 Before retrofitting	Calibrated model 2 After retrofitting
Bending stiffness EI_{yy}	4.6×10^6 Nm^2	6.5×10^6 Nm^2 (+42%)
Axial stiffness EA	32.6×10^8 N	39.3×10^8 N (+21%)

Table 7: Value of the calibration parameters corresponding to the FE models representing respectively the structure before retrofitting (model 1) and after retrofitting (model 2).

Fig. 20 compares the identified natural frequency shifts to those obtained from the calibrated FE model (model 2 versus model 1). Except for modes 1 and 14, a very clear correlation between the identified and predicted natural frequency shifts is observed. This demonstrates that the model is able to accurately describe the influence of the retrofit.

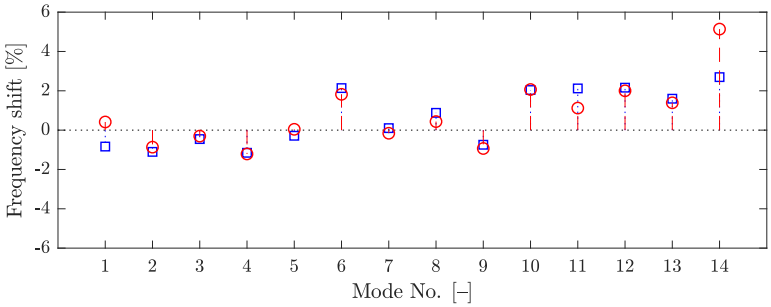

Fig. 20: Comparison of the identified (blue squares) and calculated (red circles) natural frequency shifts attributed to the retrofitting. A positive frequency shift corresponds to an increase of the natural frequency and vice versa.

4.4. Simulation of subtle structural modifications and validation of damage detection

It was previously observed that the actual retrofit resulted in relatively large shifts of the natural frequencies. Since these frequency shifts exceed the observed level of variation due to changes in environmental conditions for several modes (Fig. 5), the misfit obtained for both linear regression and robust PCA clearly indicates the transition from the original state of the railway bridge to its modified state after retrofitting.

In this section, it is investigated to what extent both methods succeed in detecting less pronounced structural modifications. Six different modification cases are considered. Each case corresponds to a situation where – in contrast to the actual retrofit, where all connections are modified – only one or two of the connections between the diagonals and the bridge deck (indicated in Fig. 21) are retrofitted by welding a steel box around the original bolted connection (Fig. 3c). As known from Section 4.3, the FE model is able to reproduce the frequency shifts when all connections are modified. As such, it is reasonable to assume that the model can also be used to predict the natural frequency shifts when only a limited number of connections (here 1 or 2) is modified. Although these modifications do not represent damage as such, they are more subtle when compared to the actual retrofit and may better reflect local changes corresponding to realistic damage.

Starting from the calibrated model for the original state, the characteristics of the modified connections are adjusted in accordance to what was obtained from the second model calibration

in Section 4.3: a local point mass of 333 kg is added at the location of the steel box, meanwhile increasing the bending stiffness EI_{yy} from 4.6×10^6 Nm^2 to 6.5×10^6 Nm^2 and the axial stiffness EA from 32.6×10^8 N to 39.3×10^8 N. In the following analysis, the actual retrofit is considered as a reference case and has therefore also been included. In the selection of the six modification cases in Fig. 21, a distinction is made between a modification at the bottom or the top connection of the diagonal. In addition, the distance of the modified diagonal with respect the center of the bridge is varied.

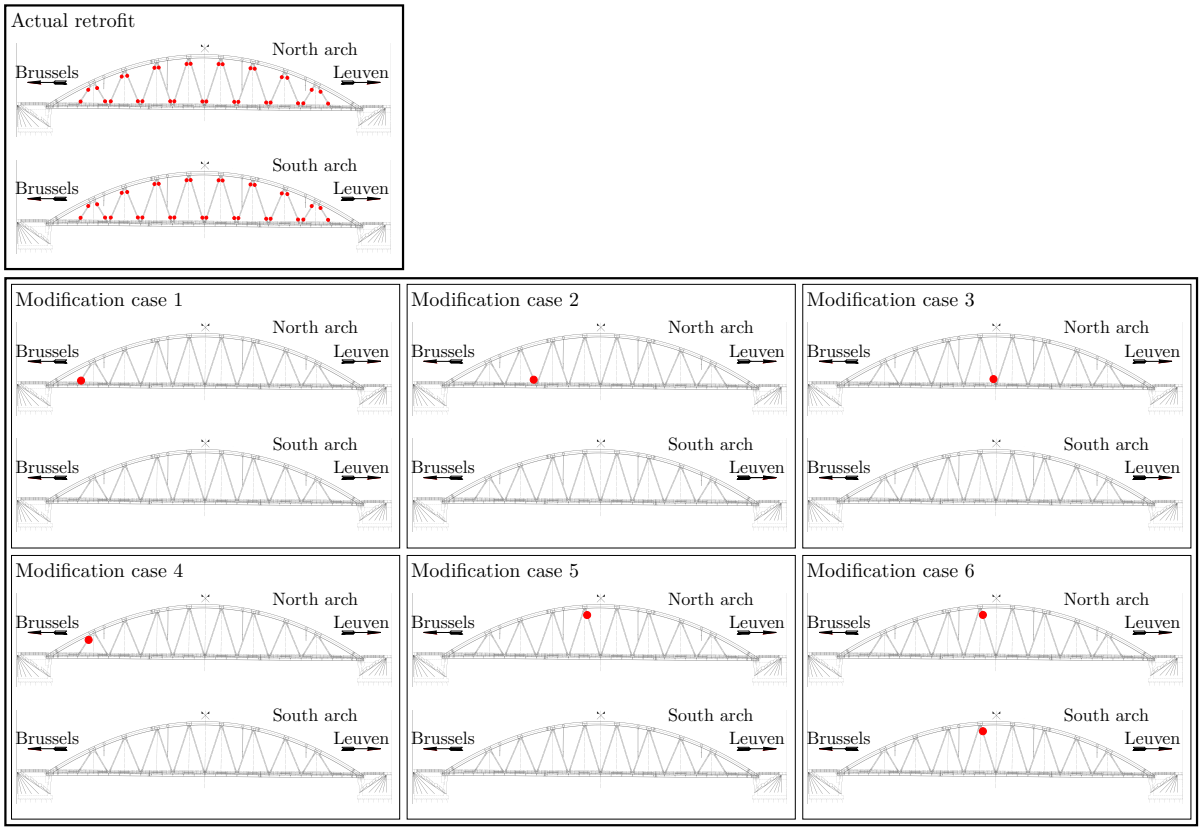

Fig. 21: Indication of the modified connections between the diagonals, the bridge deck, and the arches (red dots) for the actual retrofit (top) and the six simulated modification cases (bottom).

Table 8 presents the frequency shifts attributed to the actual retrofit and the six modification cases for each of the fourteen previously considered modes. For the actual retrofit, the identified frequency shifts are provided (Section 4.3), whereas for the simulated cases the natural frequency shifts are predicted using the FE model of the structure as explained before. As expected, the frequency shifts corresponding to the simulated cases are significantly smaller than those attributed to the actual retrofit. Furthermore, it should be noticed that the assigned frequency shifts are all smaller than the observed variation of the natural frequencies in the period before the retrofit (see Fig. 5). In the following, it is investigated to what extent the two damage detection methods presented in Section 3 enable detecting the changes in structural state attributed to each of the six considered cases.

Following the approach presented by Magalhães et al. in [13], the modeled frequency shifts corresponding to the six simulated scenarios are introduced in the actual validation data (before

Mode No.	1	2	3	4	5	6	7
Actual retrofit	-4.28	-13.70	-8.60	-28.16	-7.37	62.27	3.62
Modification case 1	0.11	0.18	-0.01	0.15	0.08	-0.07	-0.03
Modification case 2	0.17	0.16	-0.20	0.08	0.18	-0.06	0.00
Modification case 3	0.22	0.02	-0.39	0.10	-0.33	2.11	1.48
Modification case 4	0.00	-0.27	0.01	-1.26	0.10	-0.10	-0.92
Modification case 5	-0.28	-0.10	0.04	-1.58	-0.26	2.03	-0.17
Modification case 6	-0.56	-0.20	0.09	-3.17	-0.52	4.07	-0.35
Mode No.	8	9	10	11	12	13	14
Actual retrofit	34.12	-30.48	87.36	102.11	114.51	100.35	184.87
Modification case 1	-0.21	0.00	0.23	-0.57	1.82	3.72	-0.46
Modification case 2	-0.09	-0.16	1.65	-0.92	1.71	0.99	1.38
Modification case 3	1.99	-0.84	-0.07	4.25	1.72	0.31	-17.83
Modification case 4	-1.61	0.01	0.41	-0.88	2.70	3.01	-1.00
Modification case 5	4.20	-0.71	0.08	3.88	1.86	0.73	-17.60
Modification case 6	8.89	-1.47	0.17	7.82	3.77	1.45	90.59

Table 8: Frequency shifts in [mHz] identified (actual retrofit) and modeled (modification cases 1 – 6). The number of the modes (Mode No.) that are included in the damage detection are underlined.

the retrofit). The validation period of 12 non-consecutive weeks (see for example Fig. 7) is split in two blocks of 6 weeks each. For the first blocks of 6 weeks, the original data are used without modification. For the second block of 6 weeks, the frequency shifts listed in Table 8 are added to the original data, as such imposing a frequency shift. The resulting natural frequency data are then passed to both the linear regression model and the robust PCA model, discussed in Sections 3.1 and 3.2, respectively. Only data corresponding to temperatures above 2°C are considered.

In the remainder of this section, the results are presented by means of so-called Receiving Operator Characteristic (ROC) curve analyses. ROC curves provide a tool to express a model's performance based on the trade-off between Type I errors (false positives) and Type II errors (false negatives) [29]. In this particular case, when comparing the misfit of a given prediction model (linear regression or robust PCA) to a certain threshold (TH), the outcomes can be grouped into the following four categories:

1. True negative (TN) – The structure is in its original state (cfr. undamaged) and the misfit falls below the threshold.
2. False positive (FP) – The structure is in its original state and the misfit exceeds the threshold (Type 1 error).
3. True positive (TP) – The structure is in its modified state (cfr. damaged) and the misfit exceeds the threshold.
4. False negative (FN) – The structure is in its modified state and the misfit falls below the threshold (Type 2 error).

The ROC curve is constructed by plotting the false positive rate (FPR) on the horizontal axis against the true positive rate (TPR) on the vertical axis, for a varying threshold. As such, every point on the curve corresponds to a certain threshold value that defines the safety region. The true

positive rate (also called sensitivity) and the false positive rate (1-specificity) are defined as:

$$\mathrm{TPR} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}, \qquad \mathrm{FPR} = \frac{\mathrm{FP}}{\mathrm{FP} + \mathrm{TN}} \tag{15}$$

The optimal threshold results in the maximum number of correctly classified data points. Although many criteria exist, the Youden index (YI) is adopted here as a measure for the overall optimality of (1) the applied prediction model and (2) the adopted threshold [30]. The Youden index is defined as difference between the true positive rate and the false positive rate:

$$\mathrm{YI} = \mathrm{TPR} - \mathrm{FPR} \tag{16}$$

Conceptually, the Youden index corresponds to the vertical distance between the 45 degree line and the point on the ROC curve for a given threshold TH. A large Youden index (close to 1) indicates a better classification and vice versa.

Fig. 22 presents the ROC curves for the reference case, considering the actual retrofit, and for the six simulated modification cases. The results are summarized in Table 9. As previously discussed in Section 3, both linear regression and robust PCA enable a perfect detection of the actual retrofit (Fig. 22a). As can be observed from Figs. 8 and 16, there is an entire range of threshold values for which a perfect classification is achieved (8.6 to 22.2 for linear regression and 18.9 to 43.6 for robust PCA, both defined on an hourly basis). A maximum Youden index equal to 1 is obtained both for the original data on an hourly basis and those obtained after 24 hour averaging. When adopting a threshold TH = 4.67, no false negatives occur (TPR = 1), but some false positives are obtained. In this case, the use of a larger threshold is preferred.

				Optimal threshold (max. YI)				Threshold = 4.67			
		max. YI		TPR [-]		FPR [-]		TPR [-]		FPR [-]	
		LR	RPCA	LR	RPCA	LR	RPCA	LR	RPCA	LR	RPCA
Actual retrofit	1h	<u>1.00</u>	<u>1.00</u>	<u>1.00</u>	<u>1.00</u>	<u>0.00</u>	<u>0.00</u>	<u>1.00</u>	<u>1.00</u>	0.03	0.10
	24 h	<u>1.00</u>	<u>1.00</u>	<u>1.00</u>	<u>1.00</u>	<u>0.00</u>	<u>0.00</u>	<u>1.00</u>	<u>1.00</u>	0.03	0.13
Modification case 1	1h	0.05	0.09	0.67	0.46	0.61	0.38	0.03	0.11	0.03	0.08
	24 h	0.22	0.39	0.98	0.66	0.76	0.27	0.07	0.17	0.06	0.12
Modification case 2	1h	0.04	0.07	0.70	0.45	0.66	0.38	0.03	0.11	0.03	0.08
	24 h	0.14	0.34	0.90	0.85	0.76	0.52	0.05	0.17	0.06	0.12
Modification case 3	1h	0.06	0.08	0.61	0.46	0.55	0.38	0.03	0.11	0.03	0.08
	24 h	0.17	0.36	0.66	0.76	0.48	0.39	0.07	0.17	0.06	0.12
Modification case 4	1h	0.11	0.30	0.66	0.59	0.55	0.30	0.03	0.26	0.03	0.08
	24 h	0.47	0.72	0.95	0.90	0.48	0.18	0.12	0.73	0.06	0.12
Modification case 5	1h	0.15	0.35	0.68	0.70	0.53	0.35	0.04	0.31	0.03	0.08
	24 h	0.54	0.77	0.76	0.83	0.21	0.06	0.12	0.85	0.06	0.12
Modification case 6	1h	0.40	0.69	0.78	0.84	0.38	0.15	0.10	0.70	0.03	0.08
	24 h	0.90	<u>1.00</u>	0.90	<u>1.00</u>	<u>0.00</u>	<u>0.00</u>	0.95	<u>1.00</u>	0.06	0.12

Table 9: Summary of the ROC curve analysis for the actual retrofit and the six simulated modification cases, considering linear regression (LR) and robust PCA (RPCA) and distinguishing between the results for 1 hour intervals (1 h) and those obtained after 24 hour averaging (24 h). For those cases where a correct decision is made for all observed samples, i.e. Youden index YI = 1.00, TPR = 1.00, and/or FPR = 0.00, the results are underlined.

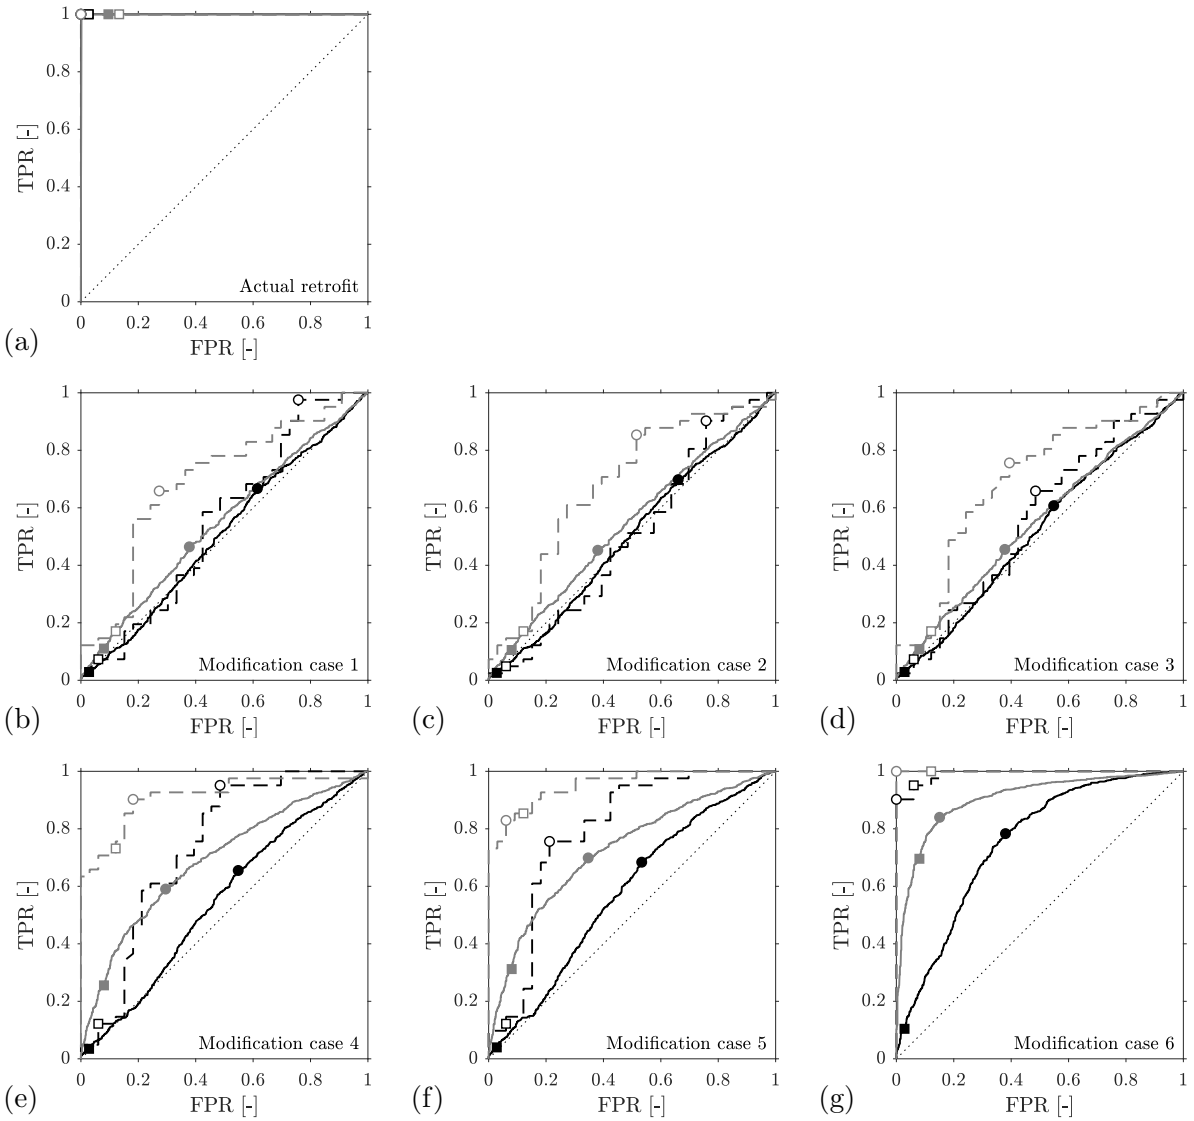

Fig. 22: ROC curves for the actual retrofit and the six simulated modification cases. The black curves correspond to linear regression, the gray curves to robust PCA. The solid lines correspond to the results for 1 hour intervals, the dashed correspond to the results obtained after 24 hour averaging. The circular markers indicate the point on the ROC curve corresponding to the maximum Youden index. The square markers indicate the point on the ROC curve for a threshold $c_{\chi^2}^{0.997} = 4.67$.

For the first three simulated cases, which concern the modification of one of the connections between a diagonal and the bridge deck, the linear regression does not enable a clear detection of the state transition. Indeed, both for the data on an hourly basis and those obtained after 24 hour averaging, the ROC curve lies around the 45 degree line, indicating a poor detection capability for every possible threshold value. Robust PCA, on the other hand, enables a better detection of the state transition, especially after 24 hour averaging. For example, for damage case 1, the threshold value corresponding to the maximum Youden index for the case of 24 hour averaging (TH = 2.01) results in 27% false positives and 34% false negatives. The success ratio and therefore the Youden index can be further improved by increasing the period of averaging. For the three considered cases (1 – 3), the imposed frequency shifts are small. As a result, adopting a threshold TH = 4.67

keeps the number of false positives low, meanwhile not enabling the detection of possible (small) structural modifications.

For cases 4, 5, and 6 (Figs. 22e, 22f, and 22g, respectively), the detection capability of both linear regression and robust PCA is much higher than for the first three cases. For each of the considered cases, the Youden index for robust PCA is significantly higher than for linear regression. Meanwhile, the benefit of the adopted 24 hour averaging is further demonstrated, leading in case 6 for the robust PCA model to a perfect distinction between the original and the modified state of the structure (YI = 1), similar to the case of the actual retrofit (Fig. 22a). In this case, where relatively large frequency shifts occur, choosing the threshold TH = 4.67 to keep the number of false positives low does not impede the damage detection. Finally, notice that the threshold resulting in the maximum number of correctly classified data points, i.e. the maximum Youden index, depends on the damage case considered and therefore cannot be determined in advance. If one wants to avoid a large number of false-positive detections and still achieve a high detection capability, it is advised to adopt a rather high threshold based on engineering judgment, for example TH = 4.67, meanwhile using a sufficiently long window for averaging the model misfit, as to remove remaining random variations.

5. Conclusions

In this work, two black-box modeling techniques are adopted to remove the effect of environmental variations on the natural frequencies obtained from a long-term monitoring campaign on a railway bridge in Leuven, Belgium. The removal of environmental variations is essential to enable damage detection in an early stage, where only small changes in the natural frequencies of the structure are expected. The first technique is linear regression, commonly adopted for temperature compensation in structural health monitoring applications. This technique is very straightforward to use and implement, but comes with the (often restricting) assumption that all environmental variables which drive the natural frequency variations are measured. The second technique, which overcomes this assumption, is robust PCA. This technique is relatively new to the field of structural health monitoring and serves as an extension of the commonly adopted linear PCA technique. Whereas the estimation of the principal components for classical linear PCA is sensitive to outliers in the training data, the robust PCA technique includes a separate outlier removal step prior the the actual principal component extraction.

The bridge monitoring data used in this study are particularly interesting since they involve an actual modification of the structure, by means of a retrofitting, that gave rise to a visible shift in the observed natural frequencies. The two modeling techniques are compared by means of a ROC curve analysis, hereby considering the actual retrofit and six structural modifications, which have been additionally modeled using a calibrated finite element model of the railway bridge. In general, it can be concluded that robust PCA outperforms linear regression for each of the considered cases. Furthermore, it is shown how averaging the model misfits over a time window containing multiple samples enables a further improvement of the damage detection capacity, even for small structural modifications that result in relatively small natural frequency shifts.

Finally, it should be noticed that the results obtained for robust PCA for the presented case are very similar to those obtained from classical linear PCA. Nevertheless, the robust PCA technique is preferred over the linear PCA technique, since it is enables diminishing the effect of potential outliers, which can lead to a misalignment of the principal directions of the model and therefore a lower damage detection capacity.

Acknowledgements

Kristof Maes is a postdoctoral fellow of the Research Foundation Flanders (FWO), Belgium (grant number 12Q9218N). FWO also provided additional funding for the measurements on railway bridge KW51 by means of research grant 1511719N. The financial support by FWO is gratefully acknowledged.

References

[1] A. Rytter, Vibration based inspection of civil engineering structures, Ph.D. thesis, Aalborg University, 1993.
[2] A. Teughels, G. De Roeck, Structural damage identification of the highway bridge Z24 by FE model updating, Journal of Sound and Vibration 278 (2004) 589–610.
[3] B. Peeters, G. De Roeck, One-year monitoring of the Z24-bridge: environmental effects versus damage events, Earthquake Engineering and Structural Dynamics 30 (2001) 149–171.
[4] E. Reynders, G. Wursten, G. De Roeck, Output-only structural health monitoring in changing environmental conditions by means of nonlinear system identification, Structural Health Monitoring 13 (2014) 82–93.
[5] E. Reynders, G. De Roeck, Robust structural health monitoring in changing environmental conditions with uncertain data, in: G. Deodatis, B. Ellingwood, D. Frangopol (Eds.), Proceedings of the 11th International Conference On Structural Safety And Reliability: ICOSSAR 2013, New York, NY, USA, pp. 2421–2428.
[6] C. Rainieri, F. Magalhaes, D. Gargaro, G. Fabbrocino, A. Cunha, Predicting the variability of natural frequencies and its causes by second-order blind identification, Structural Health Monitoring 18 (2019) 486–507.
[7] E. Cross, K. Koo, J. Brownjohn, K. Worden, Long-term monitoring and data analysis of the Tamar Bridge, Mechanical Systems and Signal Processing 35 (2013) 16–34.
[8] W. Shan, X. Wang, Y. Jiao, Modeling of temperature effect on modal frequency of concrete beam based on field monitoring data, Shock and Vibration 2018 (2018).
[9] Y. Ni, X. Hua, K. Fan, J. Ko, Correlating modal properties with temperature using long-term monitoring data and support vector machine technique, Engineering Structures 27 (2005) 1762–1773.
[10] A.-M. Yan, G. Kerschen, P. De Boe, J.-C. Golinval, Structural damage diagnosis under varying environmental conditions – part I: a linear analysis, Mechanical Systems and Signal Processing 19 (2005) 847–864.
[11] A. Deraemaeker, E. Reynders, G. De Roeck, J. Kullaa, Vibration based Structural Health Monitoring using output-only measurements under changing environment, Mechanical Systems and Signal Processing 22 (2008) 34–56.
[12] R. Azzara, G. De Roeck, M. Girardi, C. Padovani, D. Pellegrini, E. Reynders, The influence of environmental parameters on the dynamic behaviour of the San Frediano bell tower in Lucca, Engineering Structures 156 (2018) 175–187.
[13] F. Magalhães, A. Cunha, E. Caetano, Vibration based structural health monitoring of an arch bridge: from automated OMA to damage detection, Mechanical Systems and Signal Processing 28 (2012) 212–228.
[14] J. Maeck, G. De Roeck, Description of Z24 benchmark, Mechanical Systems and Signal Processing 17 (2003) 127–131.
[15] K. Maes, G. Lombaert, Monitoring railway bridge KW51 before, during, and after retrofitting, ASCE Journal of Bridge Engineering 26 (2021) 04721001.
[16] K. Maes, G. Lombaert, Monitoring data for railway bridge KW51 in Leuven, Belgium, before, during, and after retrofitting (Version 1.0) [Data set], Zenodo (2020).
[17] D. Anastasopoulos, G. De Roeck, E. P. B. Reynders, One-year operational modal analysis of a steel bridge from high-resolution macrostrain monitoring: influence of temperature vs. retrofitting, Mechanical Systems and Signal Processing 161 (2021) 107951.
[18] B. Peeters, G. De Roeck, Reference-based stochastic subspace identification for output-only modal analysis, Mechanical Systems and Signal Processing 13 (1999) 855–878.
[19] F. Magalhaẽs, A. Cunha, E. Caetano, Online automatic identification of the modal parameters of a long span arch bridge, Mechanical Systems and Signal Processing 23 (2009) 316–329.
[20] W. Heylen, S. Lammens, P. Sas, Modal analysis theory and testing, Department of Mechanical Engineering, Katholieke Universiteit Leuven, Leuven, Belgium, 1997.
[21] P. Mahalanobis, On the generalized distance in statistics, Proceedings of the National Institute of Sciences (Calcutta) 2 (1936) 49–55.

[22] A.-M. Yan, G. Kerschen, P. De Boe, J.-C. Golinval, Structural damage diagnosis under varying environmental conditions – part II: local PCA for non-linear cases, Mechanical Systems and Signal Processing 19 (2005) 865–880.
[23] E. Candès, X. Li, Y. Ma, J. Wright, Robust principal component analysis?, Journal of the ACM 58 (2011).
[24] S. Brunton, J. Kutz, Data Driven Science & Engineering – Machine Learning, Dynamical Systems, and Control, Cambridge University Press, 2017.
[25] F. Gharibnezhad, L. Mujica, J. Rodellar, Applying robust variant of principal component analysis as a damage detector in the presence of outliers, Mechanical Systems and Signal Processing 50–51 (2015) 467–479.
[26] S. François, M. Schevenels, D. Dooms, J. Jansen, J. Wambacq, G. Lombaert, G. Degrande, G. De Roeck, Stabil: an educational Matlab toolbox for static and dynamic structural analysis, Computer Applications in Engineering Education (2021) 1–18.
[27] Stabil: a MATLAB toolbox for structural mechanics (version 3.0), https://bwk.kuleuven.be/bwm/stabil, accessed on 13/08/2020.
[28] R. J. Allemang, The Modal Assurance Criterion – Twenty Years of Use and Abuse, Sound and Vibration Magazine (2003) 14–21.
[29] T. Fawcett, An introduction to ROC analysis, Pattern Recognition Letters 27 (2006) 861 – 874.
[30] W. Youden, Index for rating diagnostic tests, Cancer 3 (1950) 32–35.