A New Approach to Damage Detection in Bridges Using Machine Learning

A.C. Neves[1(✉)], Ignacio González[2], John Leander[1], and Raid Karoumi[1]

[1] Structural Engineering and Bridges, KTH-Royal Institute of Technology, Stockholm, Sweden
acneves@kth.se
[2] Sweco AB, Stockholm, Sweden

Abstract. At the same time that civil engineering structures are increasing in number, size and longevity, there is a conforming increasing preoccupation regarding the monitoring and maintenance of such structures. In this sense the demand for new reliable Structural Health Monitoring systems and damage detection techniques is high. A model-free damage detection approach based on Machine Learning is presented in this paper. The method performs on the collected feature measurements on a railway bridge, which for this study were gathered in a numerical experiment using a three dimensional finite element model. The first step of the approach consists in collecting the dynamic response of the structure, simulated during the passage of a train over the bridge, in both the healthy and damage states of the structure. The next step consists in the design and unsupervised training of Artificial Neural Networks that use as input accelerations and axle loads and compute a novelty index, called prediction error, based on a novelty detection approach. The distribution of the obtained prediction errors is statistically evaluated by means of a Gaussian Process and, after this process, damage indexes can be defined. Finally, the efficiency of the method is assessed in terms of Type I (false positive) and Type II (false negative) errors using Receiver Operating Characteristic curves. The promising results obtained in the case study demonstrate the capability of the presented method.

Keywords: Structural Health monitoring · Machine Learning · Damage detection · Model-free based method · Artificial Neural Networks

1 Introduction

1.1 General

Civil engineering structures are aging and being used past their life expectancy, at the same time carrying heavier traffic loads due to the increasing demand for transport capacity. Bridges in particular are a critical link in modern transport networks and, thus, this is probably the most appropriate time for the development of robust and reliable structural damage detection systems that ensure the operation of bridges in safe conditions.

J.P. Conte et al. (eds.), *Experimental Vibration Analysis for Civil Structures*,
Lecture Notes in Civil Engineering 5, https://doi.org/10.1007/978-3-319-67443-8_5

Structural Health Monitoring (SHM) consists in the implementation of a damage detection and classification strategy for engineering structures, making it a concept shared in many areas of research such as Aerospace, Civil and Mechanical engineering. The aptitude to monitor a structure and eventually detect damage at the earliest possible stage supports clever maintenance strategies and provides accurate remaining life predictions.

A review of some of the most recent developments within SHM divulged in published articles is here presented. A topic of relevant research interest is the optimal sensor location, when done adequately provides the maximum information with the least number of sensors, thus allowing for cost reduction. Regarding this topic, some works worth mentioning are [1–3]. One of the dominant research topics within SHM is the discrimination of the changes in structural response caused by operational and environmental variability (e.g. temperature fluctuation) from the changes caused by damage. An efficient way to make this distinction of sources of variability in structural behavior is by applying algorithms such as Artificial Neural Networks [4]. Machine Learning has highly contributed to most of the new advances in the field of SHM. These algorithms normally belong to the outlier detection category, which considers training data coming exclusively from the normal condition of the structure (unsupervised learning), exposing structural abnormalities from monitoring data. It is worth pointing out the important work done by Worden and Farrar [5] in monitoring of structures using machine learning techniques such as neural networks, genetic algorithms and support vector machines. Not surprisingly, the newest proposed methods that show superior performance, with improved accuracy and stability, result from the integration of several techniques that may already exist but that were not previously used in combination with others. Some examples are: a novel damage identification technique combining Proper Orthogonal Decomposition (POD) with time–frequency analysis using Hilbert Huang Transform (HHT) and Dynamic Quantum Particle Swarm Optimization (DQPSO) [6]; a structural damage detection method based on posteriori probability Support Vector Machine (PPSVM) and Dempster-Shafer (DS) evidence theory [7]. Finally, the performance of the damage detection method based on machine learning techniques, often so-called classifier, should be evaluated. Some commonly tools used to perform that evaluation are the Receiver Operating Characteristic curves [8] or Probability of Detection curves [9].

Based on the work of González [10], a model-free damage detection approach using Machine Learning techniques is presented in this paper. The method performs on the collected feature measurements on a railway bridge, which for this study consists of vertical linear accelerations gathered in a numerical experiment using a three dimensional finite element model. The first step of the approach consists in collecting the dynamic response of the structure, simulated during the passage of a train over the bridge, in both the healthy and damage states of the structure. The next step consists in the design and unsupervised training of Artificial Neural Networks that use as input accelerations and axle loads and compute a novelty index, called prediction error, based on a novelty detection approach. The distribution of the obtained prediction errors are statistically evaluated by means of a Gaussian Process and, after this process, damage indexes can be defined. Finally, the efficiency of the method is assessed in terms of Type I (false positive) and Type II (false negative) errors using Receiver Operating

Characteristic curves. The promising results obtained in the case study demonstrate the capability of the presented method.

1.2 Structural Health Monitoring

Structural Health Monitoring is a powerful instrument to ensure structural integrity and safety and it has become vastly popular over the past few decades. SHM consists in the implementation of a strategy to detect damage in infrastructures, combining a variety of sensing technologies with an embedded measurement system to capture, log, and analyze real-time data.

Implied in Axiom II of SHM [5], the damage identification process involves the comparison between two distinct states of the system, the baseline case and the atypical case. Therefore, SHM can be seen as a problem of statistical pattern recognition which is composed of four parts: operational evaluation; data acquisition, normalization and cleansing; feature extraction and data compression; statistical-model development for feature discrimination. Table 1 depicts a general scheme of the SHM process and what is involved in each of the parts.

Table 1. Parts of a SHM system

1. Operational Evaluation	➔	**2.** Data Acquisition, Normalization and Cleansing	➔	**3.** Feature Extraction and Data Compression	➔	**4.** Statistical model Development
- Life-safety and economic justifications to perform SHM; - Definition of damage to be detected; - Operational and Environmental conditions; - Data Acquisition limitations.		- Type and amount of data to be collected; - Periodicity in data acquisition; - Data normalization procedures; - Sources of variability.		- Selection of the best features of the data from damage detection; - Statistical distribution of the features; - Data condensation.		- Damage or not damaged; - Damage location; - Damage extension; - Damage Type; - Remaining useful life of the structure; - Incorrect diagnosis of damage (FP and FN).

Diagnosis of damage in structural systems primarily involves the identification of damage, followed by the identification of its location, type and severity. A robust SHM system is composed of the following stages of damage identification, accordingly to [5]:

- Level 1: Is there damage present in the structure?
- Level 2: What is the geometric location of the damage?
- Level 3: What is the type of damage?
- Level 4: What is the severity of the damage?
- Level 5: What is the prediction of the remaining service life of the structure?

The idea behind this hierarchy is that the higher the level of assessment the more information one will have about the structural condition but it also raises the difficulty in acquiring that information. Hence, the several levels will have different requirements on the types of sensors, the type of algorithm used for monitoring damage and number of model parameters.

The first stage of damage identification uses methods that provide a qualitative indication of the presence of damage in the structure, which can be accomplished without prior knowledge of how the system behaves when damaged – unsupervised learning. These algorithms are referred as outlier or novelty detection methods and to solve the task of novelty detection one can use learning algorithms such as Artificial Neural Networks (ANNs).

2 Method

2.1 Artificial Neural Networks

Artificial Neural Networks are a family of mathematical models that is inspired by the structure of biological neural networks in which the basic processing unit of the brain is the neuron (Fig. 1). Neurons interact with each other by summing stimuli from connected neurons (Fig. 2). Once the total stimuli exceed a certain threshold, the neuron fires - a phenomenon called activation - and it generates a new stimulus that is passed on into the network. Knowledge is encoded in the connection strengths between the neurons in the brain.

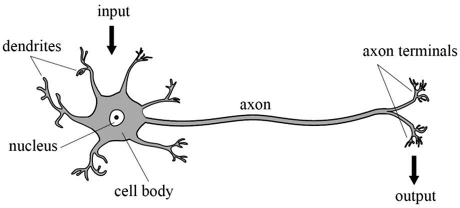

Fig. 1. The biological neuron.

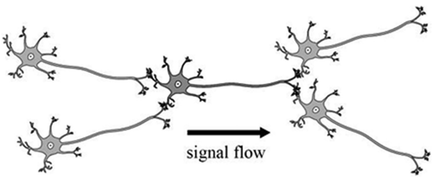

Fig. 2. Small network of neurons: the dendrites receive the input signals, the body of the cell (nucleus) is responsible for processing the input signals and the axon sends the signal from the body of the cell towards the neighbor neurons.

Mundane examples of applications of ANNs are: speech recognition and generation, optimization of chemical processes, manufacturing process control, cancer cell analysis, transplant time optimizer, recognition of chromosomal abnormalities, solution of optimal routing problems such as the Traveling Salesman Problem, etcetera. ANNs are a powerful tool for SHM in the aid of problems in sensor data processing that require parallelism and optimization due the high complexity of the variables' interactions. Generally, the ANNs offer solutions to four different problems: auto association, regression, classification and novelty detection.

The idea behind using ANNs for SHM is to use a data set of signal parameters obtained from a reference structure, such as an undamaged structure or a numerical model of a structure, and to use soft computing methods to warn about damage and its characteristics. Only the data from normal operating condition of the structure is used as training data – unsupervised learning. This is normally what happens in reality as concerning civil engineering structures one lacks the data for damaged condition of the structure of interest due to costs and practicality constraints. The downside of using such unsupervised methods is that they do not have much diagnostic ability beyond simple detection. With these methods, a reference model of the normal condition is first created and then the newly acquired data (e.g. from measurements of the structure) is compared with the data obtained from the model. If there are significant deviations, the algorithm is said to indicate novelty, meaning that the structure has departed from its normal condition and damage is probably present.

2.2 Receiver Operating Characteristic Curve

One approach that enables the statistical evaluation of the errors related with false detection is the Receiver Operating Characteristic (ROC) curve [11]. A ROC curve is a two-dimensional graphic in which the True Positive rate (TPr) is plotted in the *y* axis and the False Positive rate (FPr) is plotted in the *x* axis and the graphic demonstrates relative trade-offs between these benefits and costs, respectively, depending on a threshold that is selected, for example comparing a damage index (DI) with it. Recalling the definition of DI, one that surpasses the threshold will make the system warn for damage, whereas one that falls behind the threshold makes the system to not warn for damage. It is then understandable that a very high threshold will never indicate damage since the classifier finds no positives (resulting in 0% of False and True Positives), whereas a very low threshold will always indicate damage since everything is classified positive (resulting in 100% of False and True Positives). Figure 3 illustrates the Probability Density Functions (PDFs) of the null (structure is undamaged) and alternative (structure is damaged) hypotheses, which are in the basis of the process with which the ROC curve is created. When the threshold is placed to the right of the null distribution, damage is not detected and therefore both probabilities of TP and FP are zero. By translating the threshold to the left, the area under the null distribution increases (and thus does the probability of true detection) but so will the area under the alternative distribution increase (and thus does the probability of false detection). If the threshold is pulled to the extreme left, the TP and FP probabilities will approximate the unity. Using the distributions of Fig. 3 as an example, for the fixed threshold depicted the ROC registers approximately 70% TPs (shaded area) against 15% FPs (shaded area). Well-known characteristics of the ROC curve are:

- the trade-off between sensitivity and specificity (Fig. 4). An increase in sensitivity is achieved by moving the cutpoint to a higher value – making the criterion for a positive test less strict. An increase in specificity is achieved by moving the cutpoint to a lower value – making the criterion for a positive test more strict;
- the closer the curve comes to the left and the top borders of the ROC space, the more accurate is the damage detection method (– ROC curve in Fig. 4); the closer

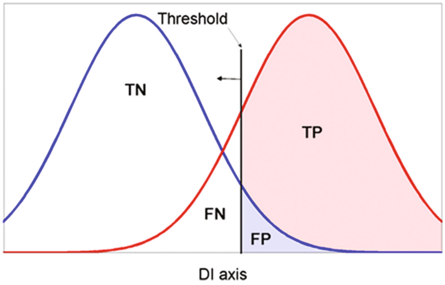

Fig. 3. ROC curve construction. – Null (Normal condition) and – Alternative (Abnormal condition) PDFs along with the detection threshold moving from the right (higher threshold) to the left (lower threshold). TP – True Positive; FP – False Positive; FN – False Negative; TN – True Negative.

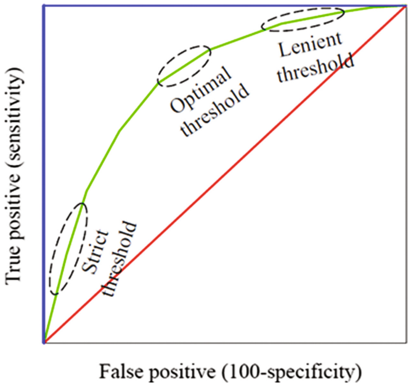

Fig. 4. ROC curves: – excellent; – good; – worthless.

the curve comes to the 45-degree diagonal of the ROC space, the less accurate is the damage detection method (– ROC curve in Fig. 4);

- the area under the curve is a convenient way of comparing the classifier methods' accuracy: an ideal perfect one as a value of 1.0 whereas a random worthless one has a value of 0.5;
- the slope of the tangent line at a certain threshold gives the likelihood ratio for that value of the test;

3 Bridge and Finite Element Model

A numerical 3D model of a single-track railway bridge was developed in the FEM software ABAQUS [12]. The structure consists of a concrete deck, two steel girder beams that support the deck and steel cross bracings that connect the girders. The deck and the girder beams were modelled as shell elements and the cross bracings were modelled as truss elements. All the elements of the bridge are assumed to be rigidly connected to each other.

Damage in the bridge is simulated considering two damage scenarios: in damage case 1, a section of the bottom flange of one girder beam is removed (Fig. 5), in an attempt to represent a damage situation where a fatigue crack exists. The cut out section has the dimensions of some longitudinal length l by the flange width, reflecting a situation when a propagating crack has reached its critical depth (about 30% of the flange's width or less) causing the sudden rupture through the whole flange width; in damage case 2 one bracing is removed (Fig. 6), which equivalently corresponds to reducing to approximately zero its Elastic modulus in the model. Assuming that girder

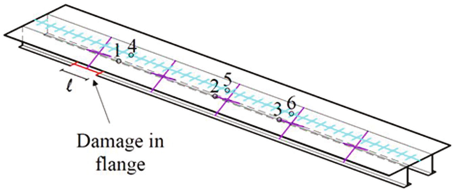

Fig. 5. Damage Case 1 (DC1). – Track alignment; – Bracing; – Damage location; **1-6** Sensor numbering and location.

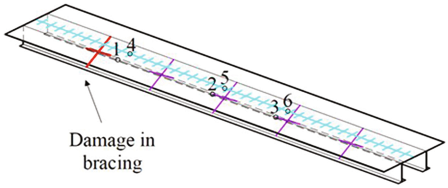

Fig. 6. Damage Case 2 (DC2). – Track alignment; – Bracing; – Damage location; **1-6** Sensor numbering and location.

beam and bracings are connected by high-tension bolts, this can reflect a situation where there is looseness in the bolted connection [13] yielding the bracing to become inefficient in its function. The numbers 1 to 6 in the figures below represent the locations of the accelerometers that are installed on the top of the bridge deck: three aligned with the train track and three aligned with the girder beam in which damage in *DC1* takes place.

The proposed method for structural assessment is intended to identify existing damage from the measured vibration of the bridge. Dynamic loads typically come from traffic, which is expected to be continuous while the bridge is in service. Traffic induced vibration was simulated in the numerical model by means of the passage of a train with a fixed configuration, crossing the bridge with a speed within the range [70–100] km/h, in increments of 0.1 km/h. A total of 300 different train passages were simulated and the corresponding measurement data sets were gathered and saved. The moving axle loads were modelled as series of constant moving forces with short time steps conforming to vehicle motion.

4 Results

Figure 7 illustrates the Root Mean Squared Error (RMSE) of the predicted accelerations by the six sensors, in the presence of an undamaged structure (in blue) and for a damaged structure (in red), reflecting Damage case 1 with the removal of a $0.9 \times 0.4\,\text{m}^2$ section from the flange. For each sensor and each train passage, in a similar way for both healthy and damaged scenarios, one can estimate the RMSE as

$$RMSE = \sqrt{\frac{\sum_{i=1}^{T} (output_i - target_i)^2}{T}} \tag{1}$$

where, for a certain instant i, $target_i$. is the expected acceleration in healthy condition of the bridge, $output_i$ is the acceleration predicted by the network and T is the time interval during which the accelerations were recorded. The train passages in the x axis are ordered by increasing speed and it is, thus, possible to observe a tendency for the error to increase with increasing train speed. Moreover, for the highest speeds (96–

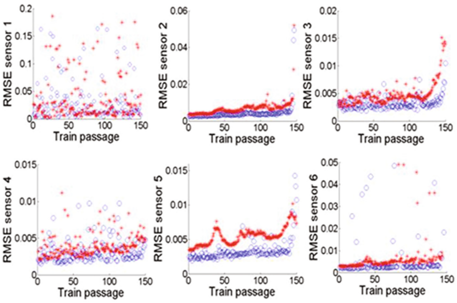

Fig. 7. RMSE against increasing speed of the train. Damage case 1: damage extension of $0.9 \times 0.4\,\mathrm{m}^2$. o Data from healthy structural condition; * Data from damaged structural condition.

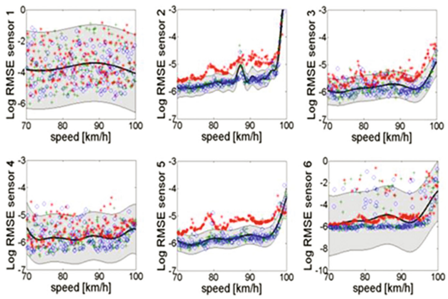

Fig. 8. GP fitted by prediction errors against increasing train speed. A Log-Normal Distribution of the error is considered. — Mean; Standard deviation; + Data to fit the GP; o Data from healthy condition; * Data from damaged condition, considering Damage Case 1 with a $0.9 \times 0.4\,\mathrm{m}^2$ section reduction.

100 km/h) it seems like even the response for healthy structural condition is poorly predicted, yielding large errors. This phenomenon can be explained by the fact that the maximum considered speeds excite the structure to frequencies close to its natural frequencies. As a consequence of that, the behavior of the structure may depart from its expected normal behavior and the trained network is not able to correctly predict the resulting accelerations anymore. Also noticeable is that the sensors positioned closer to the geometric middle of the bridge (sensors 2 and 5) seem to perform damage detection more efficiently, whilst sensors placed nearby the end supports (sensors 1 and 4) are not as efficient in the detection. This may be due to the fact that the dynamic response of the structure is more accentuated in the middle of the span than in its extremities and, accordingly, measurements registered in the middle of the span are expected to enable the network to make a clearer distinction between structural states.

It should be noted that the non-linear input-output relating function that the network uses can become quite complex and for that reason training an ANN that covers all the train load cases and speeds is extremely difficult. Fortunately, bridges are designed to be normally crossed by trains of the same configuration, very similar axle load and moving consistently within a limited range of speeds. Therefore the ANN is trained to predict accelerations only for those specific cases of speed and train type. In fact, the range of speeds (70–100 km/h) considered to train the network could actually be reduced, most likely yielding further accurate predictions of accelerations and, for example, reducing the disorder observed in the plots of sensors 1 and 4 in Fig. 7. In any case, even not making this adjustment, the results turned out to be very satisfactory.

Throughout the testing phase and after the errors in the predictions are determined, one has already qualitative indication that the network can successfully discriminate between structural states. However, making inferences based only on the plot from Fig. 7 would constitute a subjective way of judging what degree of separation is enough to suspect that damage is indeed present in the structure. Even if the bridge is found to be in good condition, the recorded dynamic response will be different for each train passage, as the magnitude of the response depends on the speed and axle load of the train plus other variations such as the operational and environmental settings. In short, this means that the prediction errors will oscillate for each train passage, even within the normal condition of the bridge. Hence there is a distribution of the errors that needs to be characterized stochastically and it is the errors that significantly deviate from this distribution that will work as a warning that damage in the structure may exist.

The prediction errors from 150 randomly selected train passages in healthy condition of the bridge are used to fit a statistical distribution that will work as a baseline for each sensor [10]. The Gaussian process (GP) [14] consists in assigning a normally distributed random variable to every point in some continuous domain. For each train speed the associated predicted errors are normally distributed and the mean and standard deviation of the error can differ between speeds. New data is then compared against the baseline: 150 other train passages in healthy condition and 150 in damaged condition. The idea is to compute discordancy measures for data and then compare the discordancy with a threshold, from which one is able to discriminate between healthy and damaged structural condition.

After the outcome of the prediction error is characterized by a GP (Fig. 8), damage detection can be performed by checking predictions that differ significantly from the expected values. A discordancy measure for normal condition data is the deviation statistic

$$z = \frac{|x_\xi - \bar{x}|}{\sigma_x} \tag{2}$$

where x_ξ is the candidate outlier and $\bar{x}$ and σ_x are, respectively, the mean and standard deviation of the data sample. The Mahalanobis distance [15] is one common measure of novelty in data and can be used in standard outlier analysis to provide a Damage Index (DI). To take into account only the train passages that give high prediction errors the distance to the mean is given in standard deviations and the error's differences, $(RMSE_n(v) - \mu_n(v))$, are signed. The DI for each train speed v is then defined as

$$DI(v) = \sum_{n=1}^{6} \frac{RMSE_n(v) - \mu_n(v)}{\sigma_n(v)} \tag{3}$$

where for each sensor n, $RMSE_n$ is the predicted error, μ_n is the mean predicted error and σ_n is the standard deviation of the predicted error. If the feature vector is related to undamaged condition, then $DI \approx 0$; otherwise, $DI \neq 0$. With the determined DIs for different train passages, the Receiver Operating Characteristic (ROC) curve can be constructed for each situation. Figure 9 depicts several ROC curves, each corresponding to a different damage severity within damage case 1. One point in the ROC space is considered better than another if it is associated to a higher *TPr* for the same *FPr*. Based on this conviction, as expected and in general, one verifies that more severe damage is related to a better detection. For example, for a fixed FPr of 8% we have

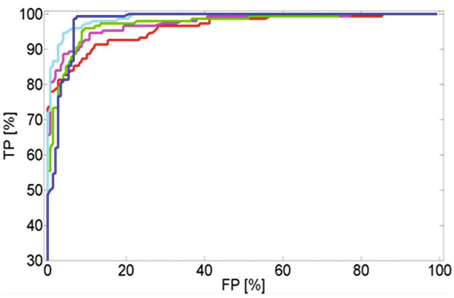

Fig. 9. ROC curves for different damage extensions l of Damage Case 1: damage resulting from cutting off a section of extension l from the bottom flange of one girder beam. – 20 cm; – 40 cm;– 70 cm; – 100 cm; – 160 cm.

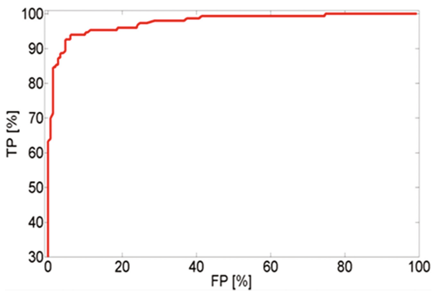

Fig. 10. ROC curve for DC2: damage resulting from a malfunctioning intermediate bracing.

associated 86, 90.7, 92, 96 and 99.3% TPr with damage severities of 20, 40, 70, 100 and 160 cm, respectively. Similarly, Fig. 10 depicts the ROC curve of the system in the presence of damage case 2. The improvement in detection capability with increasing damage is, however, not always verified as expected, since there are some operating points of DI that actually make the system more apt to detect damage in the presence of smaller rather than larger damage. For instance, we can see in Fig. 8 that the performance of the classifier for very low values of *FPr* is worse in the presence of the largest damage (160 cm), since it is related to a lower *TPr* when compared with any other smaller damage. The fact that some ROC curves respecting different amounts of damage intersect each other at certain thresholds, when ideally they should be separable over all the ROC space, may be partially explained by the fact that when the ROC curve is generated many assumptions and simplifications concerning parameters had to be previously made. Furthermore the process encompasses statistical analysis, thus yielding slightly different results every time it is repeated.

5 Conclusions

The novel methodology here proposed, based on the work presented in [10], provides a rational fashion for enhancing the damage diagnosis strategy for damaged structures, allowing for both improvements in safety and reduction of maintenance cost. The method proposes the use of past recorded deck accelerations in the bridge as input to an Artificial Neural Network that, after effectively trained, is able to predict an acceleration at a certain time in the future. The nonconformity between the measured value and the value predicted by the network will work as a primary indicator of damage. This study comprises the statistical evaluation of the prediction errors of the network by means of a Gaussian Process, after which one can select the optimal detection threshold.

From the attained results it is possible to derive some general conclusions:

- the outcome revealed to be noise sensitive, as expected, but the method seems to be robust and to perform well within typical levels of noise. It is a general conclusion that damage is more prone to be noticed in the presence of weaker noise and severer damage;
- lower vehicle speeds seem to overall provide measurements that enable better predictions by the trained network, in the sense that the prediction errors in both healthy and damaged structural condition are lower;
- the two sensors placed in the middle of the bridge seem to be the most efficient in the damage detection process, apparently disregarding where in the bridge damage takes place. This may be explained by the fact that the dynamic response of the bridge is more emphasized at half-span.
- ROC curves associated with scenarios where damage is more severe generally present a superior relation TP/FP, since for the same probability of TP one has to accept an inferior probability of FP when compared to less evident damage.

The method has although some weaknesses that can be tackled with additional research. This could concern the study of environmental and operational effects on the proposed damage detection method - other relevant parameters than accelerations may be given as input to the neural networks, such as temperature measurements. The consideration of these will most likely produce networks with higher accuracy in the prediction of the structural response, making the algorithm more shielded against the influence of other factors unrelated to damage that can induce significant changes in the behavior of the structure. The study presents a limited number of damage scenarios: a wider range of possible locations for damage in the bridge should be considered, including the impact of multiple damage scenarios, i.e. situations where damage events occur simultaneously in different parts of the structure. It would also be interesting to understand what are the limitations of the proposed method in terms of the smallest damage it can detect.

References

1. Huang, Y., Ludwig, S.A., Deng, F.: Sensor optimization using a genetic algorithm for structural health monitoring in harsh environments. J. Civ. Struct. Health Monit. **6**, 509–519 (2016)
2. Li, J., Zhang, X., Xing, J., Wang, P., Yang, Q., He, C.: Optimal sensor placement for long-span cable-stayed bridge using a novel particle swarm optimization algorithm. J. Civ. Struct. Health Monit. **5**(5), 677–685 (2015)
3. Yi, T.-H., Li, H.-N., Wang, C.-W.: Multiaxial sensor placement optimization in structural health monitoring using distributed wolf algorithm. Struct. Control Health Monit. **23**(4), 719–734 (2016)
4. Jin, C., Jang, S., Sun, X., Li, J., Christenson, R.: Damage detection of a highway bridge under severe temperature changes using extended Kalman filter trained neural network. J. Civ. Struct. Health Monit. **6**, 545–560 (2016)
5. Farrar, C.R., Worden, K.: Structural Health Monitoring. A Machine Learning Perspective. Wiley, New York (2013)

6. Rao, A.R.M., Lakshmi, K.: Damage diagnostic technique combining POD with time-frequency analysis and dynamic quantum PSO. Meccanica **50**(6), 1551–1578 (2015)
7. Zhou, Q., Zhou, H., Zhou, Q., Yang, F., Luo, L., Li, T.: Structural damage detection based on posteriori probability support vector machine and Dempster-Shafer evidence theory. Appl. Soft Comput. **36**, 368–374 (2015)
8. Figueiredo, E., Figueiras, J., Park, G., Farrar, C.R., Worden, K.: Influence of the Autoregressive Model Order on Damage Detection. Comput.-Aided Civ. Infrastruct. Eng. **26**(3), 225–238 (2011)
9. Deeb, M., Zabel, V.: The application of POD curves to damage detection in civil engineering structures – a numerical and experimental study. In: International Conference on Noise and Vibration Engineering ISMA 2012, Leuven, Belgium (2012)
10. Gonzalez, I., Karoumi, R.: BWIM aided damage detection in bridges using machine learning. J. Civ. Struct. Health Monit. **5**(5), 715–725 (2015)
11. Fawcett, T.: An introduction to ROC analysis. Pattern Recogn. Lett. **27**, 861–874 (2005)
12. Abaqus FEA: ABAQUS Inc. http://www.3ds.com/products-services/simulia/products/abaqus/. Accessed May 2017
13. White, K.: Bridge Maintenance Inspection and Evaluation. CRC Press, Boca Raton (1992)
14. Rasmussen, C., Williams, C.: Gaussian Processes for Machine Learning. The MIT Press, Cambridge (2006). ISBN 026218253X
15. Worden, K., Manson, G., Fieller, N.: Damage detection using outlier analysis. J. Sound Vib. **229**(3), 647–667 (2000)