FINAL REPORT:

Bridge Monitoring

PI: Jacobo Bielak

Research Team: George Lederman, Siheng Chen, Zihao Wang, Fernando Cerda, Piervincenzo Rizzo, James Garrett, Haeyoung Noh, Jelena Kovacevic, and Jacobo Bielak.

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Executive Summary

Already in the present, but even more so in the near future there will be computing power, communication devices (cell, WiFi, DSRC) and multiple sensors (speed, acceleration, GPS, cameras, radar, sonar, etc.) on standard vehicles. Their primary purpose is for safety and comfort, but with small additions these can be used to observe the condition of the infrastructure around the vehicle and report it to the responsible agencies. By continuously collecting this data from many vehicles throughout the road network the agency acquires a timely and comprehensive view of the status of its infrastructure at a low cost and plan the maintenance of it accordingly.

The main objective of the present project is to develop one such *indirect* structural health monitoring (SHM) approach, in which one makes use of vibration data collected from sensors installed on vehicles as they traverse the bridge or other infrastructure, rather than from sensors installed directly on the structure. Our work has proceeded in several stages. An initial exploratory study (funded by Traffic21) based entirely on mathematical models and computational simulations allowed us to demonstrate the feasibility of our indirect monitoring approach for these mathematical models and the desirability f testing this technique on physical models. Subsequent support from the National Science Foundation enabled us to conduct our first experimental study, based on a simple laboratory model of a bridge and a vehicle. Almost at the same time we also received funding from the UTC T-SET transportation center for the main purpose of exploring and developing different techniques for performing signal analysis and classification of the data collected during the course of the experiments. In this project, we have not inflicted actual damage to the bridge structure. Instead, we used small, added masses to the bridge deck, dampers, or varying support conditions as proxies. In addition, we considered temperature as a parameter since it is well known that bridge properties can vary significantly with temperature changes.

Most of our work in the project, especially the feature extraction and classification aspects, focused on data recorded during the experiments with the laboratory models. This allowed us to develop our detection capabilities in a controlled environment. In order to deal with more realistic systems, we selected two additional applications: (i) the 3rd level of the East Garage at CMU with a small robot as our vehicle. This set-up has the same general characteristics as the laboratory model, i.e., the structure tested consists essentially of a double-T beam with a deck, and the vehicle moves longitudinally along the beam; (ii) With the help of the UTC T-SET and the collaboration of the Port Authority of Allegheny County, we have begun to collect indirect data from Pittsburgh's Light Rail Line. This application offers the opportunity to continuously collect real data and to analyze and classify the data with the objective of detecting damage in an actual operational environment. We describe the various activities and results in the main body of the report. These are organized into tasks, starting with the laboratory model (Tasks 1, 2, 3) followed by the work on the more realistic structures (Task 4).

Task 1: Demonstrate our ability to identify the condition of a laboratory scale model from sensors placed on a model vehicle crossing over the bridge. This task looked at basic changes in the condition of the bridge, including the addition of extra mass, the addition of dampers or the modification of the boundary conditions.

Task 2 Consider changes in mass, damping and rotational restraint while subjecting the bridge to different temperature conditions. We addressed these issues by collecting additional data from our laboratory scale model, while experimenting with new algorithms for data analysis.

Task 3 (1) Quantify and localize damage on the bridge; (2) consider uncertainty about the condition of the bridge. (1) While previous work focused on classification of the state of the bridge between discrete groups, in this project we built a regression to consider an infinite number of classes with the output describing magnitude and position of the damage. To address part (2), we considered the same data sets as before but only labeled a small portion of the data. Here "labeling" refers to identification of the state of the bridge so the algorithm can learn. In real scenarios the condition of the bridge is often not known, or worse could be incorrectly labeled, so we found algorithms that could handle this uncertainty. We show that even with only 10% of the data labeled, our algorithms could learn how to label the unlabeled portions (label propagation) and achieve high accuracy.

Task 4 Collect data from operational structures, including a multistory parking garage on Carnegie Mellon's Pittsburgh Campus, and the bridges along Pittsburgh's Light Rail Line. This task forced us to employ new data collection techniques for handling uncontrolled environments and new data management techniques for storing much larger data sets.

The impact of this project has been two-fold. First we have advanced the state of the art for indirect bridge health monitoring, and have made such a low-cost technology more likely. Over the last two years, one journal paper has been accepted for publication, one has been submitted and is under review, and four peer-reviewed conference papers have been accepted. Second, this grant has helped train three civil engineering students to learn more about signal processing, while encouraging one signal processing graduate student to examine applications in infrastructure.

Contents

Executive Summary 3
List of Figures 7
List of Papers 9
Background 10
Problem statement 10
Proposed solution to the problem: an indirect approach 11
Task 1: Validation of the Indirect Approach with the Laboratory Scale Model 14
Experimental Setup and Protocol 15
Details of the Bridge and Vehicle Models 16
Motion Control and Data Acquisition Equipment 17
Protocol 17
Signal Analysis and Classification 19
Preprocessing 20
Feature extraction 20
Classification 23
Classification Results 24
Discussion 28
Task 2: Effect of temperature and boundary condition on the indirect health monitoring of a bridge model 30
Introduction 30
Experimental setup and protocol 31
Classification Experiments and Results 39
Effect on Classification Accuracy of Variations of Vehicles, Bridges, Vehicle Speeds and Severity of Change 39
Effect of Roughness and Damage Location 44
The Effect of Different Locations of Damage 47
Effect of Temperature gradient scenarios. 55
Discussion 56
Task 3: Damage Localization, Quantification and Bridge Condition Uncertainty 59
Data Analysis for SHM 59
Classification System 59
Progress 59

Feature Extractor 60
Classifier 60
Multiresolution Framework 61
Summary 61
Comparison 62
Conclusions 63
Task 4: Garage Experiments and Light Rail Data Collection 64
Garage Field Experiments 64
Laboratory Experiments 67
Data Selection 67
Experimental Results 70
Port Authority Project 75
Conclusions 80
Bibliography 81
Appendix ..... A

List of Figures

Figure 1: Direct and indirect SHM different in terms of sensor location. 11
Figure 2: Experimental setup. 16
Figure 3: Experimental vehicle CAD drawings. 17
Figure 4: Roller support - Rotational restraint 18
Figure 5: General classification system. 20
Figure 6: Time-domain signals 20
Figure 7: Discrete Fourier transform of the signal. 21
Figure 8: Mean energy distribution (normalized to unit energy). 21
Figure 9: Discriminant power (normalized to unit discriminant power). 22
Figure 10: Summary of feature extraction process. 23
Figure 11: Clustering of two scenarios. 23
Figure 12: Classification accuracy results (a) Average by damage type and (b) by speed. 25
Figure 13: Classification results for rotational restraint scenarios. 26
Figure 14: Classification results for damping increase scenarios. 27
Figure 15: Classification results for concentrated mass scenarios. 28
Figure 16: Experimental setup. General photo and Vehicle picture. 31
Figure 17: Introduced modifications to the bridge structure. 32
Figure 18: Bridge longitudinal view; normal and Infrared images. 37
Figure 19: B1: Damage detection for different vehicles, speeds, and severity scenarios 40
Figure 20: B2: Damage detection for different vehicles, speeds, and severity scenarios. 41
Figure 21: B1: Damage severity classification accuracy for different vehicles 42
Figure 22: B2: Damage severity classification accuracy for different vehicles 43
Figure 23: Rotational restraints scenarios cluster analysis. 43
Figure 24: Damage detection comparison among bridges with different roughness profiles. 46
Figure 25: Damage detection comparison among scenarios. 50
Figure 26: Location classification comparison among scenarios 55
Figure 27: Classification accuracy of damage scenarios for specific temperature scenario. 56
Figure 28: Classification accuracy of damage scenarios regardless temperature scenario 56
Figure 29 – Comparision of three methods, Fourier Discriminant Features, Local Discriminant Features and Sparse Representation for 18 scenarios 62
Figure 30 – Comparison of Fourier Discriminant Features and Multiresolution Classification .. 62
Figure 31 – Discriminant Accuracy at Various Labeling Ratios 63
Figure 32 – Experimental Setup for Field Tests 65
Figure 33 – Illustration Of Garage Experiments 65
Figure 34 – Detail of the Robot 66
Figure 35 – Slight Gap between Track and Bridge Allowing for Data Selection 67
Figure 36 – Data Selection Lab Data 68
Figure 37 – Detail of Broom Handle 68

Figure 38 – Garage Experiment Data Selection 69
Figure 39 – Waterfall Plot, Data from Back Wheel, Laboratory Experiments..... 70
Figure 40 – Waterfall Plot, Data from Back Chassis..... 71
Figure 41 – Waterfall Plot, Data from Bridge Sensor, Laboratory Experiments 71
Figure 42 – Waterfall Plot, Signals from the Robot 72
Figure 43 – Frequency Domain Waterfall Plot Deck Sensor 73
Figure 44 – Frequency Domain Waterfall plot. 73
Figure 45 – Frequency Domain Waterfall Plot Robot. 74
Figure 46 – Detail of Lower Frequencies from Robot..... 74
Figure 47 – Installed Equipment on the Train 75
Figure 48 –Schematic of the Sensors Placed on the light rail car..... 76
Figure 49 – GPS Plots..... 76
Figure 50 – Raw Train Signal. 77
Figure 51 - Example Signals over a bridge at Differnt Temperatures 78

List of Papers

Journal Papers

1. F. Cerda, S. Chen, J. Bielak, J. H. Garrett, P. Rizzo, and J. Kovačević, "Indirect structural health monitoring of a simplified laboratory-scale bridge mode," Int. J. Smart Struct. Syst., Challenge on bridge health monitoring utilizing vehicle-induced vibrations, 2013. To appear……**A1**

2. S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring," IEEE Trans. Signal Process., June 2013, Submitted…….....**A17**

Conference Papers

3. S. Chen, A. Sandryhaila, G. Lederman, Z. Wang, J.M.F. Moura, P. Rizzo, J. Bielak, J. Garrett, and J. Kovačević. "Signal Inpainting on Graphs via Total Variation Minimization." Proc. IEEE Int. Conf. Acoust., Speech Signal Process., Florence, Italy, May 4-9th, 2014. Submitted…....**A30**

4. G. Lederman, Z. Wang, J. Bielak, H. Noh, J. H. Garrett, S. Chen, J. Kovačević, F. Cerda, and P. Rizzo, "Damage quantification and localization algorithms for indirect SHM of bridges," Proc. Int. Conf. Bridge Maint., Safety Manag., Shanghai, China, July 2014. Accepted……..…..**A35**

5. S. Chen, A. Sandryhaila, J.M.F. Moura and J. Kovačević, "Adaptive Graph Filtering: Multiresolution Classification on Graphs" Proc. IEEE Glob. Conf. Signal Information Processing, Austin, TX, Dec. 2013…….**A43**

6. Z. Wang, S. Chen, G.Lederman, F. Cerda, J. Bielak, J. H. Garrett, P. Rizzo and J. Kovačević, "Comparison of sparse representation and Fourier discriminant methods: Damage location classification in indirect lab-scale bridge structural health monitoring," Proc. Structures Congr., Pittsburgh, PA, May 2013……..……..**A47**
7. S. Chen, F. Cerda, J. Guo, J. B. Harley, Q. Shi, P. Rizzo, J. Bielak, J. H. Garrett and J. Kovačević, "Multiresolution classification with semi-supervised learning for indirect bridge structure health monitoring," Proc. IEEE Int. Conf. Acoust., Speech Signal Process., Vancouver, Canada, May 2013, pp. 3412-3416…….……**A58**

8. F. Cerda, J. Garrett, J. Bielak, P. Rizzo, J. A. Barrera, Z. Zhang, S. Chen, M. McCann, and J. Kovačević, "Indirect structural health monitoring in bridges: scale experiments," Proc. Int. Conf. Bridge Maint., Safety Manag., Lago di Como, Italy, Jul. 2012……..**A63**

9. F. Cerda, J. Garrett, J. Bielak, R. Bhagavatula, and J. Kovačević, "Exploring indirect vehicle-bridge interaction for bridge SHM," Proc. International Conference Bridge Maintenance, Safety Management, Philadelphia, PA, Jul. 2010……..**A71**

Background[1]

Bridge structural health monitoring (SHM) has been an active research field for over 30 years. It is a multidisciplinary problem that involves: sensors, data acquisition systems, data analysis and data interpretation. In general terms, the objective of bridge SHM is the early detection and characterization of damage conditions before they pose a threat to the structural integrity of a bridge and the public it serves. An analogy can be made with the medical field where the early detection of pathology allows for several treatment options to be undertaken to restore health. A late diagnosis might lead to chronic illness or death. If diagnosed properly, prognosis and treatment follow diagnosis.

Problem statement

The current state of the aging bridge infrastructure in the USA and around the world requires more accurate diagnostic tools for a large stock of bridges. There has been an increase in public, as well as political, awareness of the current state of the bridge infrastructure. Old and new bridges are now under public scrutiny after catastrophic collapses. The collapse of the I-35 bridge over the Mississippi River on Aug. 1, 2007 is a noticeable example, as well as the collapse of other bridges around the globe (e.g., Shershan Bridge, Pakistan, Sep. 1 2007; Harp Road Bridge, USA, Aug 15 2007; Loncomilla Bridge, Chile, Nov 18, 2004.).

The research community has been developing structural health monitoring (SHM) techniques to aid in the ongoing bridge management efforts of local bridge authorities. The current standard bridge inspection practice is based on biannual visual inspections, which are subjective by nature. Sensor-based SHM is perceived as the technology that could improve the current visual inspection process (FHWA-2001). Monitoring bridge structural systems helps in planning different bridge intervention strategies, such as maintenance actions, repair or replacement (Frangopol et al. 2008). Moreover, the life-span of the bridge structure can be extended (even if the bridge shows deterioration) if the data shows it to be healthy.

Traditional bridge SHM techniques entail the placement of sensors on the structure for measuring physical parameters that are then used as indicators of the structural behavior. SHM and damage assessment have been very active research areas, and have motivated several excellent review and overview papers, which highlight some of the most relevant approaches (e.g., Van der Auweraer and Peeters 2003; Farrar and Worden 2007). Brownjohn (2006) describes some general and fundamental objectives for monitoring civil infrastructure and points out some historical applications. More specific review topics include wireless, structural health monitoring, design of devices, and the trend for localized processing (Lynch, 2007); vibration-based condition monitoring (Doebling et al. 1998, Carden and Fanning 2004); damage identification using inverse methods (Friswell 2006); unsupervised learning (Fulgate et al. 2000, Worden and Dulieu-Barton 2004, Worden and Manson 2007); and vibration-based condition monitoring methods (Carden and Fanning 2004).

Two main approaches of this type have been pursued in recent years. One is a global vibrational approach and the other is the local approach (Mal et al 2005). The first looks at the vibration of the whole structure, while the latter focuses on the wave propagation along structural elements. We refer to these two approaches based on sensors placed on the structure as *direct* approaches. The direct approaches are especially useful when monitoring progression of damage of a particular known damage condition, or monitoring a critical member of a bridge structure.

[1] This background section is based around from Cerda's PhD dissertation.

However, the transition of the traditional SHM techniques from the research community to the practical field implementation still needs to overcome difficult challenges due mainly to technical and economic considerations (Farhey 2005, 2007). One such real world application of a monitoring program is on the Lehigh River Bridge, SR-33. It was reported that during the three year monitoring program, the direct cost of the equipment made up 40% of the total monitoring cost. For a long term monitoring program, this direct cost must be addressed several times during the lifetime of the structure, due to software or equipment obsolescence. The other 60% of the cost is related to labor and other related items (Frangopol et al. 2008). Some of the challenges of the *direct* approach include: labor intensive sensor deployment and maintenance, sensor system cost, powering the sensors system, data transmission and data interpretation.

Some of these challenges are currently being addressed by the research community. With respect to the power supply of the sensor systems, a popular approach for direct monitoring is energy harvesting to make the system self-reliant. Among the power options, one can consider solar energy (Alippi and Galperti 2008) and vibration-based approaches (Beeby et al. 2006). However, other challenges, such as the lifespan of the sensor systems with respect to the lifespan of the sensed structure and deployment and maintenance costs, are unlikely to be solved with the same approach. Electrical systems and electronics are far more vulnerable to ambient conditions than structural elements and therefore less reliable than the system they are sensing. The threat of atmospheric conditions can also be extended to involuntary damage and vandalism. Depending on the socio-economic condition of a particular country, leaving electronic equipment unsupervised has to be carefully planned for and designed to prevent theft of the equipment.

The direct approach remains impractical to this day as a first diagnostic scan, especially when one considers the large bridge population of buildings that needs to be inspected (e.g. 600,000 bridges over 20 ft.) in the US bridge inventory.

Proposed solution to the problem: an indirect approach

The need for more efficient techniques for the SHM of bridges has led to the development of additional approaches to the direct method. One such approach is the *indirect SHM,* so-called because it makes use of vibration data collected from sensors installed on vehicles crossing the bridge, rather than from sensors installed on the bridge itself (Lin and Yang 2005).

Figure 1 depicts the direct and the indirect SHM approaches. In principle, the vehicle-bridge interaction data captured by the moving vehicle provides information about the bridge, which can be used for diagnostic purposes. In other words, the indirect approach uses vehicle vibration data $q_{vi}(t)$ for diagnosing a bridge's condition rather than data $q_b(t)$ collected directly from the bridge. The use of signal processing techniques and machine learning algorithms allow one to explore the feature space of the signals collected from the vehicle and extracting vehicle-response patterns for bridge SHM. The indirect SHM approach was conceived only in the last decade, and is currently being actively researched.

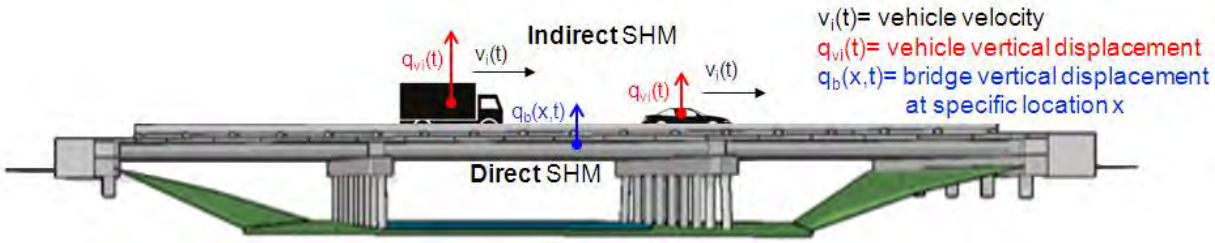

Figure 1: Direct and indirect SHM different in terms of sensor location.

The indirect approach can be viewed as being complementary to the *direct approach*. However, it has several conceptual advantages over the direct approach. The *indirect approach* allows the interrogation of a large bridge stock; it can be powered by the vehicle electrical system; it has no need to stop traffic for initial instrumentation of the structure or on-site maintenance actions; it can leverage the fleets of smart vehicles that communicate among themselves for other purposes such as: collision avoidance, traffic optimization, emergency localization or autonomous driving. The indirect approach is a data-based approach that can capture patterns of complex phenomena from large amounts of data. It takes advantage of advances in signal processing and pattern recognition algorithms.

Considering the practical future implementation of this approach, two possible scenarios might be feasible to obtain instrumented vehicle data. The first scenario that could be implemented over the medium-term considers the instrumentation of a specific fleet of vehicles that interrogate bridge structures as they go about their daily business. In this scenario, possible candidate fleets would be utility trucks or other type of public vehicles, such as transportation buses or mail trucks. They would be equipped with accelerometers, GPS, data acquisition systems and data transmission systems in order to determine their precise location and to collect their dynamic interaction data. The second scenario can be considered as a long-term vision. In the future, commercial vehicles will be readily equipped with a great number of sensors for different purposes. Some of the instrumentation that vehicles already have today are tire pressure sensors, accelerometers for adaptive suspension systems, GPS, front and rear cameras, internet connection and onboard computers. Akinci et al. envisioned taking full advantage of the opportunities that arise from having on-line vehicles making their data accessible for public interest (Akinci et al. 2003).

In the following paragraphs, we briefly review recent research efforts related to indirect SHM.

The indirect SHM approach was introduced by Yang et al. (2004), to extract the fundamental frequencies of a bridge. Yang and his colleagues derived a closed-form solution for a single-degree-of-freedom oscillator moving over a single-span, simply supported beam by assuming that the beam vibrates only in its fundamental mode. The obtained solution allowed for the identification of the two main dimensionless parameters that affected the bridge response. These parameters were S and μ, $S = \pi v/L\omega_b$, a normalized vehicle velocity, where v = vehicle velocity, L = length of beam, and ω_b = bridge's natural fundamental frequency; and $\mu = \omega_b/\omega_v$, where ω_v is the vehicle (oscillator) vertical natural frequency. Yang et al. (2005) then expanded the closed-form solution presented in 2004, to include several mode shapes as the basis for the dynamic response of the beam.

An experimental validation of the vehicle-based approach for extracting the natural frequencies of a bridge was conducted using an instrumented two-wheeled cart attached to a vehicle traveling over a simply-supported girder bridge; a heavy load truck was used to act as oncoming traffic (Lin and Yang 2005). The authors were able to identify the fundamental frequency of the bridge from the cart data, even with the simulated oncoming traffic. Toshinami et al (2010) also aimed at extracting bridge frequencies from vehicle response.

McGetrick et al. (2009) developed a numerical 1D model and studied the variations of dynamic parameters using the data derived from the numerical model. The authors report high sensitivity in the peak magnitude frequencies of the acceleration power spectral density to slight changes in the bridge structural damping. The same authors used laboratory model data (McGetrick et al. 2010) to identify the predominant frequencies of bridges and found good

agreement between the signals acquired directly from the bridge and those from the vehicle as the vehicle travels over the bridge.

Kim and Kawatani (2008) presented an approach for accurately identifying damage on a bridge structure. This approach requires data from both the vehicle and the bridge. It was first explored with numerically simulated data and later through experiments (Kim et al. 2010). A hypothesis-testing scheme that looks for patterns in the bridge response from a laboratory experiment identified successfully only severe damage conditions on the bridge structure (Isemoto et al 2010).

Based on these studies, it is reasonable to conclude that vehicle responses collected from a vehicle while travelling over a bridge contain useful information of the structural condition of the bridge. However, past studies of the indirect approach concentrated mainly on identifying certain interaction properties and indexes, such as the fundamental frequency of the bridge, power spectral density magnitude variations from vehicle data, the agreement between vehicle and bridge data predominant frequencies or the identification of severe structural changes. Thus, there is a need to complement the ongoing efforts with more robust signal processing and machine learning techniques that will allow one to detect small variations in the vehicle signature response and, therefore, allow exploration of the use of the indirect approach for diagnostic purposes.

The main objective of this work is to explore the use of the indirect method for developing a practical bridge SHM approach enhanced by signal processing and pattern recognition.

Task 1: Validation of the Indirect Approach with the Laboratory Scale Model[2]

We hypothesize that an array of sensors, mounted on moving vehicles that travel across the bridge of interest, can be helpful in identifying structural damage and thus serve as an indicator for more detailed analysis using a physical model. This approach is referred to as *indirect health monitoring.*

Yang and Chang (2009) reported results associated with field experiments where the first two natural frequencies of a bridge were extracted from the vehicle response by using empirical mode decomposition. Bu et al. (2006) measured the dynamic response of a vehicle moving on top of a simply supported Euler–Bernoulli beam. The vehicle served as a sensor and force transducer to detect damage defined in terms of the reduction of flexural stiffness. The model incorporated noise measurements, road surface roughness, and model errors such as underestimating vehicle parameters or bridge flexural stiffness.

Kim and Kawatani (2008) developed a pseudo-static damage detection method that makes use of the coupled vibration of a vehicle-bridge system. It requires data collected from both the bridge and the vehicle to characterize the damage. A numerical model that included the roadway roughness effect was used to test the approach. It was subsequently validated experimentally for different vehicle speeds and different amounts of reduction of the moment of inertia of the girders. McGetrick et al. (2009) modeled a simplified quarter car-bridge interaction to extract the fundamental natural frequency and corresponding damping of the bridge from the spectra of the vehicle accelerations. They found that better accuracy was achieved at lower speeds and smoother road profiles. Moreover, the magnitude of the acceleration power spectral density's peaks decreased with increasing bridge damping and this decrease was easier to detect with a smoother road profile. This work was validated experimentally by observing the effects of a vehicle moving across a steel girder that included a road surface profile. The effects of varying vehicle model mass and speed were investigated as well (McGetrick et al. 2010).

Isemoto et al. (2010) developed a hypothesis-testing scheme for damage detection based on the bridge vertical acceleration data induced by a passing vehicle. An experimental vehicle-bridge model, including roadway roughness, was used and only severe damage scenarios were identified. Miyamoto and Yabe (2011) exploited the vibration induced by a public bus for the indirect health monitoring of existing short- and medium-span reinforced/prestressed concrete bridges. The tests demonstrated a correlation between the vehicle vertical acceleration and the bridge vibration at midspan. By means of a numerical 3D finite element model, the distribution of characteristic deflection values was found for a particular driving speed and two severe damage scenarios.

Yin and Tang (2011) proposed a finite-element method to simulate the interaction of a vehicle and a cable-stayed bridge. The vertical displacement from the vehicle was used to identify tension loss and deck damage. The relative displacement of a passing vehicle of a bridge with known damaged conditions is used to generate a vector basis. The proper orthogonal decomposition on the relative displacement of a vehicle passing a bridge with an unknown damage condition is optimized with the known basis, and parameters of the unknown damaged bridges are reconstructed. Finally, Sirigoringo and Fujino (2012) proposed an indirect approach to estimate the fundamental natural frequency of a bridge using the response of a passing instrumented vehicle. The method was validated experimentally on a full-scale simply supported

[2] Task 1 section is based on Cerda's PhD dissertation.

short span bridge by using a light commercial vehicle instrumented with accelerometers. The spectra of the vehicle's dynamic responses while crossing the bridge were analyzed to reveal the first natural frequency of the bridge when the vehicle moved with constant velocity. The experimental study considered traveling speed ranging from 10 to 30 m/s.

All of the above indirect monitoring studies aimed at identifying the dynamic parameters of a bridge by either observing a single damage-sensitive feature or by optimizing model-based damaged sensitive vectors. The authors of these studies, however, did not report on detecting various types of damage conditions, the confidence with which the detection was ascertained, or the effects of different boundary conditions, weather patterns, traffic loads, or different vehicle speeds on the structural diagnosis.

Results are presented from a study that aims at bridging this gap. The long-term objective is to create a decentralized monitoring approach using fleets of vehicles that can continuously store or send data about the bridges over which they travel. In the laboratory, a bridge model is subjected to different vehicle speeds, damage scenarios, and structural boundary conditions. In addition to the indirect measurement of the bridge motion obtained through the vehicle vibration, other three sensors were placed directly on the bridge. In contrast to previous studies, the indirect and the direct data were used independently. This allows comparing the indirect and the direct strategies and to evaluate the effectiveness of our indirect damage detection algorithm.

Experimental Setup and Protocol

A laboratory experimental setup was built to collect data from a vehicle, which could be used to detect changes in the condition of the bridge. Using this model, acceleration data from a bridge structure and a vehicle passing over it can be collected and later analyzed for characterizing vehicle-bridge interaction patterns. The complete experimental setup consists of mechanical components that make up the bridge and vehicle system, a vehicle motion control system, and data acquisition equipment. The different mechanical components resemble a simply supported bridge structure and a four-wheeled vehicle with an independent suspension system at each wheel. The motion control equipment is able to move the vehicle over an acceleration ramp, a bridge and a deceleration ramp, causing the vehicle to reach a target speed before the end of the acceleration ramp, then maintain that velocity over the bridge, followed by a deceleration of the vehicle so that it stops at the end of the deceleration ramp. The data acquisition system records accelerations at different locations on the vehicle and the bridge, as well as the position of the vehicle. This experimental setup was inspired by the work of Kim et al. (2010).

An overview of the setup is shown in Figure 2. The vehicle, approximately in the middle of the figure, is pulled by a belt system. The vehicle was instrumented with Vibra-metrics accelerometers (Model 5102) powered by cables supported by a cable delivery system that moves parallel to the vehicle. The cables and the vehicle are propelled by a motor at the leftmost part of the experimental setup.

The travelling path of the vehicle corresponds to the acceleration/deceleration ramps and the bridge as labeled in Figure 2a. The simply supported bridge structure is in the middle of the travelling path. Below the simply supported bridge there are two reaction beams that are used to support the bridge reactions and the added dampers.

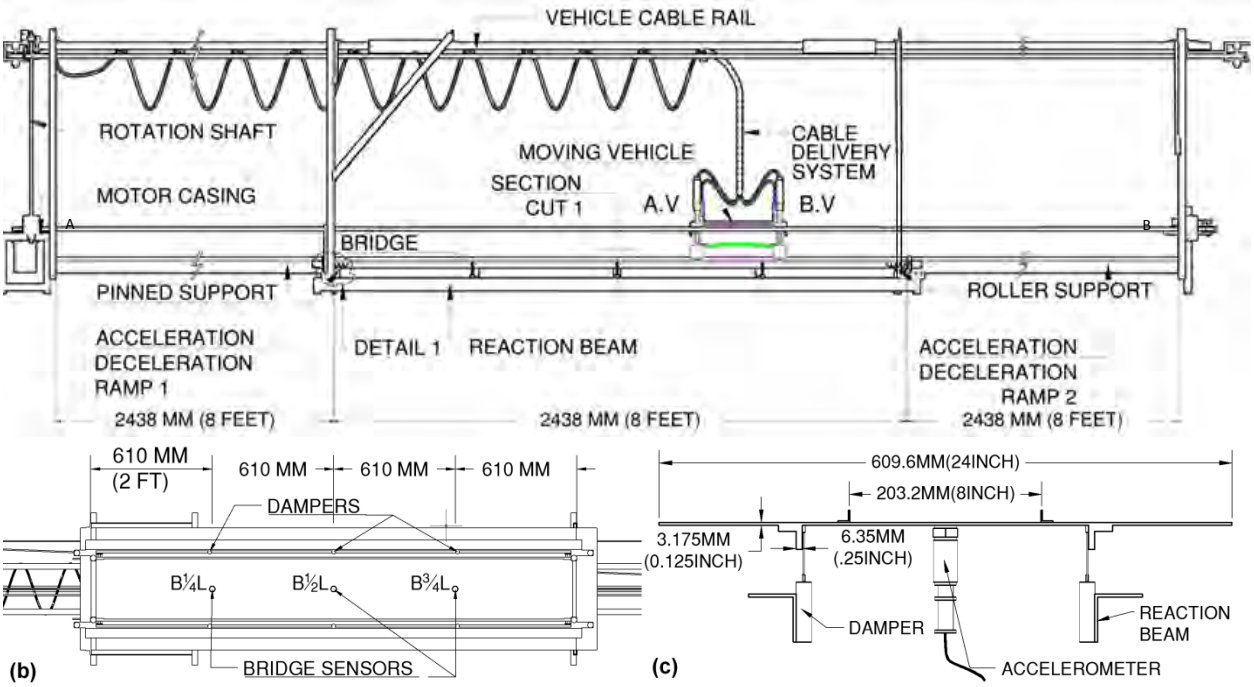

Figure 2: Experimental setup.
a) General lateral view. b) Bridge bottom view. c) Bridge cross section.

Details of the Bridge and Vehicle Models

The bridge is simply supported by a roller support at the left and a pinned support at the right. The vehicle enters the bridge from the left. The whole system is constructed to act as a closed force loop system. The longitudinal forces generated by the motor to move the vehicle are transmitted between the two supports by two connecting beams labeled as "reaction beams" in Figure 2a. The bridge structure is instrumented with three 5102 Vibra-metrics accelerometers as in Figure 2b. The sensors are equally spaced along the longitudinal direction of the bridge and named accordingly as B1/4L, B1/2L and B3/4L where L=2438 mm. The reaction beams act as a support for localized dampers that connect to the bridge structure as in Figure c. The bridge deck consists of an aluminum plate, and two angle beams act as the bridge girders. On top of the plate two angle beams serve as rails for the travel path of the vehicle. Detailed dimensions of the bridge section are shown in Figure 2c.

The bridge has a total mass of 18.3 kg, a fundamental natural frequency of 7.23 Hz, a fraction of critical damping of 3.6 percent. These are the properties of the bridge in the pristine condition, later referred to as Scenario 1.

Figure 3a shows a 3D view of the vehicle constructed for the experimental setup with the main components labeled. The vehicle was instrumented with two accelerometers connected to the suspension shafts in order to record the acceleration at the wheel level and with two accelerometers placed on the suspension to acquire data filtered by the suspension system. To keep the symmetry of the vehicle, two calibrated weights were placed on top of the un-sensed wheel shafts. Similarly to the bridge structure, the vehicle was built mainly with aluminum parts.

Two reference points are labeled on the longitudinal direction of the vehicle as points A Vehicle (A.V) and B Vehicle (B.V). A top view of the vehicle is shown in Figure 3b. The length and width of the vehicle as well as the labels assigned to the four sensors are also indicated. The

sensor labels are defined by their position and location to the reference point. Suspension A.V and Suspension B.V are labeled S.A.V and S.B.V, respectively, and the two wheel level sensor locations are labeled W.A.V and W.B.V. A picture of the model vehicle is shown in Figure 3c.

The frequency of the A.V and B.V axles of the vehicle were determined through free vibration experiments on the suspension. These experiments were performed on the vehicle when separated from the belt of the vehicle motion control system. The frequencies were obtained by averaging the power spectra of five free vibration experiments; the results are summarized in Table 2.

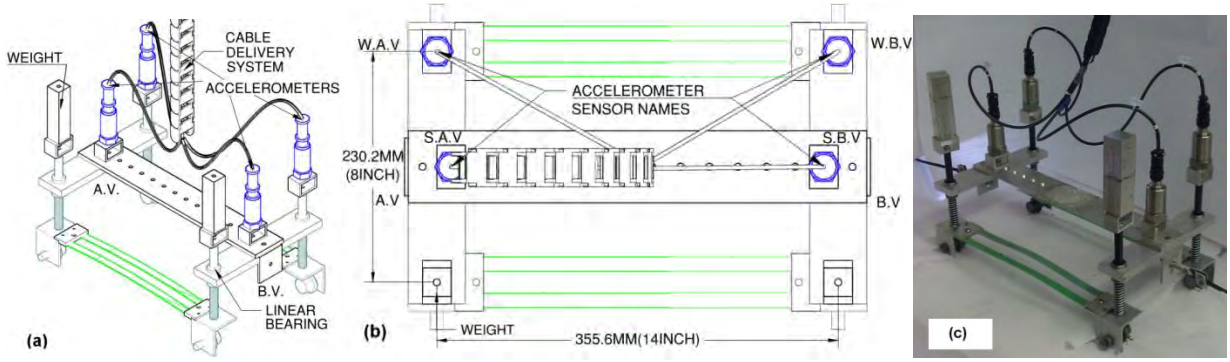

Figure 3: Experimental vehicle CAD drawings. (a) Parametric view. (b) Top view. (c) Vehicle picture.

Table 1: Vehicle properties.

Vehicle weight [kg]	4.8
A.V axle frequency[Hz]	5.0
B.V axle frequency[Hz]	5.5

Motion Control and Data Acquisition Equipment

A National Instruments® PXI system running in LabView® was assembled to operate the instrumented vehicle and to allow for data acquisition and storage. The system consisted of a PXI Chassis (NI PXI 1031) with a motion control card (NI PXI 7342), a motion interface (UMI 7772), a stepper drive (P70360) and a dual shaft stepper motor (NEMA 34). A feedback loop for position was achieved with an encoder. The acceleration data were digitized and stored for post-processing using two digitizers (NI 9234).

Protocol

Three different types of "damage" scenarios were designed: 1) variations on the support condition by imposing rotational restraints, 2) increase of damping at different locations, and 3) a mass increase at the midspan. For each kind, four levels of severity were devised in order to obtain a total of 12 different damage scenarios. Table 2 shows the twelve conditions of damage (thirteen in total). For each case the resonance frequency and the critical damping are reported and compared to the baseline, i.e. Scenario 1. For all cases, the fundamental natural frequency of vibration and the damping coefficient are determined by means of conventional free-vibration experiments.

The rotational restraint mechanism was built into each of the four beam supports of the bridge model. As shown in Figure 4a, an aluminum bar (12x1x1/8 inch) was attached to the main

girder of the bridge at one extreme and connected to the support at the other extreme to provide vertical restraint. The plate was drilled down to 6.35 mm [1/4 in] with 12.7 mm radius [1/2 in] to provide only a partial restraint.

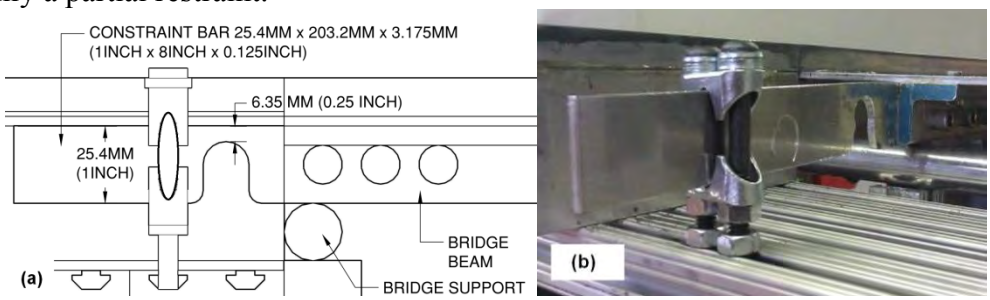

Figure 4: Detail 1 – Roller support - Rotational restraint (a) technical drawing and (b) picture.

Variation of the rotational restraints simulates the case of rubber bearings becoming stiffer in time or steel corrosion occurring on rocker supports. This condition is denoted as frozen bearings and it is a common cause of undesired stress in the structure, and, therefore, a reduction in the load capacity. In Table 2, the variations of the rotational restraints are described as scenarios 2 to SC5. In Scenarios 2-5, one, two, three and all four supports are restrained, respectively. As expected, the greater the number of rotational restraints, the higher the fundamental frequency, which provides an indicator of change to the bridge structure.

The variation of localized damping is achieved by adding dampers to the bridge at locations schematized in Table 2 (Scenarios 6 to 9). A set of AIRPOT adjustable dampers were used. They were calibrated to provide the same damping coefficient. In Scenarios 6-9, one, two, four and six dampers, respectively were attached to the bridge structure as depicted in the schematics of Table 2.

Table 2: Damage scenarios.

Schematics	SC	f [Hz]	% f shift	$\%\zeta_{crit}$	% ζ shift
	1	7.23	-	3.63	-
	2	7.46	3.17	6.34	74.9
	3	7.66	6.00	6.45	77.8
	4	8.11	12.2	7.97	119
	5	8.56	18.4	9.43	158
	6	7.24	0.17	8.52	134
	7	7.25	0.28	11.3	211
	8	7.28	0.73	26.4	629
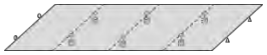	9	7.30	0.98	31.4	766
	10	7.19	0.56	4.44	22.4
	11	7.18	0.66	4.34	19.5
	12	7.14	1.29	4.07	12.2
	13	7.09	1.85	4.37	20.5

Finally, the concentrated mass at the midspan of the structure consisted of weights equal to 50 g, 100 g , 200g and 300 g. As expected, the presence of the mass decreased the fundamental frequency of vibration of the structure.

Eight different vehicle speeds, varying from 1 m/s to 2.75 m/s, were considered for each damage scenario.

Signal Analysis and Classification

In general terms, the task of distinguishing various bridge conditions is a signal-processing task of classification. The classification task is first described in general, and then we explain how it was used in our setting. Assume a real signal x of length N, i.e., $x \in R^N$ (see Original signal in Figure 5). The problem, then, can be formulated as that of designing a map from the signal space of vibrational signals $X \subset R^N$ to a response space of class labels $Y \subset \{1,2,...,C\}$ (in Figure 4 these are Damaged and Pristine labels). That is, the decision $d: X \rightarrow Y$ is the map that associates an input signal with a class label.

A general classification system consists of a feature extractor and a classifier (see Figure 5). Since the dimensionality of the input space is typically large, the feature extractor is introduced to reduce this dimensionality by setting up a feature space $F \subset R^k$ where $k \leq N$ between the

input space and the response space. The feature extractor is the map defined as $f : X \to F$ and the classifier as a map $g : F \to Y$.

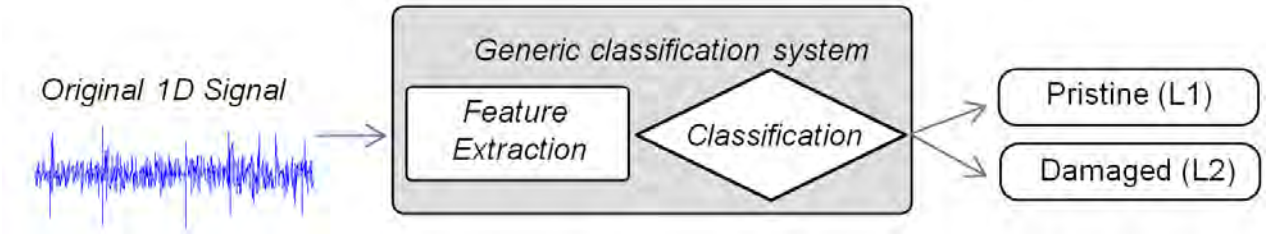

Figure 5: General classification system.

Preprocessing

Figure 6 shows the signal obtained from the vibration of the vehicle, from the time when it starts moving, through its motion across the bridge, until it is brought to a stop. The only relevant information for the bridge characterization, however, is that of the vehicle moving across the bridge. The reference start time was chosen as the moment when the rear wheels enter the bridge and as the end time the instant when the front wheels exit the bridge. That portion of the signal is highlighted in Figure 6a-d by the two vertical lines. The extracted portions of the signals are then normalized to have zero mean and unit variance.

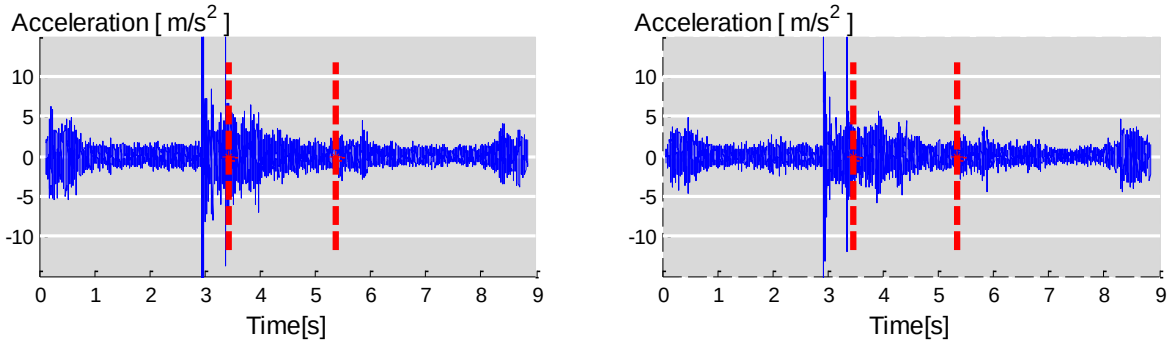

(a) Sensor B ½ L, Pristine.

(b) Sensor B ½ L, Damaged.

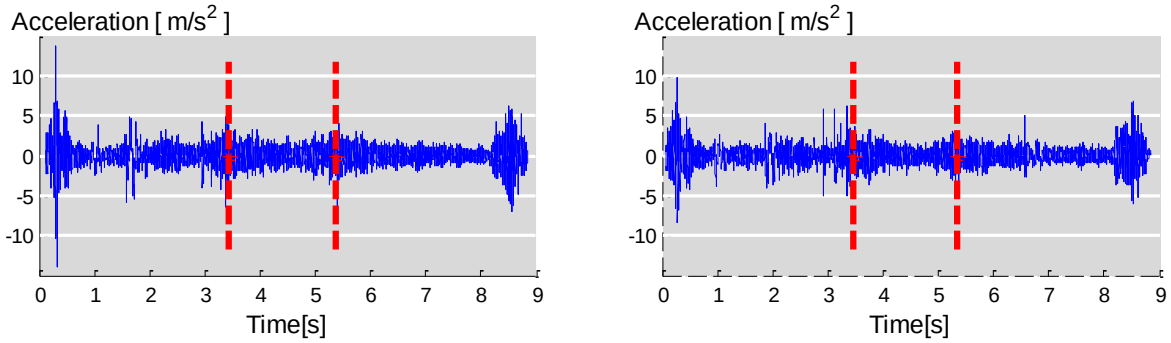

(c) Sensor WBV, Pristine.

(d) Sensor WBV, Damaged.

Figure 6: Time-domain signals.

Feature extraction

A linear structural system can be characterized in the frequency domain by its predominant natural frequencies, and their corresponding mode shapes and damping values. We explore therefore the use of frequency spectra characteristics for damage detection.

Considering the fundamental natural frequencies of the damage scenarios and the vehicle main bouncing frequency, we limited the analysis frequency spectrum to up to 33 Hz. For example, looking at the spectrum of the signal in Figure 7 for signals from two different scenarios, a pristine and a damaged one, we see that potentially we could tell responses apart from separate sensors by looking at magnitudes at certain characteristic frequencies. We thus decided to use frequencies as features, hoping to distinguish among different scenarios. Our task is then to find a set of features to maximize differentiation between classes.

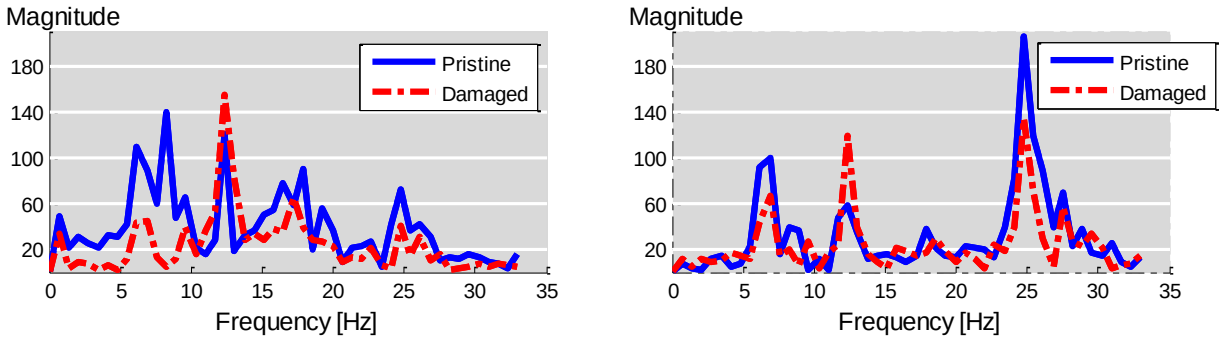

(a) Two scenarios, sensor B ½ L. (b) Two scenarios, sensor W.B.V.

Figure 7: Discrete Fourier transform of the signal.

Since the acceleration signal contains a large number of spikes and other transient signals, the spectra are noisy with little consistency between runs. To reduce noise and keep non-transient frequencies of interest, a typical approach is to average the spectra across frequency. Averaging the frequency spectra is a well-known technique used in noisy signal processing. For example, on radar signal analysis, a redundant number of antennas capture noisy signals from the same source and average them to increase the signal to noise ratio.

After averaging, we calculate the frequency-domain energy distribution for each scenario. This technique relies on the assumption that each scenario has its unique energy distribution in the frequency domain. Since we wish to tell classes apart and not individual runs, we average all the energy distributions from the same class and use the mean energy distribution as the representative member of that class (see Figure 8).

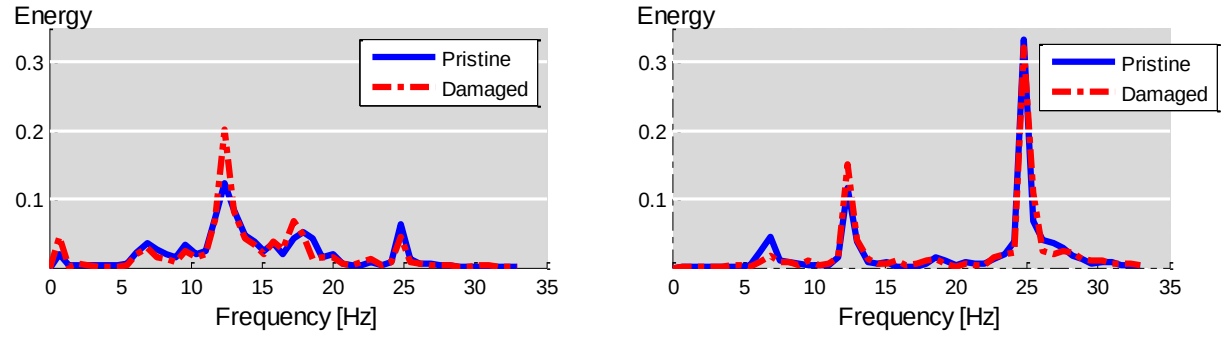

(a) 2 scenarios, sensor B ½ L. (b) 2 scenarios, sensor W.B.V.

Figure 8: Mean energy distribution (normalized to unit energy).

Let $\{x_i^{(c)}\}_{i=1}^{N_c}$ be a set of signals with N_c samples belonging to Class c. The Fourier energy map is

$$\Gamma_c(j) = \frac{\sum_{i=1}^{N_c} \| w_j^T x_i^{(c)} \|^2}{\sum_{i=1}^{N_c} \| x_i^{(c)} \|^2} \quad (1)$$

where w denotes Fourier basis vector and j denotes the frequency band. To evaluate the power of discrimination of every Fourier basis vector, we need a discriminant measure D to evaluate the power of discrimination. We will assume that higher discriminant power provides higher discrimination between classes.

For the j th Fourier basis vector, the power of discrimination Δ is denoted by

$$\Delta_j = D(\{\Gamma_c(j)\}_{c=1}^{C}). \quad (2)$$

There exist numerous choices for the discriminant measure; we use J-divergence (Kullback and Leibler 1951). Let $p = \{p_i\}_{i=1}^n$, $q = \{q_i\}_{i=1}^n$ be two nonnegative sequences with $\sum p_i = \sum q_i = 1$, J-divergence between p and q,

$$J(p,q) = \sum_{i=1}^{n} p_i \log \frac{p_i}{q_i} + \sum_{i=1}^{n} q_i \log \frac{q_i}{p_i} \quad (3)$$

Figure 9 shows a graph of the discriminant power between the frequency signals previously depicted in Figure 8.

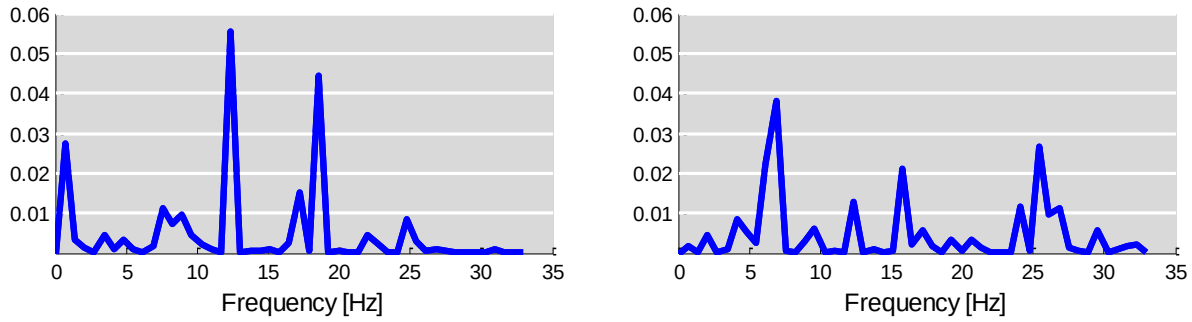

(a) Sensor B ½ L. (b) Sensor W.B.V.

Figure 9: Discriminant power (normalized to unit discriminant power).

To help understand the feature selection method, we summarize our assumptions and conclusions thus far: 1) to differentiate signals from different scenarios, we use frequencies as features. 2) if the discriminant power is higher, it is easier to discriminate between classes. The discriminant power will thus predict how well a feature will perform during classification; 3) a small number of frequencies provide most of the discriminative power; in other words, the frequency feature set is sparse. Only those frequencies that have large discriminative power are selected (this is called nonlinear approximation); see Figure 10.

Algorithm (Fourier Discriminant Basis Vectors Selection)
Task: Find $k(\leq n)$ most discriminant Fourier basis vectors
Given a dataset consisting of C classes of signals $\{\{x_i^{(c)}\}_{i=1}^{N_c}\}_{c=1}^{C}$
Step 1: Take the DFT of x.
Step 2: Construct Fourier energy map Γ_c for $c=1,\ldots,C$
Step 3: Determine the power of discrimination $\Delta_j = D(\{\Gamma_c(j)\}_{c=1}^{C})$ for every Fourier basis vector w_j
Step 4: Order Fourier basis vectors by their power of discrimination.
Step 5: Use $k(\leq n)$ most discriminant Fourier basis vectors for constructing classifier

Figure 10: Summary of feature extraction process.

This selection method performs nonlinear approximation in the Fourier domain and is data adaptive. Different data may give different frequency information and different discriminant powers. Since this method learns from the data and always chooses the frequencies with large discriminant power, it is more robust than traditional linear approximation.

Figure 11 shows the first 3D feature space. Blue circles denote the pristine scenario and red asterisks the damaged scenario. We see that with just three Fourier discriminant basis vectors, it is easy to separate the two classes.

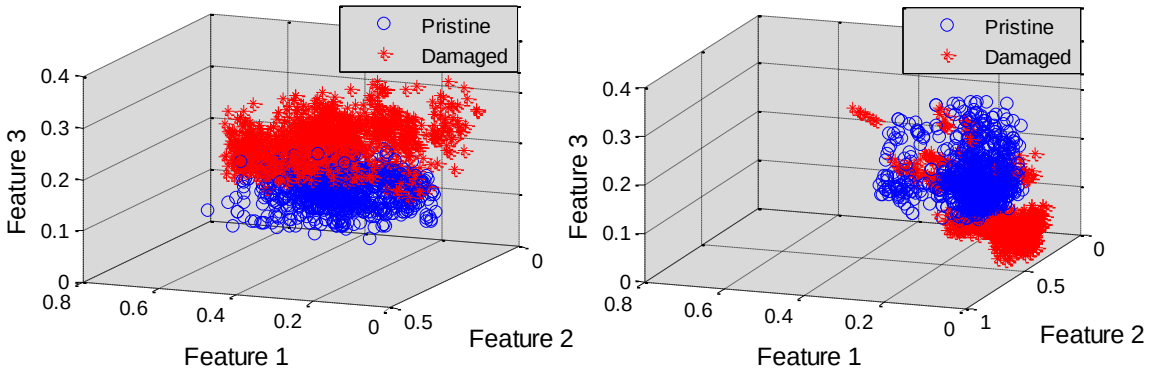

(a) 2 scenarios, sensor B ½ L. (b) 2 scenarios, sensor W.B.V.

Figure 11: Clustering of two scenarios.

Classification

The second part of a classification system is the classifier itself. In this process we take as input a feature vector and output a class label. The classification problem here is called supervised learning, as a labeled training set is given. Different classifiers are available, such as naïve Bayes, neural networks and many others (Duda et al 2000). In this work the support vector machine (SVM) was selected, which is described next.

When looking for the best boundary between classes, it is desirable to achieve two things: 1) find the boundary that gives high classification accuracy; 2) avoid overfitting. To satisfy these two requirements, SVM maximizes the margin, which means the distance between a decision boundary and a data point, and expresses it as a function of the weight vector and bias of the separating hyperplane, which is used to separate the space into two half spaces. In addition, there is low risk of overfitting because SVM produces a linear boundary.

In each scenario, data for 32 runs was collected, out of which 3 were averaged, yielding C_{32}^{3}=4960 available samples for each scenario. As our dataset, 1000 out of 4960 samples were

chosen for each scenario. 20-fold cross validation was performed. Each time, 2000 data samples, consisting of 1900 training samples and 100 testing samples, were used to create and test the SVM-based classifier. A Fourier discriminant basis search algorithm was used and the top 5 frequencies that provide the largest discriminant power were selected as features. Then a kernel SVM was used as the classifier.

Classification Results

The results of the classification experiments are presented and discussed in terms of the classification accuracy, which is defined as the number of test samples correctly classified divided by the total number of test samples.

For the two classes defined, pristine and damaged, scenarios 2-13 belonged to the latter class. The data collected from all seven accelerometers were used.

Figure 12 shows the variation of the average classification accuracy for different variables. Figure 12a shows the average across the different severities, speeds and sensor locations for each damage type. The three bridge sensors, B1/4L, B1/2L and B3/4L, are averaged and referred to as "Bridge", the two sensors at the wheel level, (W.A.V and W.B.V) are averaged and referred to as "Wheel" and the two sensors at the suspension level, (S.A.V and S.B.V) are averaged and referred to as "Suspension". The standard deviation across the averaged variables is shown at the top of each bar. An average classification accuracy for all the sensors for each damage type is depicted with a black line and corresponding percentage. The baseline in Figure 12a and b is 50%, which is the expected probability of randomly choosing between two labels (pristine or damaged). Classification accuracy values of over 90% are obtained despite the subtle changes introduced in the bridge structure. The amount of change inflicted was deliberately small to test the detection capability of the combined indirect approach using the signal processing techniques described in Section 0. The signals from the sensors located at the wheel level were classified consistently across the different damage types, and more accurately than those from the sensors located on the bridge or on the vehicle at the suspension level.

The classification results in Figure 12b show how the average classification accuracy for all damage scenarios varies for different vehicle speeds. Similarly to Figure 12a, each bar represents the mean accuracy classification across the different damage scenarios. At the top of each bar, the corresponding standard deviation is shown. Looking at Figure 12b, one can see that there is a jump between the first four speeds, between 1 and 1.75 m/s and the four higher speeds from 2 to 2.75 m/s. The average across the two groups of speeds is shown by a black line and corresponding percentage. There is about a 7% difference in classification accuracy between these two speed ranges. This classification accuracy difference is consistent for the average classification accuracy of the sensors at the different locations (i.e., Bridge, Wheel and Suspension).

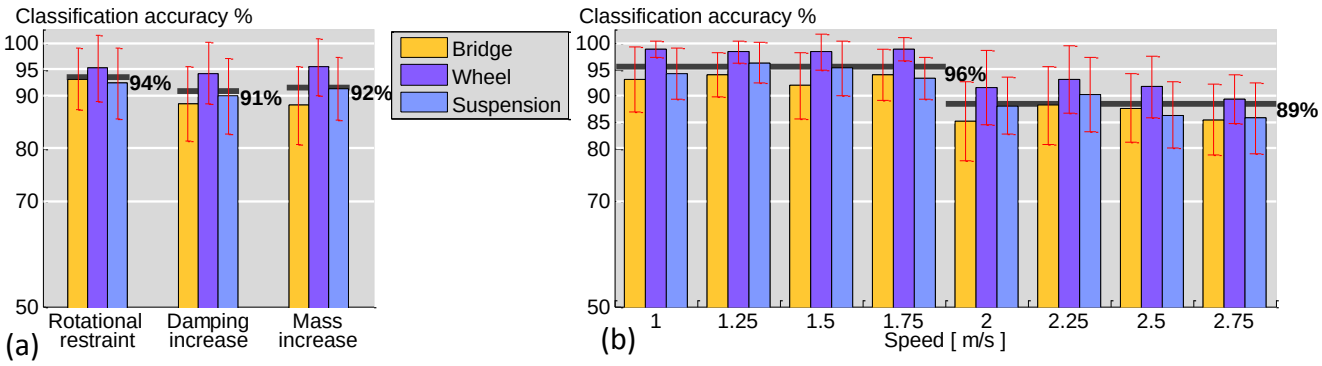

Figure 12: Classification accuracy results (a) Average by damage types. (b) Average by speeds.

Figure 13 illustrates the sensitivity of the classification method to different severity levels of damage for each scenarios. Figure 13a, b and c shows the average classification for the different damage severity levels for the rotational restraint damage type for different vehicle speeds. For all of these graphs, the thickness of the line depicts the level of damage, with the thinnest line indicating the least amount of damage inflicted (e.g., only one of four rotational restraints invoked in SC2) and the thickest line indicating the maximum amount of damage inflicted (e.g., all four rotational restraints invoked in SC5).

Figure 13a, b and c shows the average classification accuracy for each rotational restraint damage severity level for the signals from all the sensors on the bridge, all the sensors on the suspension, and all the sensors on the wheel, respectively. Figure 13d, e, and f break the results down for each sensor and show the average classification accuracy for each rotational restraint damage severity level for each signal from the three sensors on the bridge, B1/4L, B1/2L and B3/4L, from the two sensors on the suspension, S.A.V and S.B.V, and from the two sensors on the wheel, W.A.V and W.B.V, respectively.

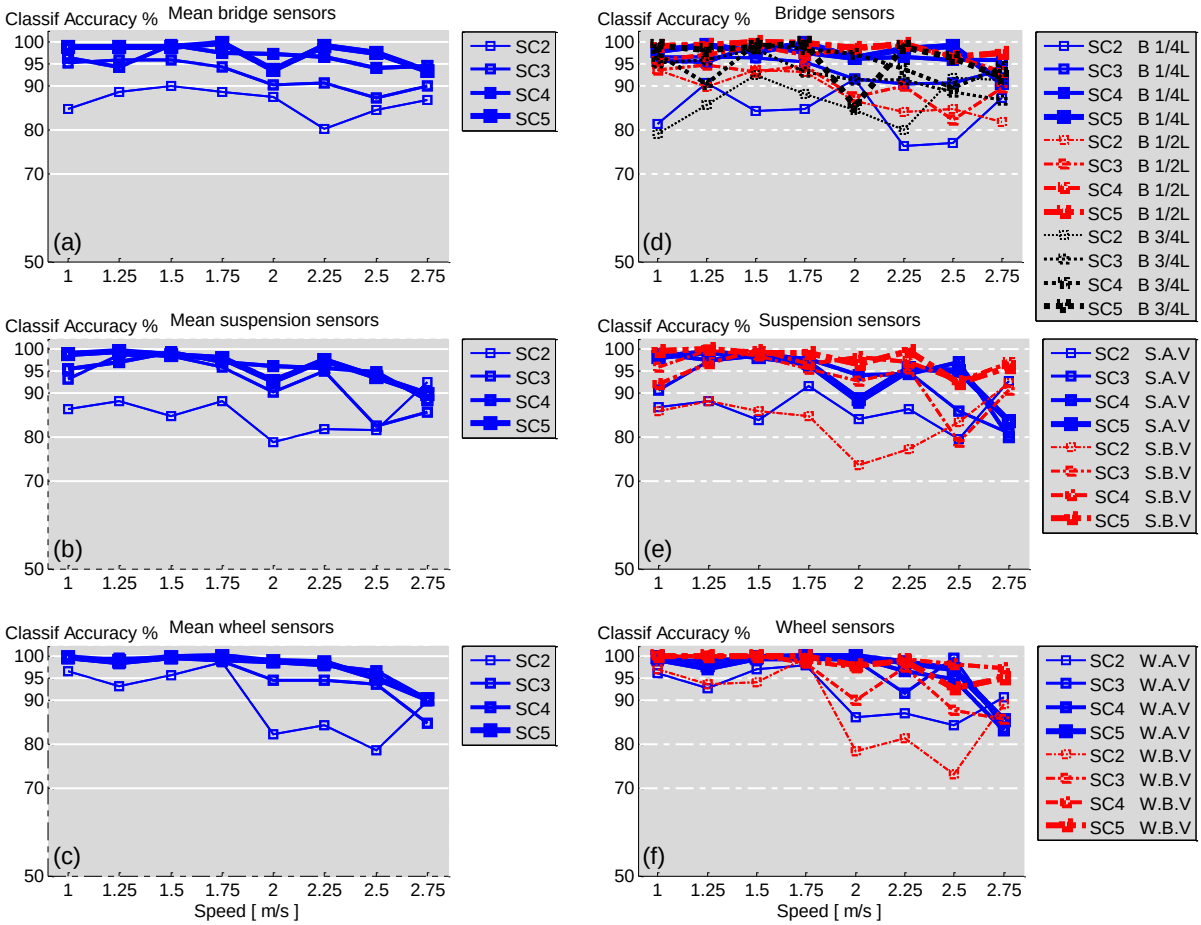

Figure 13: Classification results for rotational restraint scenarios (a,d) Bridge sensors, (b,e) Suspension sensors (c,f) Wheel sensors.

As can be seen from the graphs in Figure 13, the classification accuracy for SC2 is lower than for the other rotational restraint scenarios (SC3, SC4 and SC5), and increases as the severity of damage increases. For SC2, there is a variation in the classification accuracy with respect to the speed. However, the more severe rotational restraint scenarios seem to be more independent of the speed with high classification accuracy for low speeds and a slight parabolic decrease for higher speeds. No significant difference in the classification accuracy is apparent in Figure 13 regarding the sensor location. This shows that in terms of classification accuracy, the signal processing approach performs as well with sensor data from the vehicle (sensor or wheel) as with sensor data directly measured on the bridge. In other words, the results would indicate that in this particular set of experiments, the indirect approach has a classification accuracy that is as good as that of the direct approach.

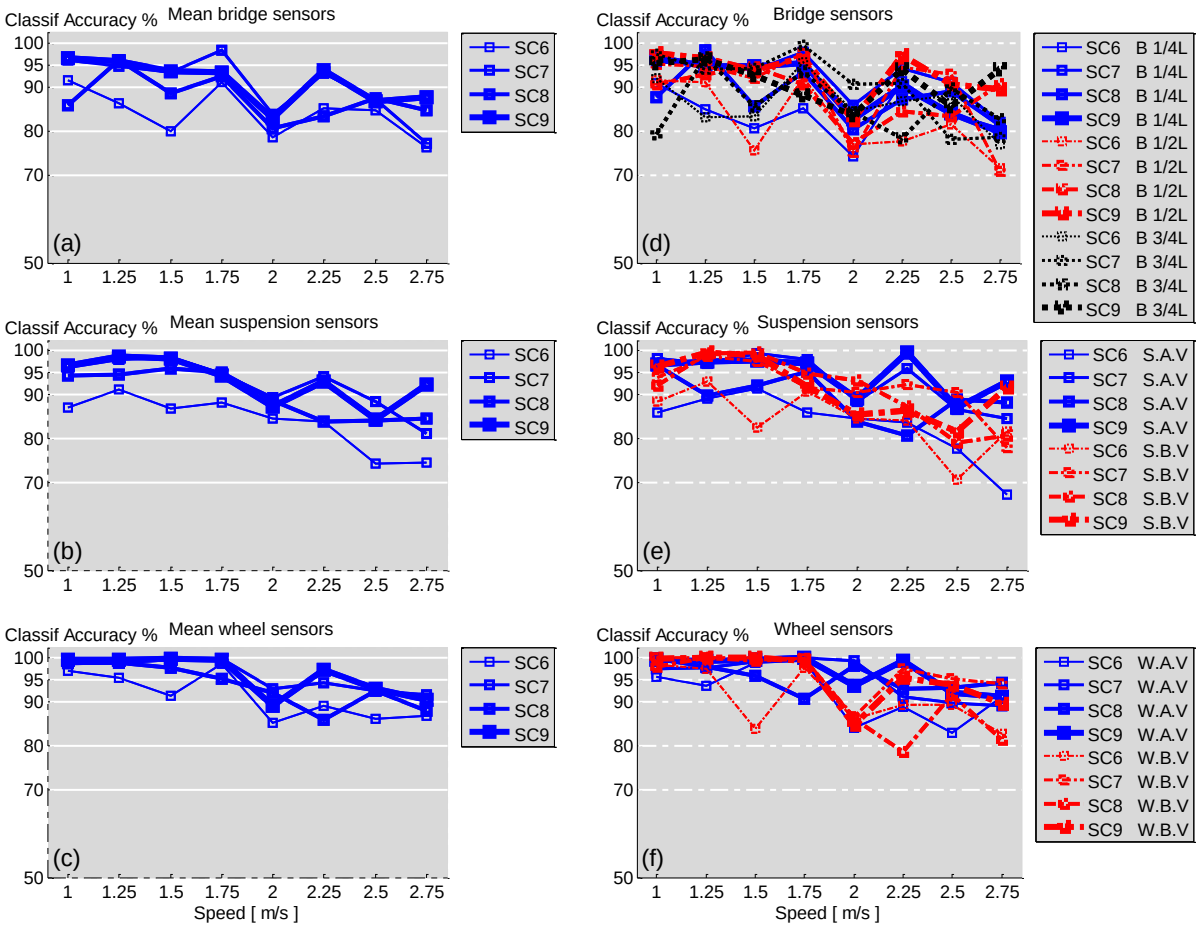

Figure 14: Classification results for damping increase scenarios. (a,d) Bridge sensors, (b,e) Suspension sensors (c,f) Wheel sensors.

Figure 14 and Figure 15 are similar in nature to Figure 13, but display the classification accuracies for the two other damage types explored: the scenarios with increasing amounts of damping and with increasing amounts of mass. Figure 14a b and c shows Scenario 6, the scenario with the single damper, as the one with the least classification accuracy across all speeds. In terms of effects due to changes of vehicle speed, a decrease in the classification accuracy appears to occur when the vehicle velocity is 2 m/s as shown in Figure 14a. In general terms, the same observations from Figure 13 apply to Figure 14. There is a slight decrease of the classification accuracy with higher speeds, and the classification accuracy seems to be independent of the sensor location; that is, there is little difference in the classification capability between the direct and the indirect approaches. Figure 15 shows the classification results for the scenarios with a mass increase at the midspan. Even though the inflicted change in the bridge structure is quite subtle, the classification accuracy is high, especially for the lower speeds. The same observations made for Figure 14 can be made for Figure 15 regarding the variation of classification accuracies with respect to the vehicle speed and the sensor locations.

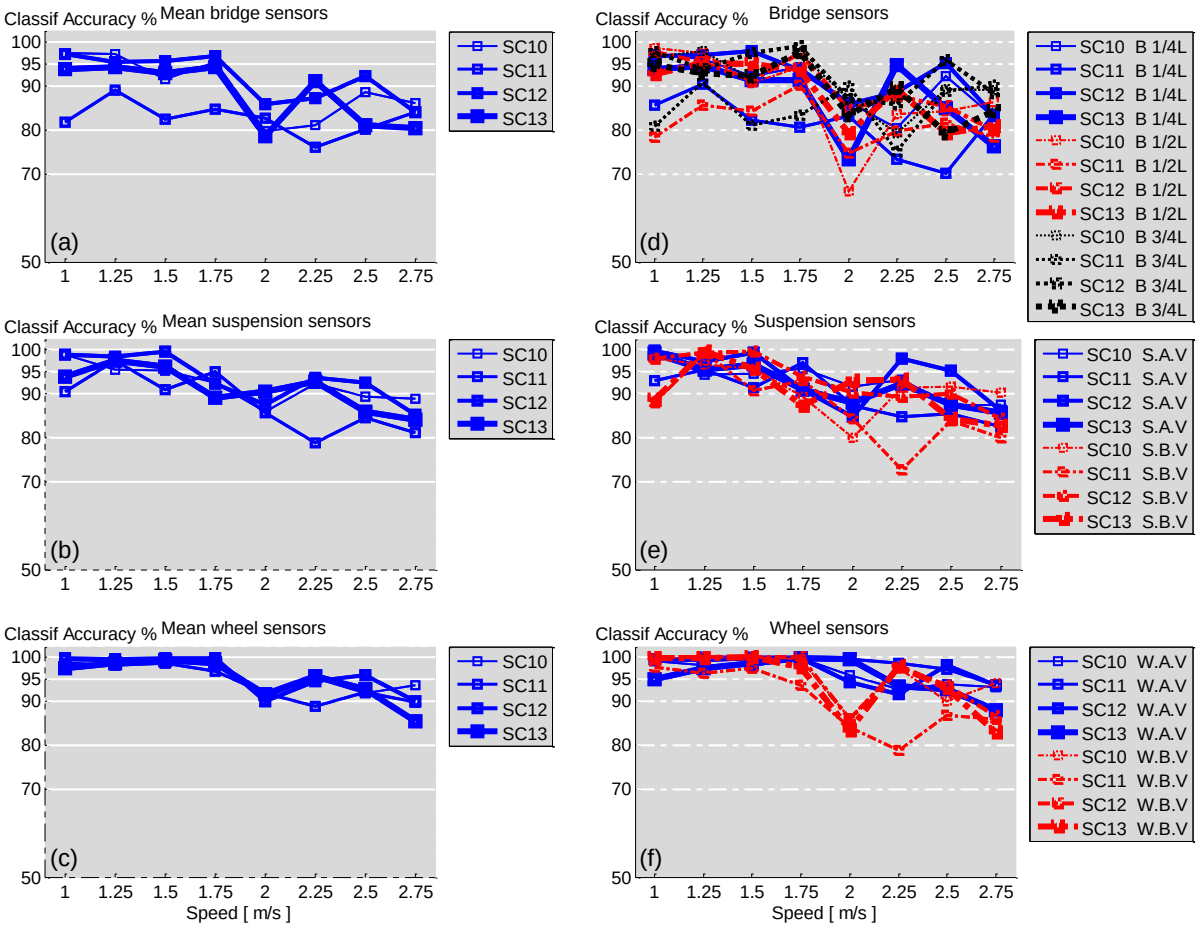

Figure 15: Classification results for concentrated mass scenarios. (a,d) Bridge sensors, (b,e) Suspension sensors (c,f) Wheel sensors.

Discussion

The damage detection capability of an indirect monitoring approach based on data collected from moving vehicles was proposed and evaluated for bridge monitoring by applying classification to experimental data taken from a laboratory vehicle-bridge physical model.

To perform numerous test repetitions, a fully automatic vehicle-bridge model needed to be built. In this study, the amount of experimental data samples is significantly greater than that from previous experiments. Each scenario was run 32 times at eight different velocities. Three different damage types were built into the experimental setting, and each damage type had four different severity scenarios.

The synchronized acceleration data from the bridge and the vehicle, and the vehicle position data, allowed for the comparison of the *direct* and *indirect* approaches in terms of the accuracy with which each could classify the existence of damage for different extents of damage.

A feature extraction technique based on averaging the power spectrum from a set of data was used to achieve very high noise reduction. Then, features extracted from the Fourier domain were automatically chosen from the denoised data samples based on their significance and classified using an SVM classifier.

High classification accuracy was achieved across three *distinct types of changes* inflicted into the bridge structure: 1) a change in the support conditions obtained by introducing rotational restraints at the supports; 2) an increase in the damping of the bridge structure; and 3) a localized mass increase at the midspan of the bridge.

The *severity* of the changes inflicted in the bridge structure was consistent with higher classification accuracy. For example, SC3, SC4 and SC5 imposed more significant changes into the bridge structure than SC2, and consistently higher classification accuracy was obtained. Nonetheless, the, classification accuracy achieved for the subtle change inflicted on SC2 is on average above 85%.

The detection of the various changes in the bridge structure was quite insensitive to the vehicle *speed.* This effect can be important for practical applications where vehicle speeds cannot be readily controlled. However, a small jump was observed between the lower and higher speeds, where the classification accuracy decreases by about 7 percent at the higher speeds.

Independent of the sensor *location,* high classification accuracy was achieved across all the sensors. The *indirect* and *direct* approaches seem to be equally effective for damage detection when applying the proposed signal processing techniques. Of the two sensor locations considered in the *indirect* approach, the wheel level and the suspension level, the sensors at the wheel level performed better than the sensors at the suspension level.

Given the simplicity of the model considered, the results presented are strictly applicable only to the particular experimental setup and cannot be generalized for full-scale structures at this time. On the other hand, we observed a high degree of consistency in the classification accuracies across the very different types and severity of damage and for different vehicle speeds. This gives us hope that our approach might be applicable to more general systems. In the next task we attempt tp validate the robustness of these results for more realistic systems and conditions. These scenarios include different roadway roughness profiles, atmospheric conditions and other bridge interaction variables such as different vehicle/bridge mass ratios, the effect of ongoing traffic and torsional effects on the bridge by non-symmetric loading from the vehicle path.

Task 2: Effect of temperature and boundary condition on the indirect health monitoring of a bridge model

Introduction

The indirect approach to the health monitoring of bridge structures uses a smaller number of sensors, mounted on moving vehicles that travel across the bridge of interest, to collect data that may help in identifying structural damage and thus serve as an indicator for more detailed analysis.

In the previous task we presented the initial results of a study that proposes to employ fleets of vehicles to routinely acquire and send data about the bridges over which they travel. A laboratory-scaled (i.e., 8 ft. (2438 mm) long) bridge model was subjected to different vehicle speeds, damage scenarios, and structural boundary conditions. In addition to the indirect measurement of the bridge motion obtained through the vehicle vibration, sensors were installed on the bridge to compare the effectiveness of the proposed indirect approach to a conventional direct SHM paradigm. The indirect and direct approaches seem to be at least equally effective for damage detection when applying the proposed signal processing techniques. The average damage detection accuracy obtained was over 90% for three different types of damage scenarios and four different severities of damage. Of the two sensor locations considered in the indirect approach, the wheel level and the suspension level, the sensors at the wheel level performed better than the sensors at the suspension level. In both cases, the indirect sensors outperformed the direct sensors. However, the previous results were constrained to one vehicle and one bridge.

This task (2) builds upon work in the previous task (1) and presents a study where the effect of three instrumented moving vehicles was evaluated on two Bridges; the effect of temperature; and simulated road roughness were also considered. With respect to the previous work in the area of indirect bridge health monitoring, this presents four new evaluations of the effectiveness of this indirect SHM approach: 1) the effect of temperature and roughness on the accuracy of the damage classification; 2) the generality of the approach when three different vehicles instead of one are used to collect data; 3) the generality of the approach when two slightly different bridge models are tested; and 4) how well classification of damage scenarios can be performed for different locations and severities.

Although the environmental and operating conditions have been addressed in the literature for direct SHM approaches (Farrar et al. 1998; Peeters and Roeck 2001; Cornwell et al. 1999; Sohn 2007; Meruane and Heylen 2011; Yan et al. 2005a-2005b, Deraemaeker et al. 2008, Serker Kamrujjaman et al 2010, Zhu and Rizzo 2011), this project presents new insights concerning the effect of certain environmental conditions on the indirect SHM approach using vehicle accelerations.

This task is organized into 4 sections, with the first being the introduction. The second section presents a description of the experimental setup and protocol used for the experiments described in this paper. The third section presents the experimental results in four subsections: a discussion of the effects of vehicle and bridge variations; a discussion of the effects of roadway roughness; a discussion of the effect of varying the location of damage and the effects of temperature gradients. The fourth section presents a brief discussion.

Experimental setup and protocol

The indirect SHM proposed in this study was tested on two bridge models. A picture of the bridge setup, hereafter designated as B1, is presented in Figure 16a. A picture of the vehicle is presented in Figure 16b. Details and technical drawings of the experimental setup were presented in Section 3.2 and shown in Figure 13. The model consisted of three segments of 8ft. (2438 mm) long each: an acceleration ramp made of 2 in. x 1 in. x 1/8 in. (50.8 mm. x 25.4 mm. x 3.175 mm.) 6063-T52 Aluminum Arch. Channel section, a bridge deck consisting of a simply supported aluminum plate with angle sections as girders as previously shown in Figure 13c, and a deceleration ramp identical to the first segment. The deck was a 8 ft.x 2 ft. (2438 x 610 mm) aluminum plate with thickness equal to 1/8 in. (3.175 mm). The second bridge herein indicated as B2, differed from B1 by having smaller (3.175 mm instead of 6.35 mm) angle sections as longitudinal girders.

Figure 16: **Experimental setup. a) General angle photo of experimental setup, b) Vehicle picture. W.A.V. is the wheel level on the A axle, and S.A.V is the suspension level on the A axle. W.B.V. and S.B.V. are the corresponding wheel and suspension levels on the B axle.**

To create three different vehicles for testing the generality of our concept for different vehicles, three different masses were placed at the midspan of the bar connecting the front and rear axle of the vehicle. As such three vehicles with different dynamic properties were able to be used in testing the indirect SHM approach. The fundamental frequencies of the three vehicles are summarized in Table 3.

Table 3: Vehicle properties.

	V1	V2	V3
Vehicle weight [kg]	4.8	5.2	5.7
A.V axle frequency[Hz]	5.0	4.6	4.5
B.V axle frequency [Hz]	5.3	4.8	4.7

In this study, the effect of damage scenarios, temperature variation, and surface roughness on the response of the sensors was evaluated. To simulate different damage scenarios, the following boundary conditions were varied: (a) rotational restraints, (b) damping and (c) mass. Rotational restraints were added to the bridge at the each of the four supports of the bridge girders as detailed in Figure 17a. Six dash-pot dampers (Figure 17b) were progressively connected to the deck to change the bridge damping characteristics. An additional mass (from 50g up to 300g) was positioned at different locations on the deck in order to simulate a change in

the natural frequencies and mode shapes of the structure. Two roughness conditions, hereafter indicated as R1 and R2, were simulated when testing the indirect SHM approach for classifying each of the damage scenarios. R1 represents the case of accidental imperfections of nominally smooth guide rail over which the vehicle travels. R2 represents the case of a built-in roughness profile for a particular bridge. R2 was achieved by modifying the rail that guides the moving vehicle with an angle section that had been randomly filed down to generate a roughness profile as shown in Figure 17c. The roughness profile measured for B1R2 and B2R2 is shown in Figure 17d, where track 1 is the right rail on the A-B motion direction and track 2 the left rail on the A-B motion direction. Finally, six electric heat sheets were attached underneath the deck in order to model the effect of changes in temperature. The heat sheets location is shown in Figure 17e. A circuit was built to discretely turn on and off the heat sheets and produce different temperature distribution. An infrared camera was used to record and report the bridge temperature scenarios.

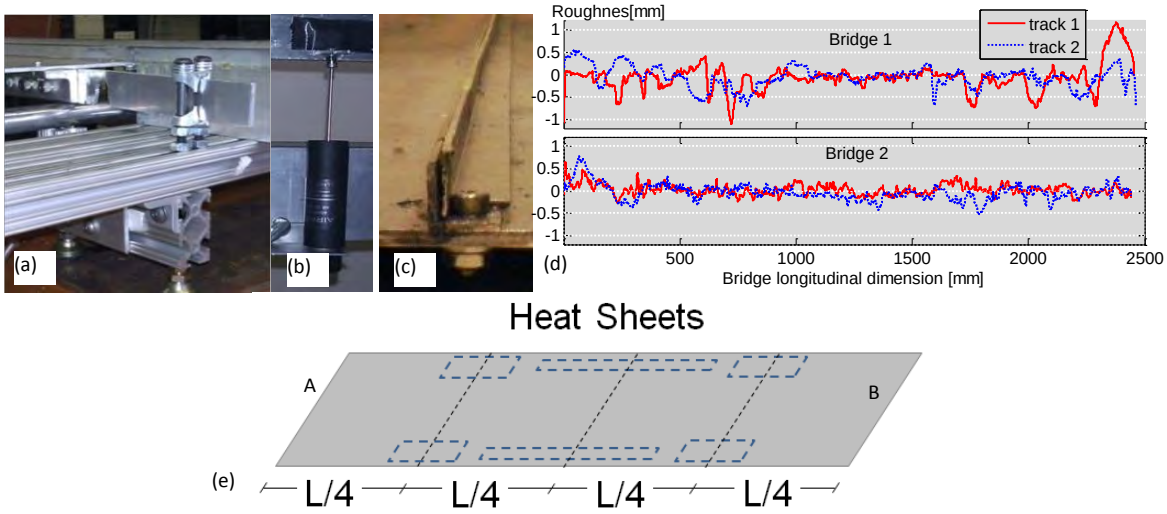

Figure 17: Introduced modifications to the bridge structure. (a) Rotational restraints (b) Additional damping and (c) Roadway roughness profile built on top of guide rail (d) Measured additional roughness (e) Scheme for Heat Sheet locations. A and B are reference points for the forward travelling direction.

Three sets of experiments were conducted. In the first set, 13 different boundary conditions were created and the three vehicles were used to monitor both bridges B1 and B2. Table summarizes the location of the rotational restraints (scenarios SC020BiR1-SC050 BiR1), dashpots (SC060 BiR1-SC090 BiR1), and masses (SC100BiR1-SC130BiR1), where i=1 for Bridge 1 and i=2 for Bridge 2. The table also presents the corresponding fundamental frequency (f [Hz]) and % critical damping ($\%\,\zeta_{crit}$) for both bridges. We used conventional free vibration testing to determine these dynamic parameters. The variation with respect to the baseline condition (SC010BiR1, the *undamaged* scenario) is presented as well in terms of %f shift and $\%\,\zeta_{crit}$ shift in Table 4. These dynamic parameters in Table 4 are shown to illustrate the magnitude of the change imposed to the bridge structure for the different damage scenarios. Then, classification analysis was performed with data collected from vehicle-bridge interaction experiments. Each of the three vehicles crossed the bridge at 8 different speeds ranging from 1 to 2.75m/s at 0.25 m/s increments. To assess repeatability, 32 iterations were performed for each combination of vehicle, damage scenario and vehicle speed. Data was recorded when the vehicle was moving in both the forward (A-B) direction and backward (B-A) direction. As such, a total of 13 (scenarios) x 2 (bridges) x 3 (vehicles) x 8 (speeds) x 32 (iterations) x 2 (motion directions) = 39936 individual experiments were conducted. However, in this paper only the forward motion

is considered from this point forward, which means that 19968 experiments composed the first set of interaction experiment data set.

Table 4: Damage scenarios: types and severity.

Schematics	B1 R1					B2 R1				
	SC	f [Hz]	% f shift	$\%\zeta_{crit}$	% ζ_{crit} shift	SC	f [Hz]	% f shift	$\%\zeta_{crit}$	% ζ_{crit} shift
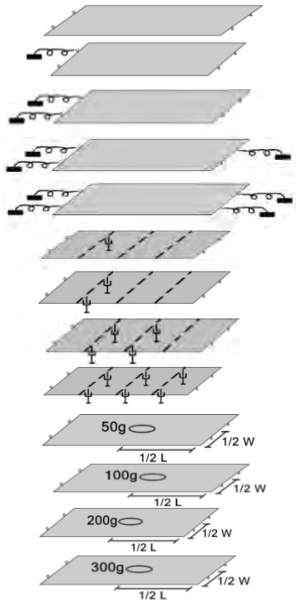	010B1R1	7.23	0.00	3.63	0.0	010B2R1	5.75	0.000	7.70	0.00
	020B1R1	7.46	3.17	6.34	74.9	020B2R1	6.07	5.56	8.59	11.6
	030B1R1	7.66	6.00	6.45	77.8	030B2R1	6.36	10.5	11.3	46.0
	040B1R1	8.11	12.2	7.97	120	040B2R1	6.72	16.8	13.9	79.9
	050B1R1	8.56	18.4	9.37	158	050B2R1	7.11	23.7	14.4	86.7
	060B1R1	7.24	0.17	8.52	135	060B2R1	5.80	0.87	14.3	86.3
	070B1R1	7.25	0.28	11.3	212	070B2R1	5.81	1.04	15.8	105
	080B1R1	7.28	0.73	26.4	629	080B2R1	5.96	3.66	30.2	293
	090B1R1	7.30	0.98	31.4	767	090B2R1	5.94	3.22	46.7	507
	100B1R1	7.19	-0.56	4.44	22.5	100B2R1	5.73	-0.43	6.97	-9.52
	110B1R1	7.18	-0.66	4.34	19.5	110B2R1	5.71	-0.72	6.60	-14.2
	120B1R1	7.14	-1.29	4.07	12.3	120B2R1	5.69	-1.02	6.42	-16.6
	130B1R1	7.09	-1.85	4.37	20.5	130B2R1	5.66	-1.60	6.41	-16.8

In the second set of experiments, the effect of roughness and location of damage was studied. The roughness condition R2 was compared to the smooth rail condition R1 for some scenarios and after only the roughness condition is used as it is a more realistic condition.

Table 5 summarizes the configurations and dynamic characteristics of both bridges under roughness scenarios R2. The scenarios in

Table 5 allowed comparing the effect of roughness on scenarios SC020BiR2, SC030BiR2, SC070BiR2, SC080BiR2, SC110BiR2 and SC120BiR2, with their corresponding SC020BiR1, SC030BiR1, SC070BiR1, SC080BiR1, SC110BiR1 and SC120BiR1, where i=1 for Bridge 1 and i=2 for Bridge 2.

Table 5: Damage scenarios: locations and roadway roughness

Schematics	B1 R2					B2 R2				
	SC	f [Hz]	% f shift	$\%\zeta_{crit}$	% ζ_{crit} shift	SC	f [Hz]	% f shift	$\%\zeta_{crit}$	% ζ_{crit} shift
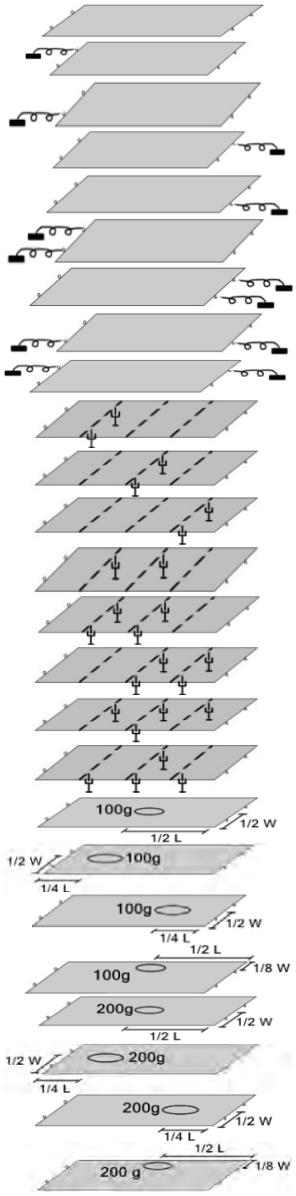	010B1R2	7.77	0.00	13.04	0.00	010B2R2	5.80	0.00	5.81	0.00
	020B1R2	7.86	1.06	15.63	0.20	020B2R2	5.95	2.62	9.85	0.69
	021B1R2	7.87	1.30	19.64	0.51	021B2R2	5.95	6.10	8.53	0.47
	022B1R2	8.00	2.99	13.46	-0.03	022B2R2	6.15	6.90	13.55	1.33
	023B1R2	8.17	5.20	24.85	0.91	023B2R2	6.20	12.07	9.23	0.59
	030B1R2	8.10	4.27	16.17	-0.24	030B2R2	6.50	13.79	8.46	0.46
	031B1R2	8.30	4.27	21.05	0.61	031B2R2	6.60	8.62	11.61	1.00
	032B1R2	8.40	8.14	19.53	-0.5	032B2R2	6.30	6.90	9.32	0.60
	033B1R2	8.10	6.85	9.98	-0.23	033B2R2	6.20	-0.86	13.14	1.26
	070B1R2	7.77	0.00	18.93	0.45	070B2R2	5.75	-0.86	19.00	2.27
	071B1R2	7.77	0.00	27.06	1.08	071B2R2	5.75	-0.86	33.13	4.70
	072B1R2	7.77	0.00	18.54	-0.42	072B2R2	5.75	-0.86	18.82	2.24
	073B1R2	7.77	0.00	23.40	0.79	073B2R2	5.75	-0.86	23.42	3.03
	080B1R2	7.67	-1.30	34.52	1.65	080B2R2	5.75	-0.86	40.47	5.97
	081B1R2	7.67	-1.30	33.80	1.59	081B2R2	5.75	-0.86	37.91	5.52
	082B1R2	7.67	-1.30	35.51	1.72	082B2R2	5.75	-0.86	39.85	5.86
	083B1R2	7.70	-0.88	37.03	1.84	083B2R2	5.75	-0.86	40.36	5.95
	110B1R2	7.67	-1.30	11.82	-0.09	110B2R2	5.65	-2.59	10.16	0.75
	111B1R2	7.67	-1.30	12.28	-0.06	111B2R2	5.65	-2.59	10.87	0.87
	112B1R2	7.67	-1.30	11.70	-0.10	112B2R2	5.65	-2.59	9.64	0.66
	113B1R2	7.67	-1.30	13.81	0.06	113B2R2	5.65	-2.59	8.99	0.55
	120B1R2	7.60	-2.16	11.25	-0.14	120B2R2	5.60	-3.45	9.04	0.56
	121B1R2	7.70	-0.88	12.55	-0.04	121B2R2	5.65	-2.59	9.41	0.62
	122B1R2	7.77	0.00	12.62	-0.03	122B2R2	5.65	-2.59	8.86	0.52
	123B1R2	7.60	-2.16	11.28	-0.14	123B2R2	5.60	-3.45	8.17	0.41

To conduct experiments related to the effect of damage location, four cases for each damage scenario were conducted. For example, in the damage scenario where a single rotational restraint is applied (SC020BiR2), it is applied at four different locations; for these four locations, the labels SC020BiR2, SC021BiR2, SC022BiR2 and SC023BiR2 were used. The same was done for scenarios SC030 BiR2, SC070 BiR2, SC080 BiR2, SC110 BiR2 and SC120 BiR2 with corresponding scenarios where the location of damage was changed. Vehicle 1 (V1 from Table 3), operated at four speeds ranging from 1.5 to 2.25m/s in 0.25 m/s increments in the second experiment set. As such, 25 (scenarios) x 2 (bridges) x 1(vehicle) x 4 (speeds) x 32 (iterations) = 6400 experiments composed the second set of experiments.

In the third set of experiments, the effect of a temperature gradient on the bridge was investigated. The bridge B1 with roughness R2 was subjected to eleven different temperature patterns, as summarized in the leftmost column of Table 6. For illustrative purposes, Figure 18 shows one of these patterns. In order to make the infrared measurements accurate, the surface of the deck was painted using a standard gray paint. For each pattern four different damage scenarios were devised. One is the pristine condition and the other three represent one of each of the three damage types. Vehicle V1 was used at speeds ranging from 1.5 to 2.25m/s at 0.25 m/s increments. A total of 11 (temperature gradients) x 4 (bridge scenarios) x 4 (speeds) x 1 (vehicle) x 1 (bridge) x 32 (iterations) = 5632 experiments composed the third set of experiments.

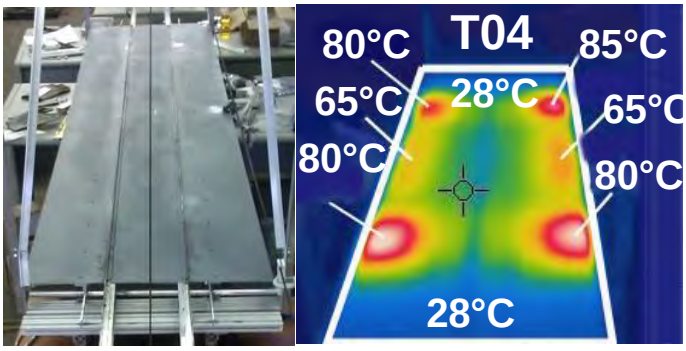

Figure 18: Bridge longitudinal view; normal and Infrared images.

The different scenarios described in this section were processed for feature extraction and classified using the same methodology as described in Task 1. The Following describes the classification experiments performed results obtained.

Table 6: Damage scenarios: temperature B1 R2

Infrared Images	Schematics	SC	f [Hz]	% f shift wrt T00	% f shift wrt SC010	ζ	$\%\zeta_{crit}$ wrt T00	$\%\zeta_{crit}$ wrt SC010
	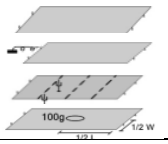	010B1R2	7.26	0.00	0.00	9.30	0.00	0.00
		020B1R2	7.40	0.00	1.92	6.32	0.00	-32.06
		070B1R2	7.16	0.00	-1.39	13.45	0.00	44.67
		110B1R2	7.16	0.00	-1.39	5.78	0.00	-37.85
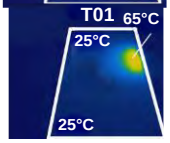	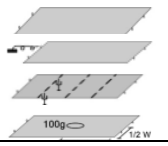	010B1R2	7.63	4.98	0.00	13.06	40.48	0.00
		020B1R2	7.71	4.08	1.05	12.15	92.35	-6.98
		070B1R2	7.68	7.28	0.77	14.43	7.32	10.52
		110B1R2	7.61	6.22	-0.22	9.03	56.27	-30.86
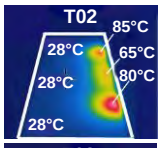	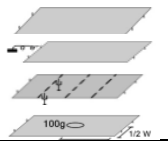	010B1R2	7.46	2.66	0.00	14.49	55.90	0.00
		020B1R2	7.55	2.04	1.30	12.33	95.28	-14.90
		070B1R2	7.47	4.23	0.11	17.75	31.98	22.47
		110B1R2	7.44	3.87	- 0.23	12.87	122.72	-11.21
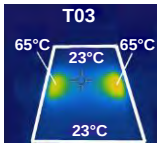	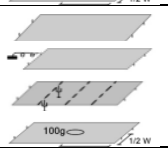	010B1R2	7.59	4.51	0.00	9.54	2.66	0.00
		020B1R2	7.91	6.79	4.14	14.79	134.25	55.02
		070B1R2	7.67	7.04	1.00	17.17	27.67	79.92
		110B1R2	7.59	5.99	0.00	11.77	103.65	23.29
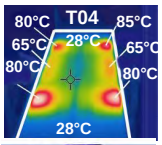	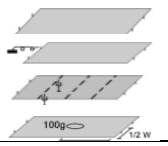	010B1R2	7.32	0.81	0.00	15.58	67.60	0.00
		020B1R2	7.50	1.35	2.46	19.0	200.86	21.95
		070B1R2	7.47	4.23	1.95	16.63	23.63	6.72
		110B1R2	7.38	3.05	0.80	12.39	114.49	-20.46
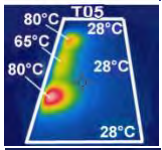	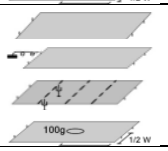	010B1R2	7.52	3.47	0.00	15.58	68.71	0.00
		020B1R2	7.60	2.72	1.17	11.45	81.25	-27.01
		070B1R2	7.57	5.63	0.67	16.26	20.91	3.68
		110B1R2	7.53	5.16	0.22	10.01	73.33	-36.15
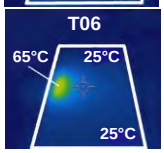	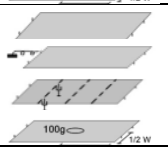	010B1R2	7.67	5.56	0.00	13.35	43.60	0.00
		020B1R2	7.66	3.40	-0.16	13.46	113.15	0.84
		070B1R2	7.47	4.23	-2.63	22.23	65.33	66.57
		110B1R2	7.42	3.52	-3.29	17.0	194.25	27.35
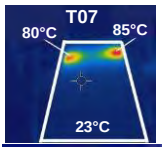	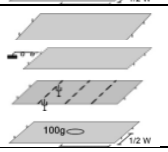	010B1R2	7.71	6.13	0.00	12.84	38.10	0.00
		020B1R2	7.86	6.12	1.90	10.2	61.55	-20.53
		070B1R2	7.67	7.04	-0.55	16.19	20.41	26.14
		110B1R2	7.67	7.04	-0.55	11.11	92.28	-13.47
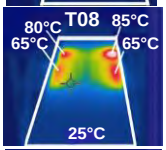	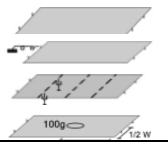	010B1R2	7.39	1.74	0.00	10.98	18.10	0.00
		020B1R2	7.66	3.40	3.59	11.35	79.78	3.41
		070B1R2	7.47	4.23	1.02	16.57	23.24	50.96
		110B1R2	7.46	4.11	0.91	11.34	96.35	3.33
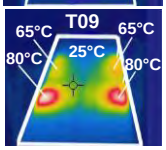	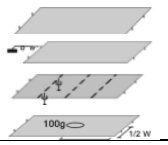	010B1R2	7.53	3.70	0.00	10.65	14.54	0.00
		020B1R2	7.66	3.40	1.62	10.51	66.50	-1.25
		070B1R2	7.48	4.46	-0.67	15.40	14.48	44.59
		110B1R2	7.48	4.46	-0.67	12.87	122.75	20.86
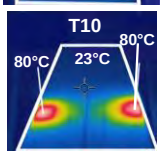	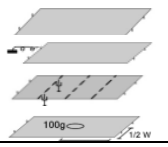	010B1R2	7.73	6.37	0.00	9.80	5.44	0.00
		020B1R2	7.91	6.79	2.33	10.0	58.31	2.00
		070B1R2	7.73	7.98	0.11	17.16	27.57	75.04
		110B1R2	7.67	7.04	-0.76	11.39	97.17	16.22

Classification Experiments and Results

Effect on Classification Accuracy of Variations of Vehicles, Bridges, Vehicle Speeds and Severity of Change

The signal processing described in Task 1 was used for damage detection and severity identification. The damage detection is a two class experiment (i.e., damaged or undamaged) while the severity classification in this task (Task 2) is a four class experiment (i.e., 020BiR1-050BiR1).

Damage detection. The first class in this two label experiment, *undamaged*, is associated with pristine condition represented by any scenario labeled as SC010BiR1 in Table 4. Scenarios labeled SC020BiR1-SC0130BiR1 were associated with the second class, *damaged*. We thus obtained, 7 (sensors) x 12 (scenarios) x 8 (speeds) x 3 (Vehicles) x 2 (Bridges) = 4032 separate assessments of classification accuracy. The classification accuracy is the proportion of true estimated results (both true positives and true negatives) in the population and is a very good indicator of the effectiveness of correctly classifying new incoming data. The results are presented for each bridge separately in Figure 19 - Figure 22.

Figure 19 show the results associated with Bridge 1. Three separate bars are used to associate the sensors to their respective location (Bridge, Wheel or Suspension). The separation of the results by location of bridge sensor allows for a comparison between the *direct* and *indirect* approach. The vertical bars represent the standard deviation associated with of the averaged results of 288 experiments for the bridge sensors 3 (sensors) x 12 (scenarios) x 8 (speeds) = 288 (detection experiments), and 192 detection experiments for the wheel and suspension sensors 2 (sensors) x 12 (scenarios) x 8 (speeds) = 192 detection experiments). At the top of each bar the standard deviation is shown. The average accuracy obtained for all the sensors is shown with a black line behind the bar graphs with numeric value of average included. Figure 19a shows the classification accuracy relative to the rotational restraints as determined using each of the three vehicles. In the framework of the binary classification performed here, an accuracy of 50% represents the probability of randomly assigning a label to one of the two classes. Overall the performance of the sensors mounted on the wheel is comparable to the direct monitoring, and superior to the accelerometers mounted on the suspensions. There is not a significant difference between Vehicle 1 and Vehicle 3 in Figure 19a. Similarly Figure 19b and 6c show the results associated with the damping and the mass increases, respectively. In Figure 19b and 19c there is a 5% to 6% decrease in accuracy between Vehicle 1 and Vehicles 2 and 3. The rotational restraints scenarios are less sensitive to vehicle changes than the damping increase or the mass increase scenarios. The consistent results of the rotational restraint can be attributed to % change in the natural frequency of the rotational restraint scenarios being more significant than the damping or mass increase scenarios (see Table 5). The fundamental frequency of the rotational restraint scenarios vary up to 18% for Bridge 1 and 23.7% for Bridge 2 while the mass and damping increase scenarios vary less than 4%. Figure 19a - c suggest that: 1) the increase of the vehicle mass has a detrimental impact on the classification accuracy that is more significant for slight changes in the bridge; and 2) the sensors at the wheel level outperform all other sensors.

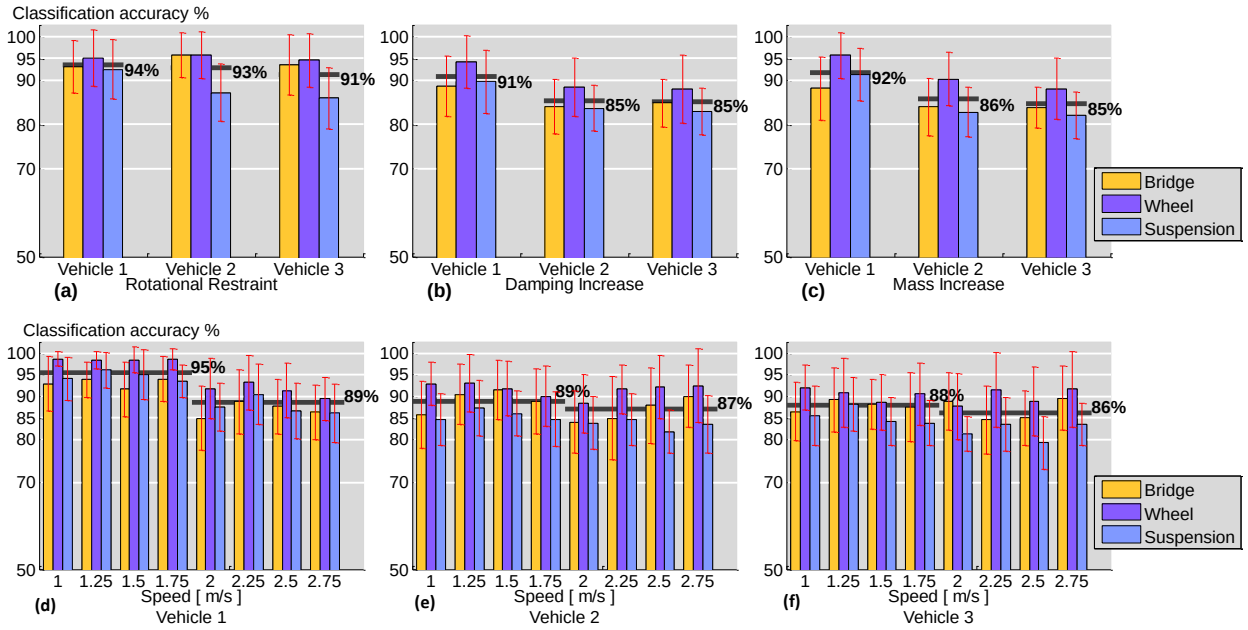

Figure 19: Bridge 1: (a-c) Damage detection accuracy for different vehicles averaged across different speeds and severity scenarios of the same damage type: (a) Rotational Restraint; (b) Damping Increase and (c) Mass Increase. (d-f) Damage detection accuracy for different speeds averaged across the different damage scenarios of the same vehicle: (d) Vehicle 1; (e) Vehicle 2 and (f) Vehicle 3.

The effect of the vehicle speed on the classification accuracy is shown in Figure 19d-f. An average of the first four speeds and the last four for all the sensors is depicted with a black line behind the bar graphs with an associated numeric value. Figure 19a shows a decrease among the first four and the last four speeds. The other two vehicles, Figure 19b and c are more stable with respect to speed variation.

Similar to Figure 19, Figure 20 reports on the results associated with Bridge 2. Figure 20a-c shows a small change in the average classification results for the different vehicles and better classification accuracy with respect to the results for Bridge 1. Figure 20a shows an average 96% classification accuracy for the Rotational Restraints scenarios. Figure 20b shows an average 93% classification accuracy for the Damping increase scenarios. Figure 20c shows an average 90% classification accuracy for Mass Increase scenarios. Figure 20d-f shows the classification accuracy across the different speeds for Bridge 2. Figure 20d-f shows a steadier trend across the different speeds for all the vehicles for Bridge 2 when compared to Bridge 1 in Figure 19d-f. The overall average classification accuracy for the three vehicles and different speeds is about 93% for Bridge 2.

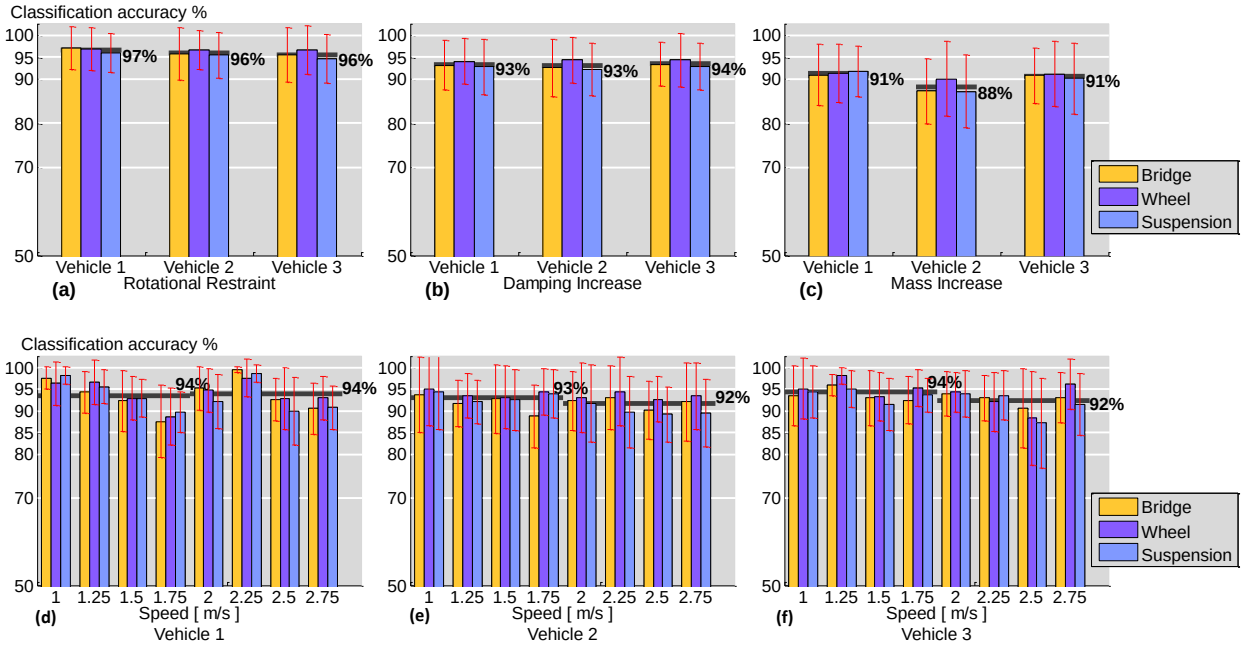

Figure 20: Bridge 2: (a-c) Damage detection accuracy for different vehicles averaged across different speeds and severity scenario of the same type: (a) Rotational Restraint; (b) Damping Increase; and (c) Mass Increase. (d-f) Damage detection accuracy for different speeds averaged across the different damage scenarios: (a) Vehicle 1; (b) Vehicle 2; and (c) Vehicle 3.

Severity classification. The severity classification experiments were defined as four label classification experiments for each damage type: Rotational Restraints, Damping Increase and Mass Increase. SC020BiR1, SC030BiR1, SC040BiR1 and SC050BiR1 are the four scenarios for the Rotational Restraint damage type as shown in Table 2. The four Damping Increase severity levels were defined by the scenarios: SC060BiR1, SC070BiR1, SC080BiR1 and SC090BiR1. Finally, the four mass Increase severity levels were defined by the four scenarios: SC100BiR1, SC110BiR1, SC120BiR1 and SC130BiR1. Each of the severity classification experiments considered was run with the corresponding data of the 4 labels, from a particular sensor and for a particular damage type (one of three), vehicle (one of three), vehicle speed (one of eight) and bridge (one of two). We ran, therefore, 1008 total classification experiments 7 (sensors) x 3 (damage types) x 8 (speeds) x 3 (vehicles) x 2 (bridges) = 1008). Just as in the case of the damage detection experiments, the results are shown for each bridge separately. Results for Bridge 1 are shown in Figure 21 and for Bridge 2 in Figure 22. The three sensor locations (Bridge, Wheel and Suspension) were represented by separate bars in the bar graphs. In Figure 21, the baseline for this case was set at 25%, which is the probability of randomly assigning a label among the four labels, P=1/4.

Figure 21a-c show the average classification accuracy results for Bridge 1 for each of the three damage types for each of the three different vehicles used (similar to Figure 19a-c and Figure 20a-c). Consistent trends as the ones observed in the two-label detection experiments can be seen in the four-label severity classification experiments. The Rotational Restraint scenarios

are on average higher in classification accuracy than the other two damage types because of a more significant shift in the natural frequency of the bridge cause by rotational restraint. The classification results for the four-label experiments across the different speeds for Bridge 1 are shown in Figure 21d-f. A significant decrease in the average classification accuracy for Bridge 1 occurs for the last four speeds in comparison to the first four speeds for vehicle 1. The classification accuracies for the other two vehicles are less sensitive to a change in speed.

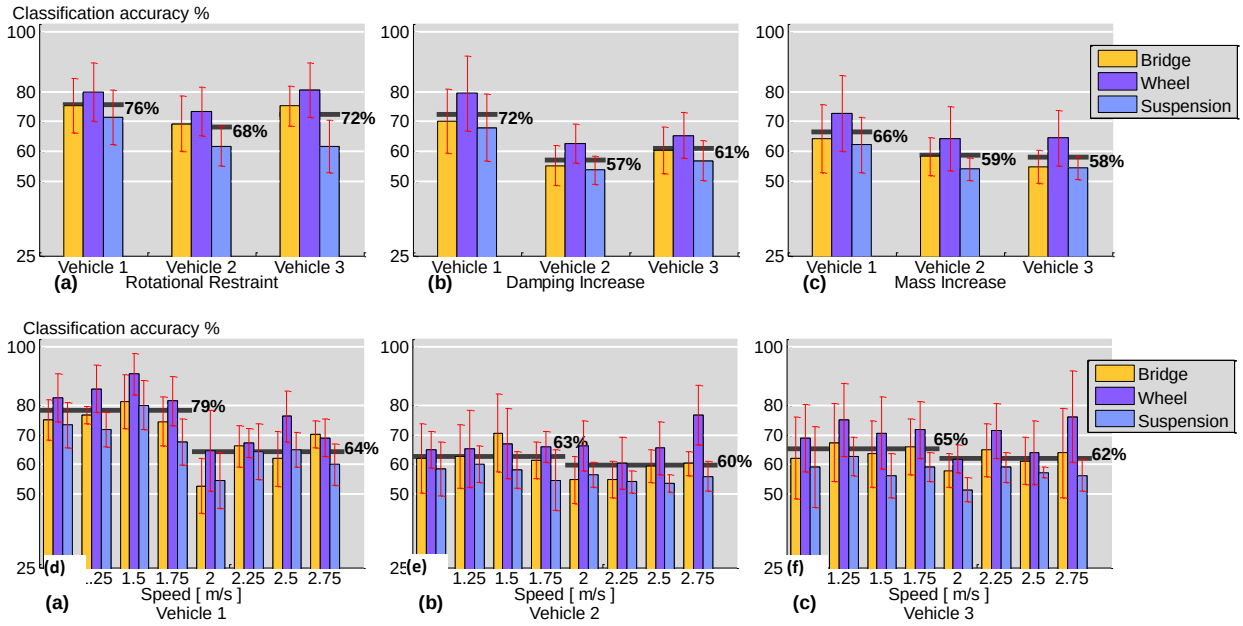

Figure 21 Bridge 1: (a-c) Damage severity classification accuracy for different vehicles averaged across different speeds and severity scenario of the same damage type: (a) Rotational Restraint; (b) Damping Increase; and (c) Mass Increase. (d-f) Damage severity classification accuracy for different speeds averaged across the different damage scenarios of the same vehicle: (d) Vehicle 1; (e) Vehicle 2; and (f) Vehicle 3.

Similarly, Figure 22 presents the classification accuracies for Bridge 2 in the same manner that Figure 21 does for Bridge 1. Figure 22a-c shows better classification accuracy with respect to the results for Bridge 1 shown in Figure 21a-c. An 80% classification accuracy was obtained for the Rotational Restraints scenarios as shown in Figure 22a, a 73% classification accuracy for the Damping increase scenarios is shown in Figure 22b, and a 67% classification accuracy for Mass Increase scenarios is shown in Figure 22c.

Figure 22d-f show the classification accuracy for damage types on Bridge 2 for different vehicle speeds. It shows a steadier trend across the different speeds for all the vehicles for the damage scenarios on Bridge 2 when compared to Bridge 1 in Figure 21d-f. The overall average classification accuracy for the three vehicles and different speeds is about 73% for Bridge 2. Overall the classification results obtained for Bridge 2 were higher than those for Bridge 1 and more stable with respect to the different speeds.

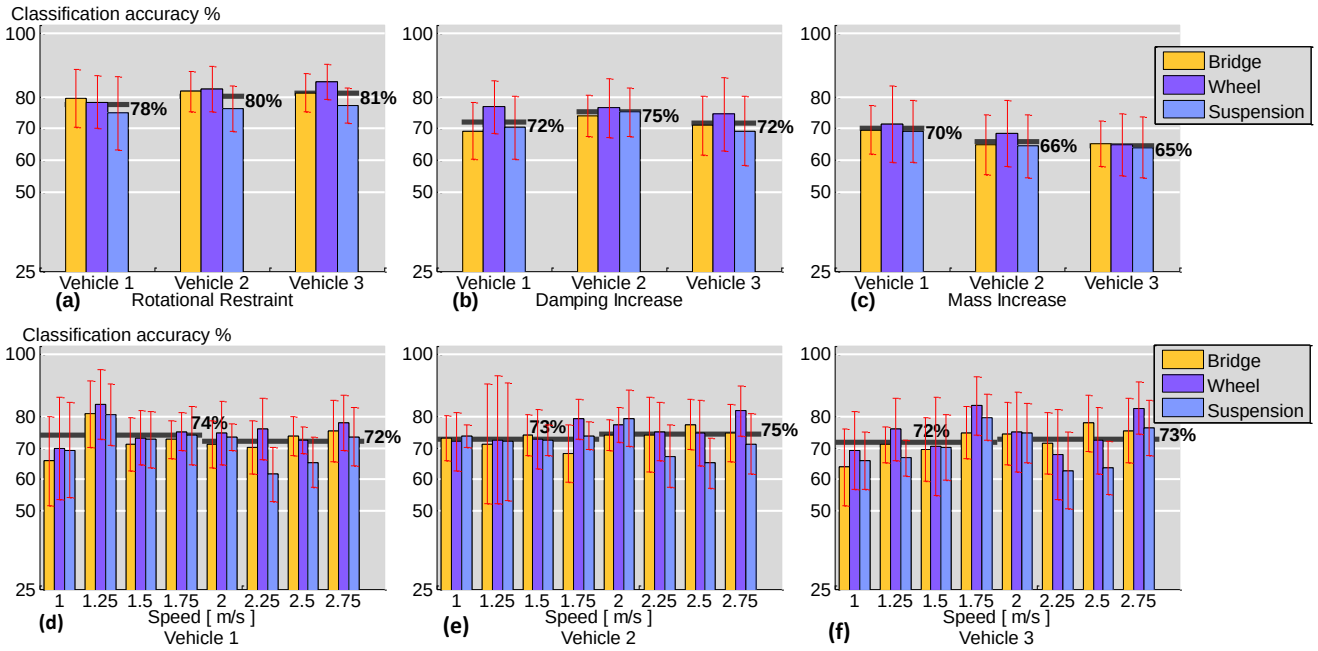

Figure 22: Bridge 2: (a-c) Damage severity classification accuracy for different vehicles averaged across different speeds and severity scenario of the same damage type: (a) Rotational Restraint; (b) Damping Increase; and (c) Mass Increase. (d-f) Damage severity classification accuracy for different speeds averaged across the different damage scenarios of the same vehicle: (d) Vehicle 1; (e) Vehicle 2; and (f) Vehicle 3.

The classification results obtained were further inspected by creating cluster plots of the four Rotational Restraint scenarios: SC020B1R1, SC030B1R1, SC040B1R1 and SC050B1R1. Figure 23 shows the cluster plot generated using the first three frequency features for the samples defined for Bridge 1, Vehicle 1, Speed 1 and without additional built-in roughness. The most significant frequencies are used as features for the different scenarios as described in Section 3.3.2. One can see that the high classification accuracy obtained for the rotational restraint scenarios can be well justified when looking at the cluster plots. The four different severity scenarios for the rotational restraints generated distinct clusters that can be distinguished by the naked eye.

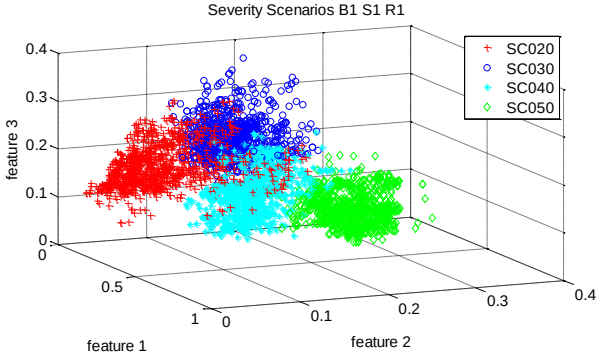

Figure 23: Rotational restraints scenarios cluster analysis. Vehicle 1, Bridge 1, Speed 1 m/s, R1 (Roughness from accidental imperfections).

Effect of Roughness and Damage Location

Tables 4 and 5 showed the second set of scenarios run for determining the effects of roadway roughness on classification accuracy. For this set of experiments, both bridges were modified to include a roadway profile as described in Section 4.2. We compared the damage detection accuracy for two different roughness conditions, R1 and R2.

Damage detection accuracy for different roughness. In a second set of experiments, we explored the accuracy of damage detection when a roadway roughness profile exists. The average classification accuracies obtained with and without roughness were plotted in Figure 24. Figure 24 a – d show the cases of the two rotational restraint scenarios Bridges 1 and 2. The caption and the scheme below in each plot in Figure 24 identifies the corresponding scenario of one or two rotational restraints. Similarly, schemes and captions identify the corresponding scenario for each plot in Figure 24. Figure 24 e – h show the cases of the two damping increase scenarios in Bridges 1 and 2, and Figure 24i – l show the mass increase scenarios in Bridges 1 and 2. Each graph in Figure 24 shows the average classification accuracy for Vehicle 1 and a specific Bridge (Bridge 1 in Figure 24a,b,e,f,I,j or Bridge 2 in Figure 24b,c,g,h,k,l), using four different vehicle speeds (1.5, 1.75, 2 and 2.25m/s), the different sensor location and the two different roughness profiles (that of the purchased angle referred to as R1 and that of the manufactured roughness referred to as R2, both were discussed in Section 4.2). The baseline was set to 50% and each bar represents the average results obtained for a particular sensor location as in the previous damage detection graphs.

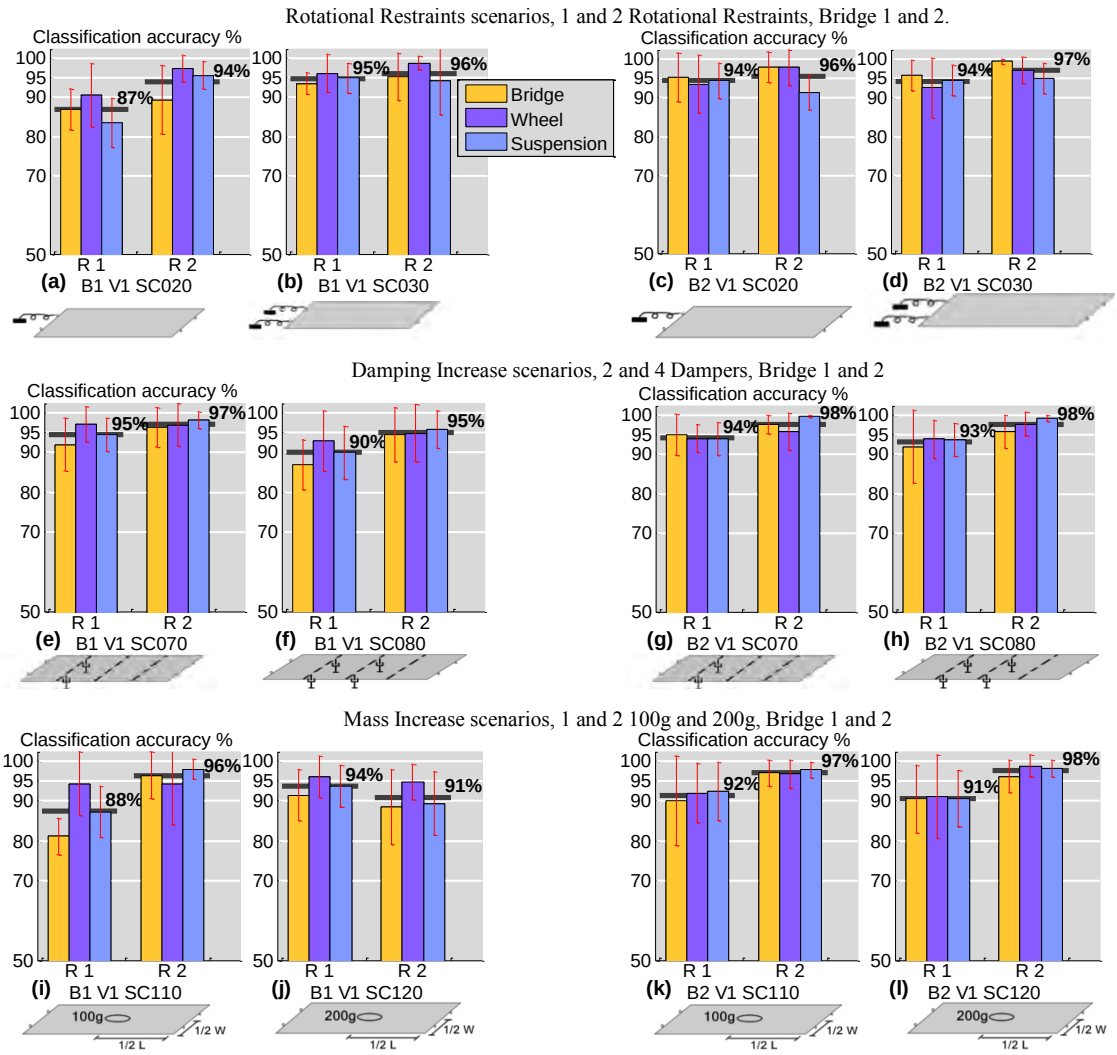

Figure 24: Damage detection comparison among bridges with two different roughness profiles. (a-d) two Rotational Restraints scenarios: (a-b) Bridge 1 and(c-d) Bridge 2; (e-h) two Additional Damping scenarios: (e-f) Bridge 1 and (g-h) Bridge 2; (i-l) two Additional Mass scenarios: (i-j) Bridge 1 and (k-l) Bridge 2.

The comparison presented in Figure 24 shows an increase in the damage detection capability when the bridge had more significant roadway roughness (case R2). These results indicate that roadway roughness may actually contribute to more accurate classification for the indirect approach, which is admittedly counter to what one might have expected when adding this variable to the problem. Intuitively, however, we can possibly explain this increase in classification accuracy as a result of the roadway roughness increasing the vibration of the bridge allowing the vehicle to better capture the dynamic characteristics of the bridge structure. However, including the roadway roughness in the manner we did in this experiment (adding an additional portion of the rail) also increased the natural frequency of the undamaged bridge (See

Table 2), and the increase in classification accuracy also could be attributed to a different vehicle/bridge frequency ratio. More investigation is necessary on this issue, but the experiment does begin to dispel concerns that adding roadway roughness would make the classification accuracy worse.

The Effect of Different Locations of Damage

Using the second set of scenarios shown in

Table 4, we setup two types of classification experiments. The first was a set of damage detection experiments for the same amount of damage at different locations, and the second was a set of four label classification experiments where the location was made to be different between scenarios with the same amounts and type of damage. These two sets of classification experiments are further described in the following two sub-sections.

Detection accuracy for damage at different locations. We explored the accuracy of damage detection for scenarios with the same amount of damage of a certain type at different locations on the bridge structure (see

Table 5). In these two-label experiments, we explore the dependence of the accuracy for detecting different types of damage on the location of that damage. As in previous sections, each of the detection experiments considered data from a particular sensor, different damage scenarios, a particular speed, a particular vehicle and a particular bridge. We obtained therefore 7 (sensors) x 24 (damage scenarios) x 4 (speeds) x 1 (Vehicle) x 2 (Bridges) = 1344 separate assessments of classification accuracy. The corresponding average results shown in Figure 25 are grouped by scenarios of the same damage type and severity level but for different locations. For example, in Figure 25a, SC020 is the scenario of a single rotational restraint at a particular location and SC021, SC022 and SC023 are also single rotational restraints but at different locations as defined in

Table 5.

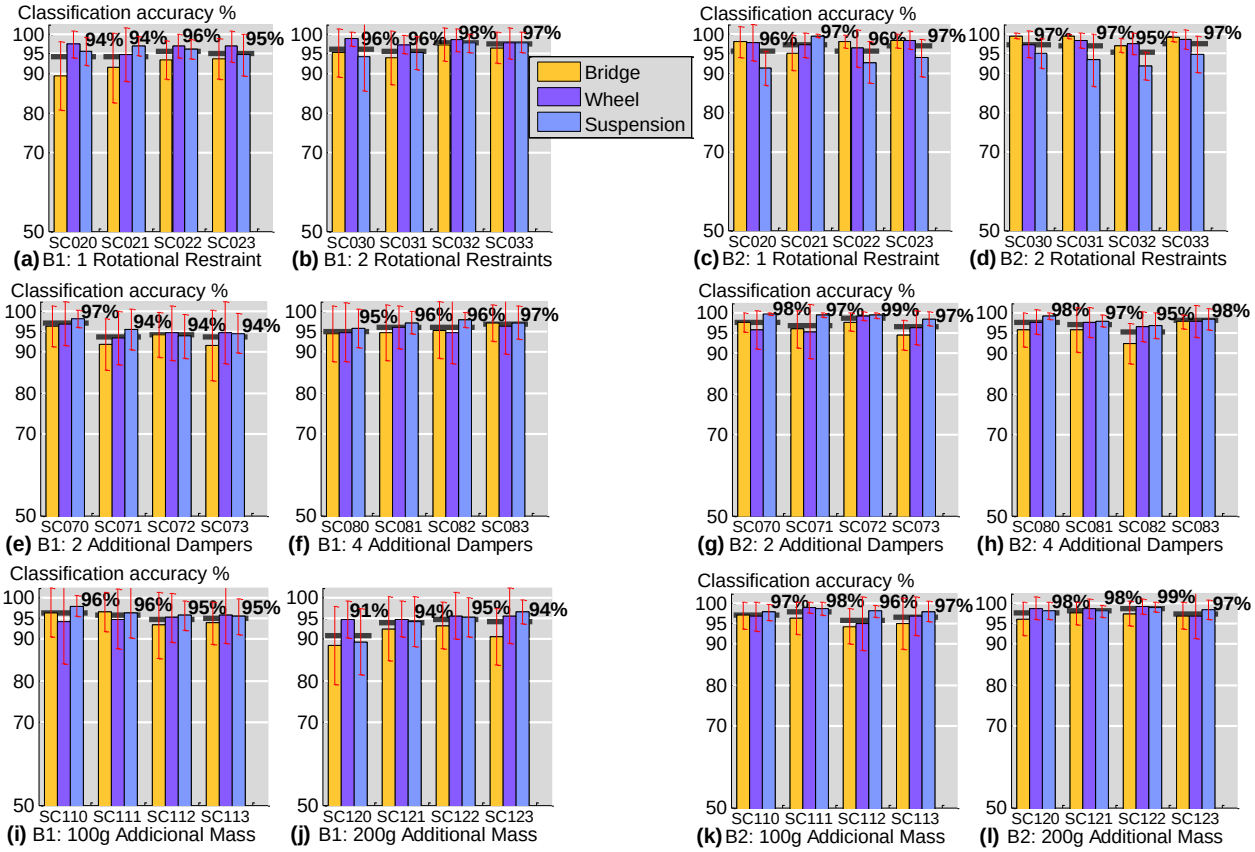

Figure 25: Damage detection comparison among scenarios with different damage locations: Bridge 1 (a,b,e,f,i,j), Bridge 2 (c,d,g,h,k,l).

Figure 25a shows the detection results for different locations of a single rotational restraint for Bridge 1, and Figure 25b shows the average detection results for two rotational restraints at different locations for Bridge 1. Among the scenarios shown in Figure 25b, SC030 and SC031 have both rotational restraints on one side of Bridge 1. The other two scenarios in Figure 25b, SC032 and SC033 had two rotational restraints, one on each side of Bridge 1. Similarly, Figure 25c shows the same scenario for Bridge 2 as that shown for Bridge 1 in Figure 25a. In both bridges, the damage detection accuracy obtained was above 95%. These results suggest that for the damage scenarios modeled, they can be detected with a very high accuracy regardless the location of the applied damage. The same is true for the other scenarios plotted in Figure 25.

Location classification. We also performed a four-label classification experiments on the data from the same set of experiments described in

Table 5. The four labels were defined as the four possible locations of damage of a particular damage type and severity applied to the bridge. In the caption of Figure 26, we named SC02- the four-label location classification experiment for classifying among the four scenarios in which a single Rotational Restraint is applied (SC020, SC021, SC022 and SC023 for Bridge 1 or 2 and Roughness 2). Similarly, we used SC03-, SC07-, SC08-, SC11- and SC12- in Figure 26 for naming the four-label location classification experiments corresponding to scenarios that vary the locations of scenarios SC030, SC070, SC080, SC110 and SC120, respectively as defined in

Table 5. Each four-label location classification experiment was run for a specific sensor, scenario type and severity, speed and bridge. We ran therefore 7 (sensors) x 6 (scenario types and severity) x 4 (speeds) x 1 (vehicle) x 2 (bridges) = 336 four-label location classification experiments. As in previous figures, we plot the different bridge results separately; Figure 26a plots the results for Bridge 1 and Figure 26b plots the results for Bridge 2. The baseline of Figure 26 was set at 25%, P=1/4 (P=direct probability).

In Figure 26a, the lowest classification accuracy obtained was 67% for SC08-, which suggests a small difference between the scenarios with four dampers. We can confirm this small variation by looking at

Table 5. SC080 through SC083 in

Table 5 share the same two middle dampers and show small variation in the fundamental frequencies. They would be therefore hard to classify by using the simple frequency-based features that we have used in this study. SC11- and SC12-, which group the location variation for a 100g and a 200g increase as defined in

Table 5, also show small or no variation of the natural frequency. However, even for this challenging case an average classification of over 70% was achieved. In Figure 26b, corresponding to Bridge 2, the lowest classification accuracies obtained were for scenarios SC08- and SC11- (four dampers and a 100g increase at different locations accordingly) there the classification accuracies for both scenarios was over 60%.

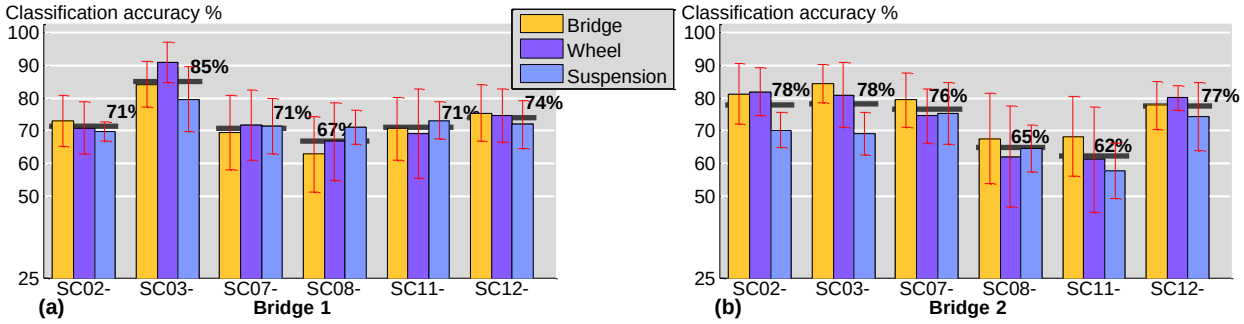

Figure 26: Location classification comparison among scenarios with different damage location: (a) Bridge 1, (b) Bridge 2.

Effect of Temperature gradient scenarios.

We now consider the third data set obtained when the thermal conditions were varied, which was described in Table 6. We ran two different classification experiments using the same signal processing and classifications approach as presented in Section 3.3. The first classification experiment was a damage detection experiment assuming the temperature scenario to be known, and the second classification experiment was a damage detection experiment assuming the temperature scenarios to be unknown. In this second classification experiment, the data from the 11 different temperature gradients for a particular bridge damage scenario is used to define a label. The next two subsections describe the obtained results.

Damage detection within same temperature scenario. As shown in Table 6, four different bridge scenarios: SC010B1R2 (Undamaged), SC020B1R2, SC070B1R2 and SC110B1R2, were run under each specific temperature gradient scenario. We preformed damage detection experiments (two labels: Undamaged vs Damaged) for each of the three damage scenarios: SC020B1R2, SC070B1R2 and SC110B1R2 at a particular temperature scenario. In other words, each damage scenario is defined by the corresponding experimental iterations that were conducted for the given temperature scenario that consist of 1 (temperature gradient scenario) x 1(bridge damage scenario) x 1(speed) x 1 (vehicle)x 1 (bridge) x 32(runs) = 32 iterations to define a particular damage scenario. Figure 27 shows the average and standard deviation of the classification accuracy obtained for the damage detection experiment for each the three damage scenarios under a specific temperature gradient conditions. That is the average of 3(damage scenarios) x 4 (speeds) = 12 classification accuracy experiments. In all these experiments Vehicle 1 was used and the speeds ranged from 1.5 to 2.25m/s at 0.25 m/s increments. The baseline of the graph was set to 50% as in all the previous two–label damage detection graphs. The obtained results show that the average classification accuracy might vary by as much as 15% for the different temperature scenarios (see the results for T02 and T03). This variation suggests that some temperature scenarios might be more favorable for damage detection than others. The results obtained from the bridge sensors and from the vehicle sensors achieved roughly the same classification accuracy.

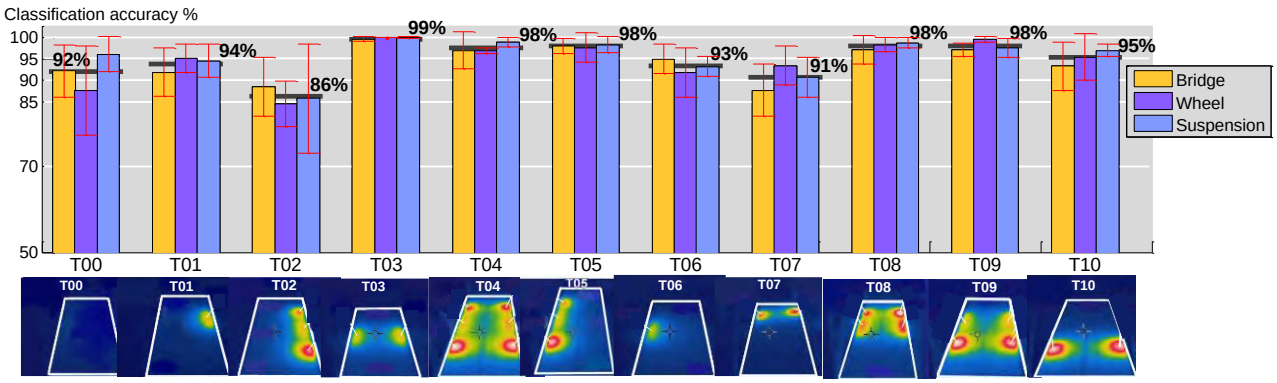

Figure 27: Classification accuracy of damage scenarios for specific temperature gradient scenario.

Damage detection without considering temperature information. The second classification experiment considered a particular bridge condition to define a particular bridge damaged label with all the temperature variations. That is, SC010 is now defined by the 32 runs at each of the 11 different temperature conditions. The other three scenarios in Table 6: SC020, SC070 and SC110 are also defined for this condition. In this case we are treating the temperature condition as an embedded variable in the data. Figure 28 shows the average classification results obtained for the detection experiments. The overall results show an interesting average classification accuracy among the three damage scenarios. The average classification results decrease with respect to the previous experiments shown in Figure 27, where the temperature was considered a known variable. This decrease was expected as this is a more challenging experiment where the embedded variable can mask the damage changes made to the bridge structure. In general terms, it can be seen that the Wheel and Suspension sensors achieved classification accuracies similar to or better than those using the sensor located on the bridge structure.

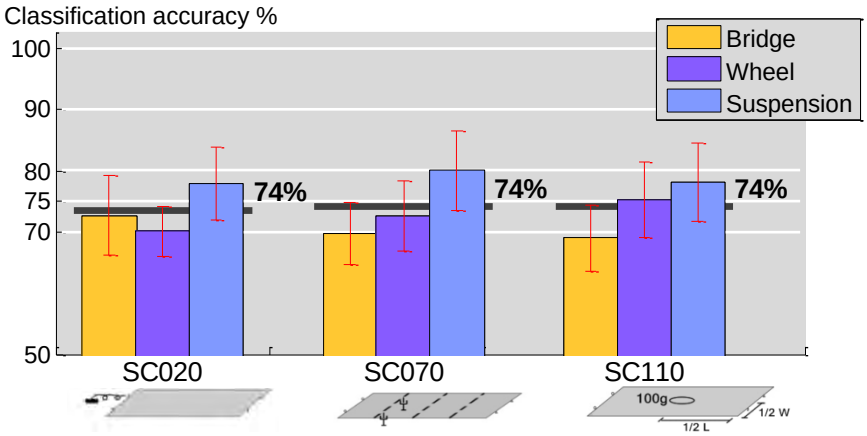

Figure 28: Classification accuracy of damage scenarios regardless temperature scenario for training and testing.

Discussion

In summary, we further validated an *indirect* bridge SHM approach for different temperatures based on vibration data collected from a vehicle as it traverses the bridge structure, and compared the results with the traditional *direct* monitoring approach that uses sensors placed

on the bridge structure. We built an automated laboratory scale experimental setup that allowed for data acquisition and accurate control of the speed of the vehicle, accurately known modifications to the bridge, and multiple repetitions. We used a simple feature extraction and classification approach to perform several classification experiments on the collected data. We explored the space of possible variations to the vehicle-bridge interaction system to test the robustness of the indirect monitoring approach versus the direct monitoring approach. The variations considered two different bridges, eight different vehicle speeds, three different masses of vehicle, two roughness scenarios, 11 different temperature gradients, and three different types of damage made to the bridge (rotational restraints, additional damping and an additional mass). The different types of damage were explored in terms of detecting the existence of the damage and also classifying among scenarios with different damage locations and different levels of damage severity.

For the vehicle variation, the introduction of a higher mass to the vehicle chassis decreased the average classification accuracy for the mass increase and damping increase scenarios in Bridge 1 by about 6%, and showed no significant difference for Bridge 2. The detection classification results for the sensors located at the wheel level and at the suspension system of the moving vehicle (*indirect monitoring*) were comparable with the results obtained with sensors located on the bridge structure (*direct* monitoring).

For the different vehicle speeds, the Bridge 1-Vehicle 1 combination showed a 6% decrease in the damage detection classification accuracy between speeds of 1.75 m/s and 2m/s; however, all other Bridge-Vehicle combinations showed low dependence on speed. The detection results obtained using the indirect and direct monitoring approaches were similar. The severity classification results show similar trends to the damage detection classification accuracy for Bridges 1 and 2, but about 20% lower than those for damage detection. The decrease is attributed to the increase of the number of labels from two labels in the detection experiments to four labels in the severity classification experiments.

The inclusion of roadway roughness increased the detection rate for almost all the damage scenarios compared. The indirect and direct monitoring approaches showed comparable results for all cases.

Different locations of the same magnitude of damage showed similar average damage detection classification accuracy for both bridge structures. The location classification accuracy, a four label classification task, showed variations from 85% to 67% for Bridge 1 and from 78% to 62% for Bridge 2. This variation can be attributed to the choice of scenarios implemented. Some scenarios were very similar to others they were being classified against. The indirect and direct monitoring approaches showed comparable results for all location classification cases.

Two types of damage detection experiments were conducted regarding various temperature gradient conditions. The first damage detection experiment compared damage scenarios within the same temperature gradient conditions. The variation of the detection accuracy of about 15% suggests that there are temperature gradient conditions more favorable than others for detecting damage scenarios. The second damage detection experiments assumed that there was no information about the temperature gradients and treated it as an embedded variable in the data. The classification results for the three scenarios explored on this case showed remarkable consistency for the three scenarios explored of about 74%. Once again the sensors considered as indirect performed just as well as those sensors located on the bridge structure.

The results showed in this section, match, in most cases, the intuition about the behavior of the vehicle-bridge interaction systems with the obtained trends in terms of the classification

accuracies. That is, scenarios that had a greater variation of the dynamic characteristics identified by free-vibration experiments had also higher detection accuracy.

The signal processing methodology used in this work was based on frequency based features and a standard classification algorithm. The numerical classification results could be greatly improved by exploring other feature space representations and other classification algorithms.

From this work, it is clear further research was needed to validate the indirect versus the direct approach for bridge SHM. The results presented thus far are constrained to the corresponding experimental setup. Task 4 in this report to will explore other systems and scales where indirect monitoring could be employed.

Task 3: Damage Localization, Quantification and Bridge Condition Uncertainty

In this section we investigate questions which will be of great interest to agencies who own bridges. Assuming there is damage, they will want to know where the damaged has occurred and the severity of the damage. We conducted experiments to answer these questions. In particular our goal was to see how changes in the severity of the damage and location of the damage would change the dynamic response of the bridge so that we could create a general framework for looking at these changes. In these experiments we defined "damage" as the placement of an additional mass on the bridge which would change the frequencies of the bridge in a similar manner to the way that frequencies of the bridge might change as the stiffness of the bridge decreased.

A full write-up of this work is presented in Lederman et al. 2014, which has been included in the Appendix.

The second part of this task was to consider cases where the condition of the bridge is not known. Our approach to structural health monitoring is model free—we do not build a finite element model of the bridge because this can be expensive and inaccurate. Instead we have the algorithms learn about the nature of the bridge. For this learning process to occur the data must be "labeled"—for example, the bridge could be labeled as healthy and then the initial signals would become the baseline. Later changes in the bridge would then be compared to this baseline.

However, in practice we will have to begin monitoring bridges at different stages in their life cycle. When the bridge inspectors give a condition assessment, signals collected during that time period can be labeled by their assessment. However this assessment might change, or worse, the inspectors could have misclassified the condition of the bridge. To accomplish this task we looked at a different approach for selecting features and for classifying the data. A full write up of this work can be found in Chen et al 2013, which can be found in the Appendix. However a brief overview of the work is given below.

Data Analysis for SHM

The goal of data analysis is to analyze patterns of various damage types and label each observation into a predefined class. In data analysis for SHM, we can detect presence or absence of damage, or go one step further and detect the severity, location and type of damage. These tasks are signal-processing problems known as classification.

Classification System

A classification system often has two core components, a feature extractor and a classifier. The feature extractor extracts numerical features from the data with the aim of discriminating classes based on those features. The task of the classifier is to label each observation into a predefined class based on those features.

Progress

So far, we have tried several different feature extractors and classifiers to (1) study the potential power of each algorithm, (2) find the suitable situations in which to use each algorithm, (3) study the difficulty of classification problem, and (4) upgrade our classification system to achieve better results.

Feature Extractor

A feature extractor often includes two approaches: representation and discrimination. Generally speaking, representation approach is robust to noise and does not need training data, while the discrimination approach is sensitive to noise and needs training data, but often achieves better classification results. In bridge SHM project, data is noisy and we have access to few labeled samples (it is expensive to physically inspect the bridge and label it as damaged/undamaged, etc). Up to this point, we have used Fourier discriminant features, principal component analysis, sparse representations, and local Fisher discriminant analysis. We are constantly studying the effectiveness of various feature extractors for use in this project as well as how to combine the advantages of both representation and discrimination approaches.

(1) Fourier discriminant features
Fourier discriminant features choose the most discriminative frequencies as features. It consists of two steps: 1. the data is represented in a Fourier basis; 2. J-divergence is used as the discriminant measurement to evaluate the discriminative power of each frequency and then sort frequencies by their discriminative power. Step 1 provides representation by a fixed basis and Step 2 adaptively discriminates data.

(2) Principal component analysis (PCA) (Abdi et al 2010.)
PCA is a standard representation algorithm. It learns an orthogonal linear transformation from the given data that transforms the data into a new coordinate system such that the first coordinates captures the greatest variance, the second coordinate captures the second greatest variance, and so on.

(3) Sparse representations (Wright 2009, Aharon 2006).
Sparse representations represent signals with a linear combination of a small number of elementary signals called atoms that account for most or all information of a signal. Often, the atoms are chosen from a so called over-complete dictionary. The aim of a sparse representation is often to reveal certain structures of a signal and to represent these structures in a compact and sparse representation. The representation dictionary is learned from data, so it is data adaptive.

(4) Local Fisher discriminant analysis (LFDA) (Sugiyama 2007)
LFDA is an advanced discrimination algorithm that improves on the popular linear discriminant analysis. It provides a linear supervised dimensionality reduction mechanism and is particularly useful when some classes consist of separate clusters. LFDA has an analytic form of the embedding matrix and the solution can be easily computed by solving a generalized eigenvalue problem. Compared to linear discriminant analysis, it considers local information and provides more than (c-1) nontrivial eigenfunctions, where c is the number of classes.

Classifier

We have used two types of classifiers. Supervised classifiers are trained on the labeled data and tested on unlabeled data. Semi-supervised classifiers are trained on both labeled and unlabeled data and tested on unlabeled and unseen data. We have used support vector machine (SVM) and graph-based semi-supervised learning.

(1) Supervised learning (Bishop 2006)
Supervised learning is the machine-learning task of inferring a function from labeled training data. Some popular supervised learning algorithms include naïve Bayes (Domingos 1997),

logistic regression (Hilbe 2009) and SVM (Burges 1997), which we used almost exclusively. When classifying, we want a classifier to learn the boundary between different classes and separate the feature space into different class spaces. Intuitively, a good separation is achieved by the hyperplane that has the largest distance to the nearest training data point of any class (so-called functional margin), since in general, the larger the margin the lower the generalization error of the classifier. SVM maximizes the margin and expresses it as a function of the weight vector and bias of the separating hyperplane.

(2) Semi-supervised learning (Zhu 2005)

Semi-supervised learning is a technique for training classifiers with both labeled and unlabeled data. It assumes that unlabeled data can provide distribution information to build a stronger classifier. Some popular semi-supervised learning algorithms include generative mixture models with expectation maximization, co-training, transductive SVM and graph-based approaches. We focus on label propagation (Zhu 2002), which is one of graph-based approaches.

Label propagation assumes that, while the measured samples exist in a high-dimensional space, they are distributed in a low-dimensional manifold. Based on this, a graph is constructed to analyze the distribution of both labeled and unlabeled samples. By analyzing how the labels propagate on this graph, classification can be achieved.

Multiresolution Framework

Multiresolution classification system (MRC) (Chebira 2007, Chebira 2008) is a generic classification framework. It decomposes images into localized space-frequency subbands using a wavelet packet tree (Coiffman 1991). In each subband, MRC extracts features, classifies them and gets a local classification result. A global weighting algorithm combines the local results to get a global decision. Its advantages are that (1) it extracts hidden features in localized time-frequency zones; and (2) the traditional generic classification system is only the first subband of MRC. It reaches a decision using information from all the subbands.

Summary

Version	Feature Extractor	Classifier
Cerda Thesis 2012	Top 5 Fourier discriminative features	SVM
ASCE Structures Conference 2013, Pittsburgh PA	Sparse representations	MSE
Damage Assessment of Structures 2013, Dublin Ireland.	LFDA	SVM
ASCE International Workshop on Computing in Civil Engineering 2013, Los Angeles CA	Multiresolution + PCA	SVM
IEEE International Conference on Acoustics Speech Sound and Signal Processing 2013, Vancouver, CA	Multiresolution + PCA	SVM

Comparison

1. Fourier discriminative features vs. sparse representations vs. LFDA
   When the labeling ratio (the number of the labeled data/the number of all the data) is 80%, we compare Fourier discriminative features with sparse representations and LFDA under 18 scenarios.

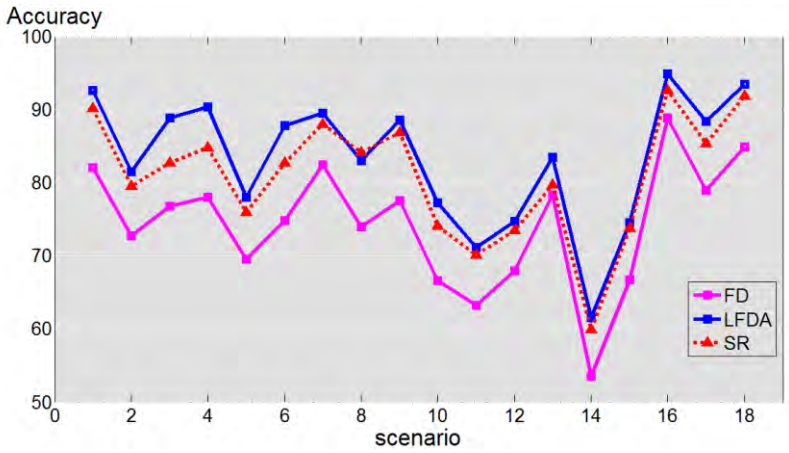

Figure 29 - Comparison of three methods, Fourier Discriminant Features, Local Discriminant Features and Sparse Representation for 18 scenarios

2. Fourier discriminative features vs. Multiresolution + PCA
   When the labeling ratio is 50%, we compare Fourier discriminative features with multiresolution classification

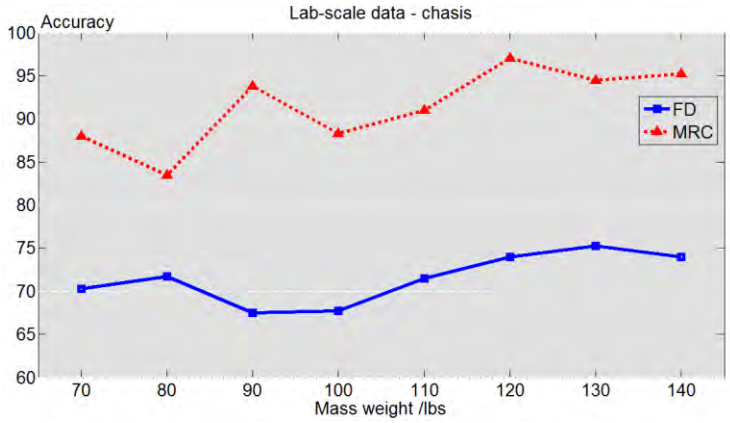

Figure 30 - Comparison of Fourier Discriminant Features and Multiresolution Classification

3. Supervised learning vs. semi-supervised learning under multiresolution framework
   We compare supervised learning with semi-supervised learning by changing the labeling ratio. SMRC stands for supervised multiresolution classification while SSMRC stands for semi-supervised MRC.

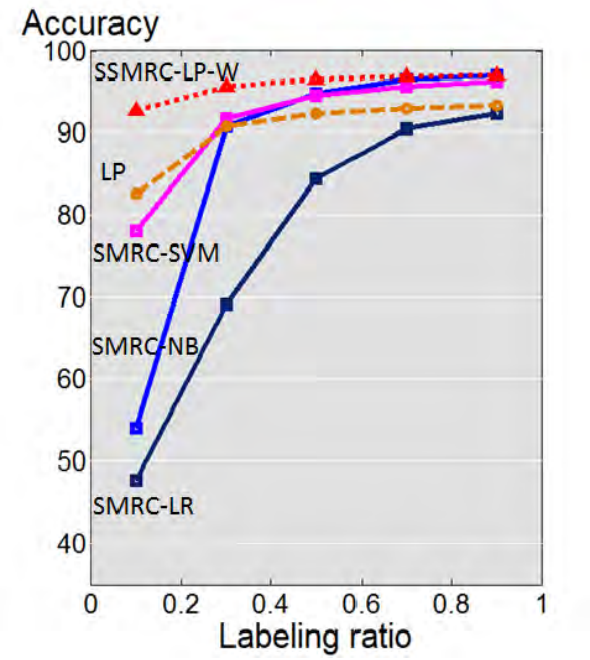

Figure 31 - Discriminant Accuracy at Various Labeling Ratios

Conclusions

Based on theoretical analysis and numerous experiments, we conclude the following:

(1) Feature extractor: When the data is sufficiently labeled, LFDA has the best performance; when there are fewer labels, sparse representations are more effective.

(2) Classifier: When the data is sufficiently labeled, SVM has the best performance; when there are fewer labels, a semi-supervised classifier is more effective as can be seen in Figure 31.

As a framework, multiresolution classification system provides a consistently better performance than a generic classifier.

Task 4: Garage Experiments and Light Rail Data Collection

In the first section we explain a set of experiments we conducted on a parking garage structure. We compare these results to those of the laboratory scale model.
In the second part of this section, we show some initial results from recent data collection we have begun in collaboration with the Port Authority of Allegheny County. This project is still in its early phases. We have not yet collected sufficient data to give results. However we will specify what hardware we have installed, describe the database we have setup to manage the data, and show some of the raw data.

Garage Field Experiments

Our goal was to find a long span structure where we could conduct uninterrupted experiments, while minimizing the impact on other users. We ultimately selected the 3rd level of the East Campus Parking Garage at Carnegie Mellon's Pittsburgh Campus. The lower levels of the garage fill first, so the 3rd floor was nearly empty during the early mornings when we conducted our experiments. Although there were relatively low levels of ambient vibration in the deck, we did occasionally observe some noise from cars on lower levels, and from heavy traffic on adjacent Forbes Avenue.

The deck is made of 32" x 8'6" (0.81m x 2.59m) precast prestressed concrete double tees with estimated 700lpf (1042 kg/m), each isolated from adjacent sections by an elastomer (Figure 32, 33). Due to the elastomer, each double tee section acts as a simply supported beam. The span where we ran our experiments is 51.5ft (15.7m) long with an approximate weight of 36,120lbs (16,383kg).

As our vehicle, we used an iRobot ATV-JR robot, with a weight of 50kg[3], traveling at a constant speed of 2m/s. We fitted the robot with three Vibra-Metrics 5102 Piezoelectric Accelerometers, one on the front right of the vehicle, one on the rear left, and one on the rear right, as shown in Figure 34. The accelerometers on the robot were bolted down, a "direct" accelerometer was placed on the deck with a layer of surfing wax to help the accelerometer adhere to the surface. (This method of adhesion was found adequate to transfer the vibration through some free-vibration tests.) This accelerometer on the deck was placed at the midspan of the bridge and was the same model as that on the robot.

Since this is an operational garage, we could not induce real damage in the structure; instead, we simulated a change in the structure by placing mass at the midspan of the precast T-section. We used 5 gallon buckets filled with coarse aggregate which we weighed in a concrete laboratory prior to delivery to the garage. Because in the simulations we consider only a one dimensional beam, we were careful to apply the load symmetrically, so that the experiment would be similar to the simulations. As shown in Figure 32 and 33, an equal number of buckets are on each side. We varied the total mass level on the bridge from 0 to 360 lbs (163kg) in 40lbs (18.1kg) increments for a total of 10 different mass levels. For the particular experiment that we will describe in this report, we ran 5 trials at each mass level. Then we recharged the robot, and

[3] This is the mass from the manufacturer's website. We did not measure the mass of the particular robot, which would be slightly different due to some modifications.

ran another 5 trials at each mass level. After some data cleansing we had a total of 9 usable trials from each mass level, which are presented in the results.

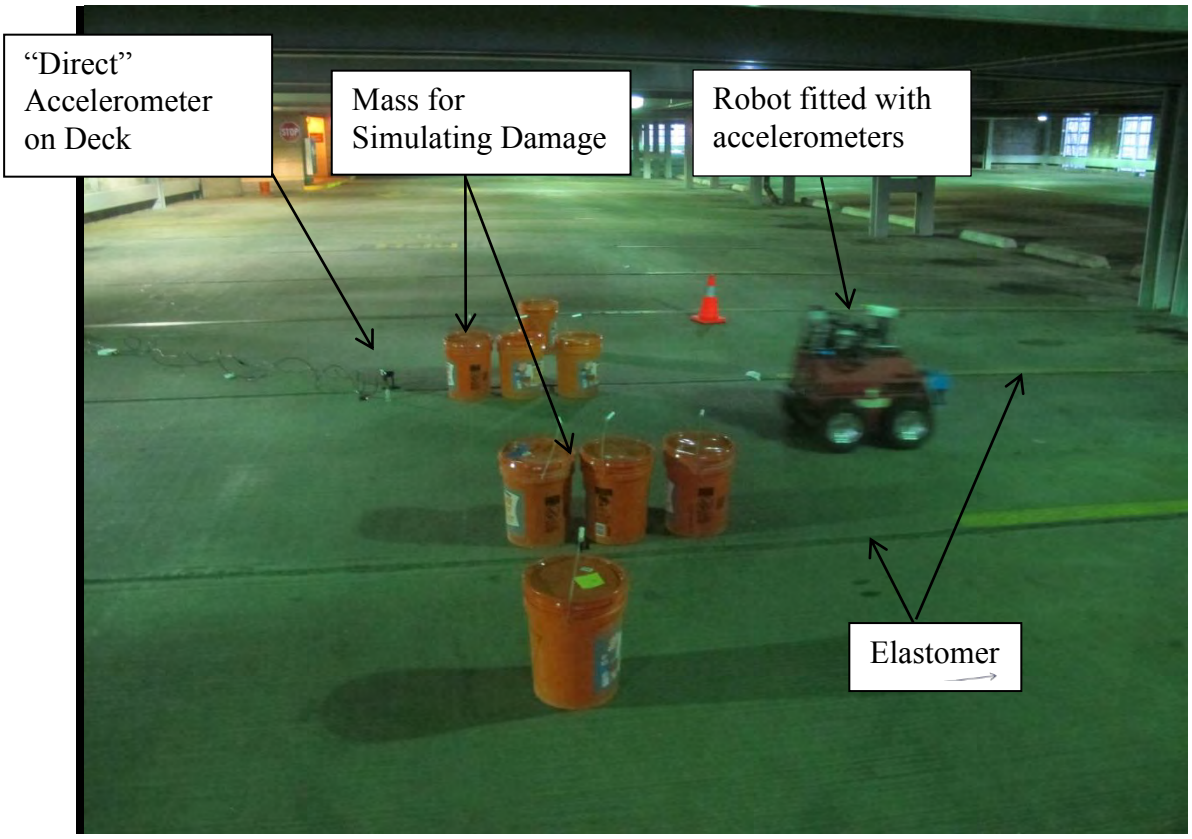

Figure 32 - Experimental Setup for Field Tests

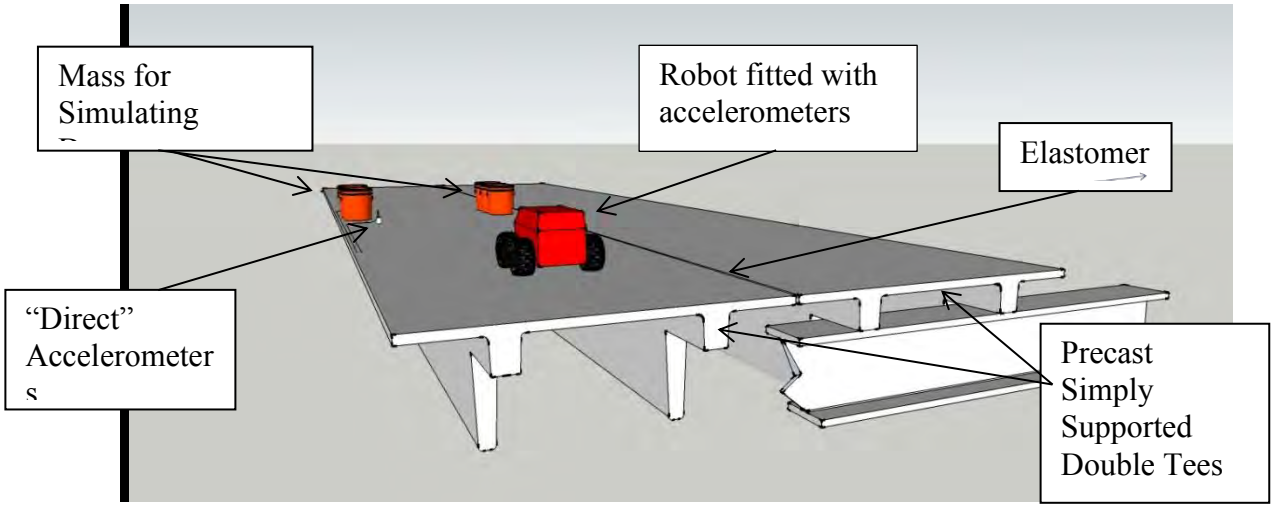

Figure 33 - Illustration Of Garage Experiments

(a) Front View (b) Rear View

Figure 34 - Detail of the Robot

Laboratory Experiments

Laboratory experiments were carried out on the scale model bridge to match the garage experiments. The ratio of added mass to the mass on the bridge deck was the same for both the garage experiment and the lab experiment. The purpose was to see if commonalities existed between the two scales (i.e. laboratory and garage deck) so that the work on the laboratory bridge could be applied more broadly.

Data Selection

In the laboratory experiments, extracting the signal when the car is on the bridge is relatively easy because of a bump as the vehicle enters the main span. Extracting the signal in the garage experiments was much more difficult. Data extraction is an important process to spatially align the data so that the computer can learn feature consistently.

On the lab scale model, we selected the signal when the entire vehicle is on the bridge. The red portion of the signal shown in Figure 36 spans the time from when the back wheel enters the bridge, until the front wheel leaves the bridge. There is a noticeable spike as the front and back wheel leave the bridge due to a small gap between the track on the bridge, and the track on the deceleration ramp, as shown in Figure 35.

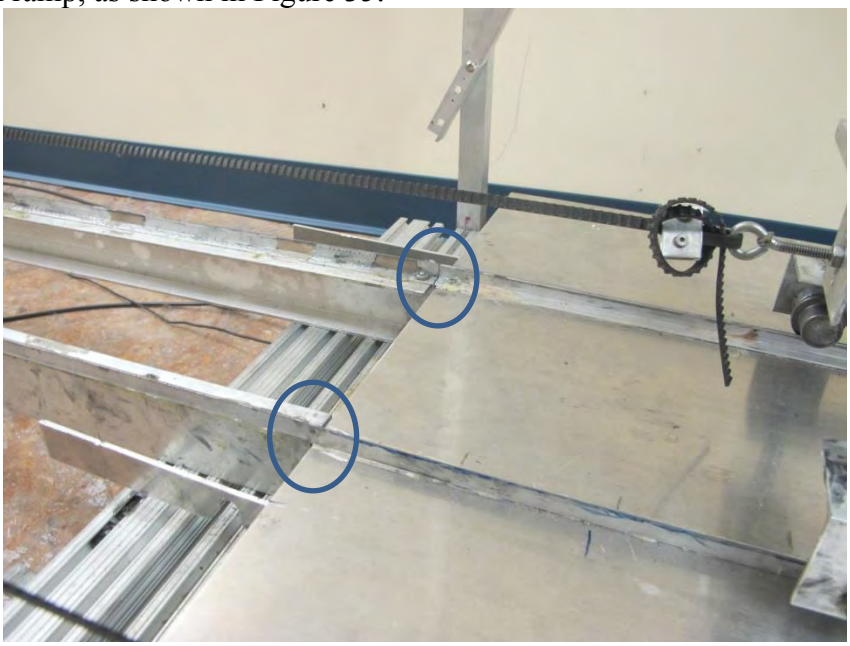

Figure 35 – Slight Gap between Track and Bridge Allowing for Data Selection

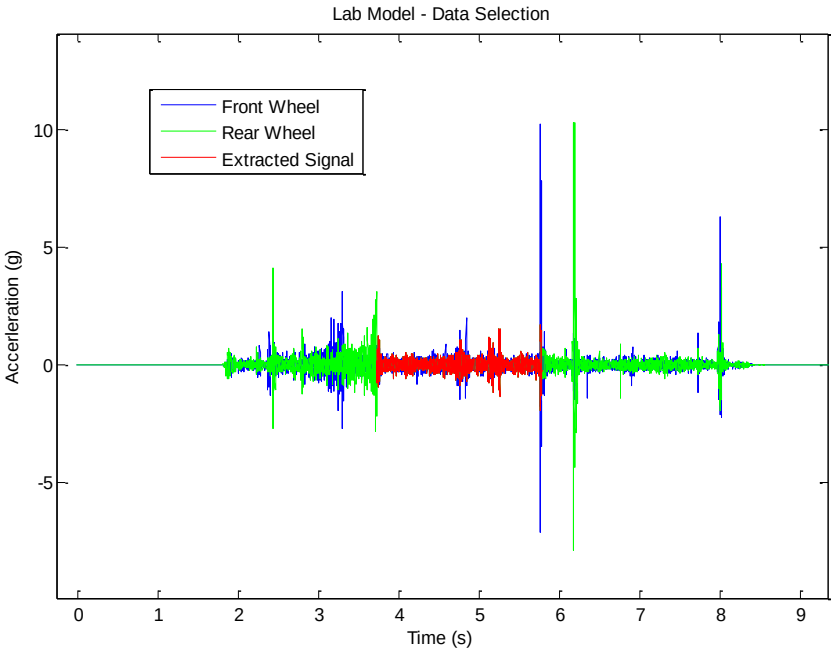

Figure 36 - Data Selection Lab Data

In the garage experiment, the vehicle is constantly on the main span and there is no sudden bump we could use to spatially align the data. In this case our goal was to select data where the vehicle was traveling at a constant speed. Given that the garage span is 15.7 meters in total, we ensured that the robot was moving at a constant speed as it traveled over the middle 7.7 meters, and only extracted that data (there was a 4m area on both sides for accelerating and decelerating.)

To facilitate data extraction we taped a broom handle to the garage deck in order that a peak would be recorded in the accelerometer signal (as shown in Figure 38). The broom handle is shown in Figure 37. The broom handle also allowed us to synchronize the data acquisition system on the garage deck with the system on the robot itself.

Figure 37 - Detail of Broom Handle

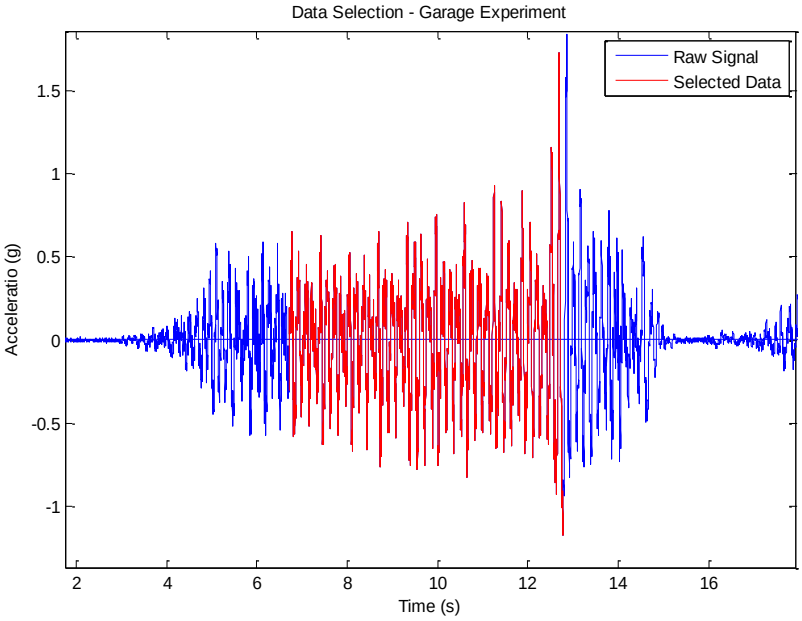

Figure 38 - Garage Experiment Data Selection

Experimental Results

Laboratory Data

The laboratory data largely exhibits changes in the magnitude of various frequencies, as shown in the waterfall plot in Figures 39, 40 and 41. For each of these plots, the signals from the 30 trials at each mass level were averaged, and then the magnitude of the discrete Fourier transform for each mass level was plotted.

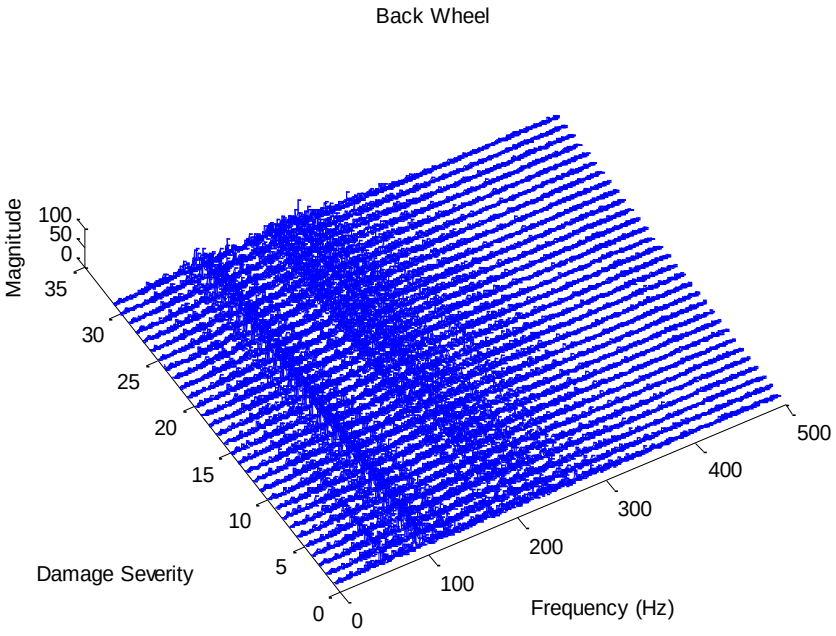

Figure 39 - Waterfall Plot, Data from Back Wheel, Laboratory Experiments (Severity Level 0 corresponds to no mass on the bridge, Level 1 correspond to 5g, up to Level 31 which corresponds to 150g)

The predominant change between damage severity levels (or more precisely, between amounts of mass placed on the structure) is changes in the magnitude of the various frequency values. Although intuitively we might expect the frequencies to shift downward as more mass is added to the structure, we do not actually observe this phenomenon in the data.

It is worth noting that the sampling rate of the accelerometer is around 1600Hz, so it only captures information up to 800 Hz. There appears to be meaningful patterns at least above 400 Hz. In fact sound recording of the vehicle crossing the bridge could also be used to classify the level of applied damage (albeit with slightly lower accuracy).

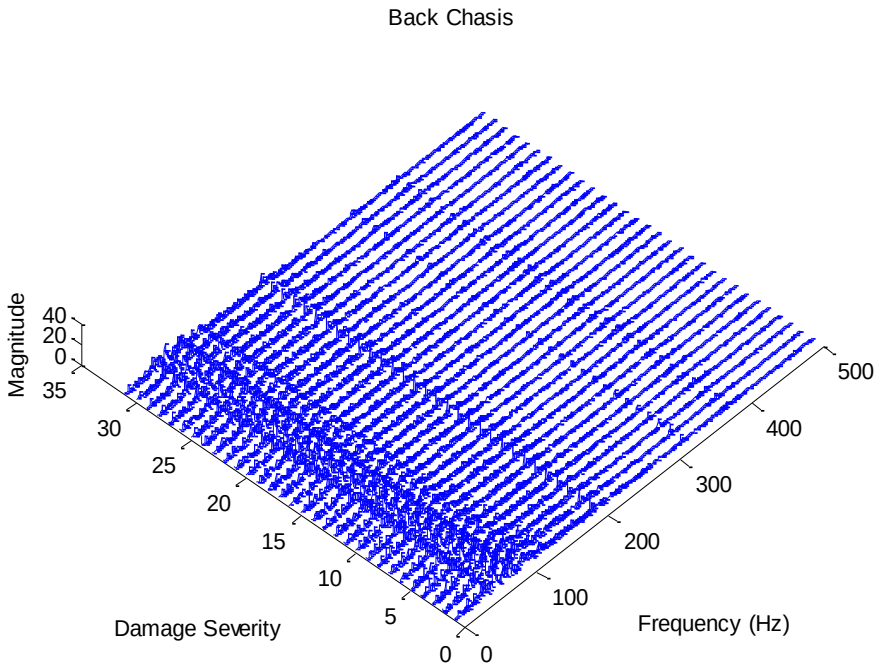

Figure 40 - Waterfall Plot, Data from Back Chassis, Laboratory Experiments (Severity Level 0 corresponds to no mass on the bridge, Level 1 corresponds to 5g… up to Level 31 which corresponds to 150g)

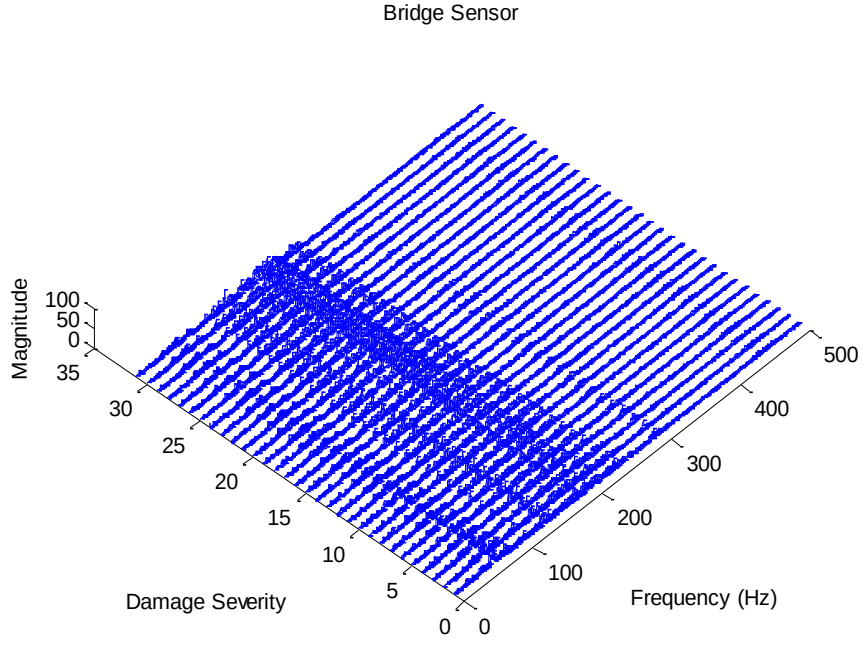

Figure 41 - Waterfall Plot, Data from Bridge Sensor, Laboratory Experiments (Severity Level 0 corresponds to no mass on the bridge, Level 1 corresponds to 5g… up to Level 31 which corresponds to 150g)

Garage Data

The signals from the robot have both a significant peak around 6Hz and a smaller peak around 400 Hz as can be seen in Figure 42. This high frequency component is particularly interesting because we do not see such high frequencies in the signal from the accelerometer on the bridge itself as seen in Figure 43. This 400Hz peak shifts downward around 3.4% as additional mass is placed on the deck, as can been seen in Figure 44. This 3.4% shift mimics the shift we see in the fundamental frequency of the bridge as seen in Figure 45, and the shift in the lower frequencies recorded by the robot as seen in Figure 46. We suspect that the shift in this high frequency component must somehow reflect the additional mass (simulated damage) placed on the bridge, although it is difficult to demonstrate causality conclusively. Because we increased the mass on the bridge in succession (rather than placing random levels of mass) it is possible this shift has something to do with the robot, for example, the battery losing charge over time, and the motor vibrating at a different frequency.

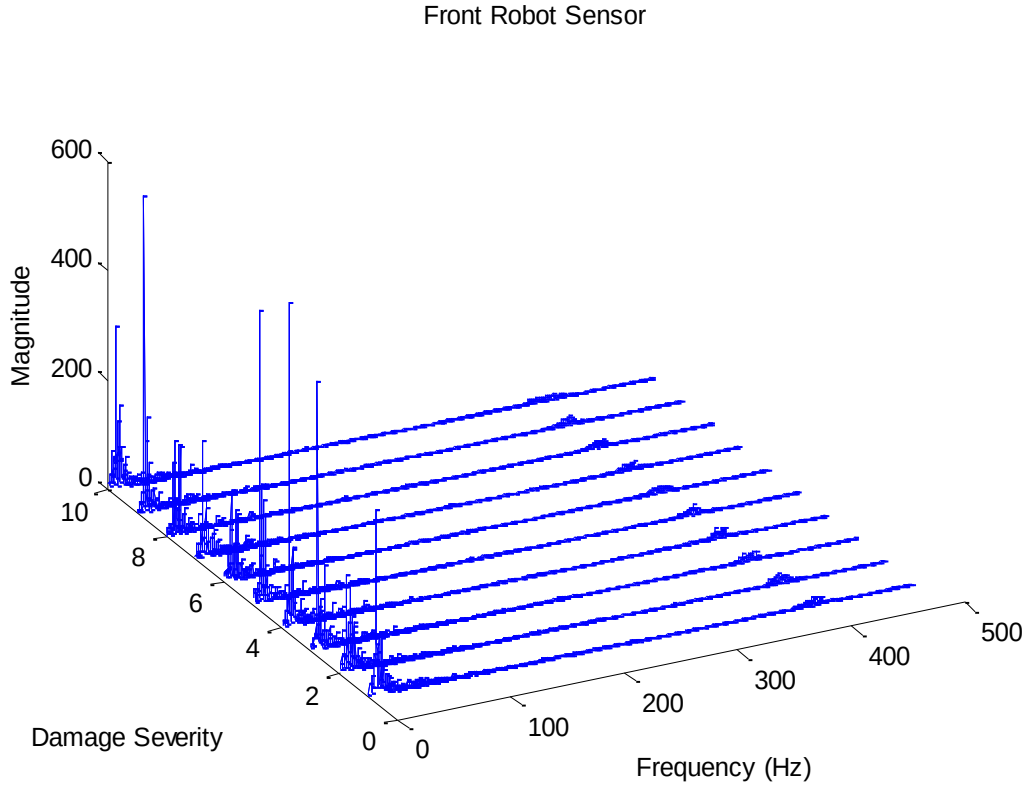

Figure 42 - Waterfall Plot, Signals from the Robot (Severity Level 0 corresponds to no mass on the bridge, Level 1 corresponds to 40lbs… up to Level 10 which corresponds to 360lbs)

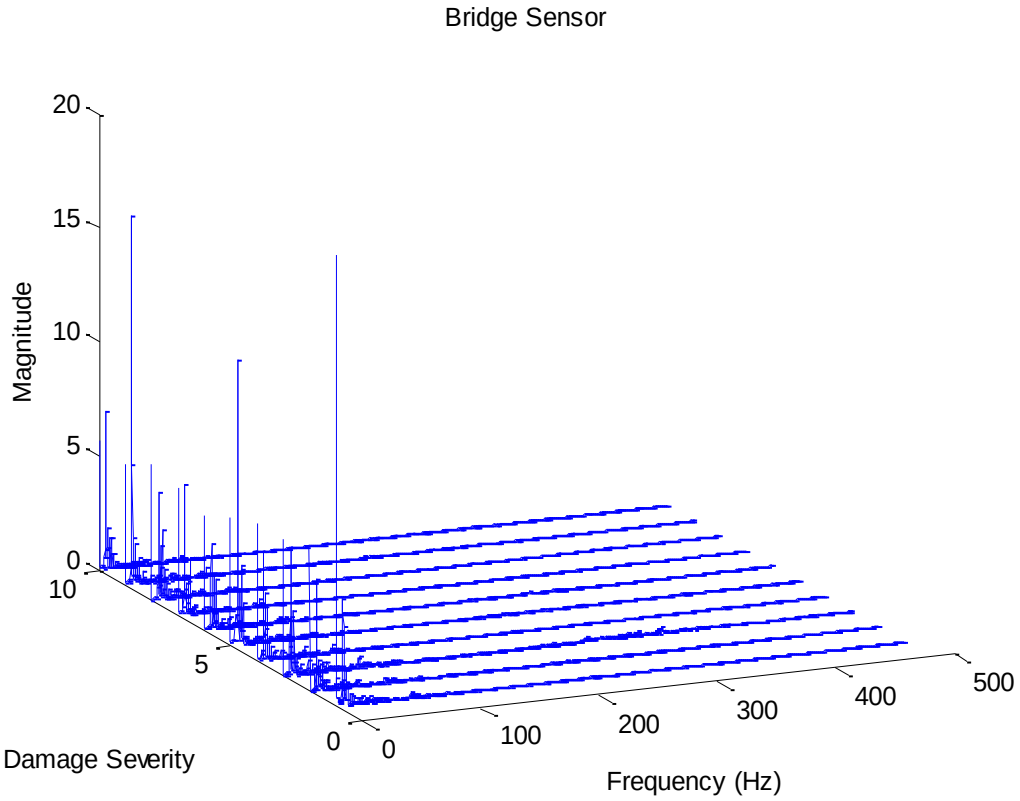

Figure 43 - Waterfall plot, signals from the bridge (Severity Level 0 corresponds to no mass on the bridge, Level 1 corresponds to 40lbs… up to Level 10 which corresponds to 360lbs)

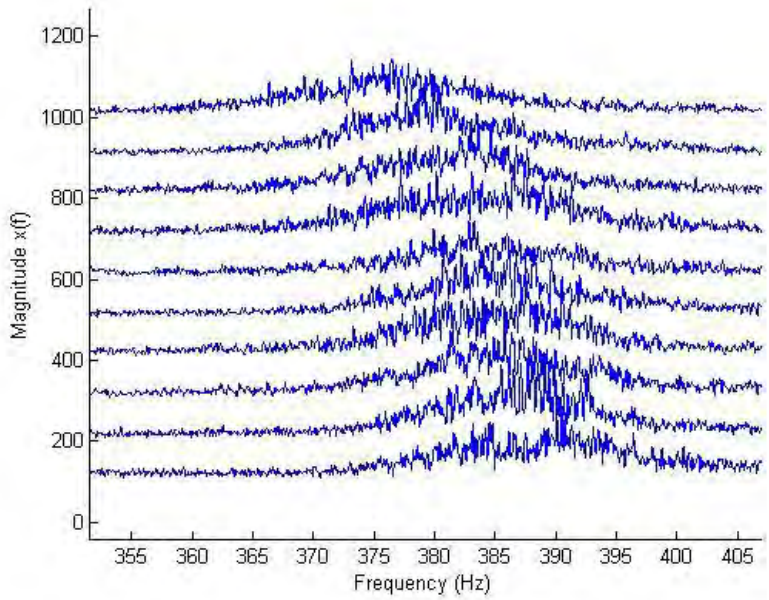

Figure 44 – Magnitude in the Frequency Domain, Each Progressive Level of Damage plotted 100 points above the previous.

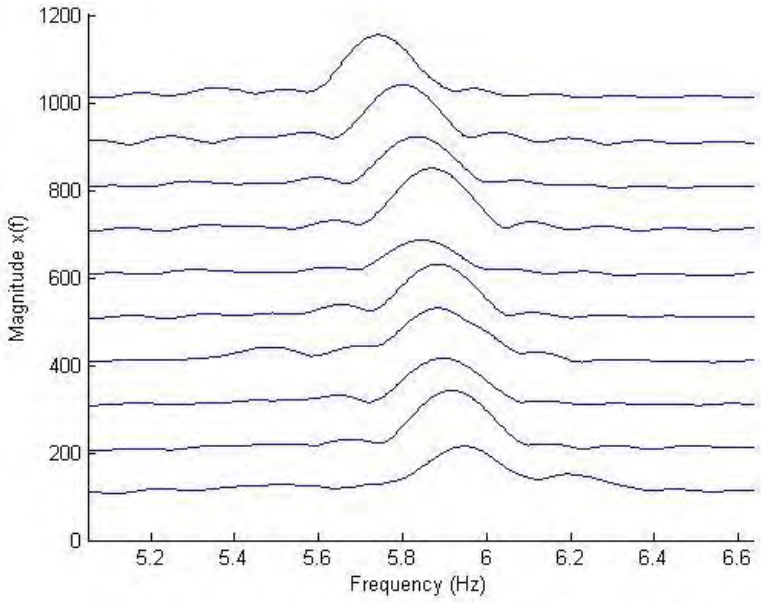

Figure 45 - Magnitude in the frequency domain of the signal directly from the bridge, with each progressive level of mass plotted 100 points above the previous. Each signal was normalized and zero padded, providing a linear interpolation of the data in the frequency domain in order to see the shift more clearly.

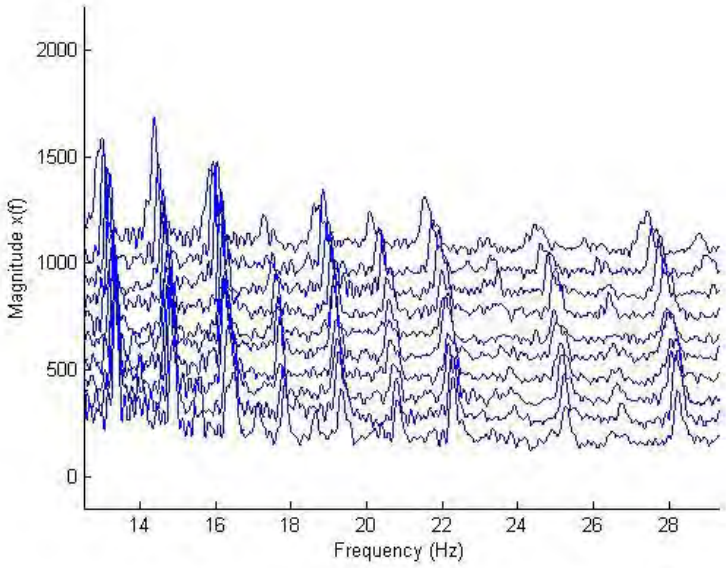

Figure 46 - Magnitude of the frequency domain from the robot, showing detail of the lower frequencies. This signal was also zero padded. A clear downward shift can be seen at many frequency levels.

Port Authority Project

With the help of the University Transportation Center (T-SET), we have begun to collect indirect data from Pittsburgh's Light Rail Line. This project has been facilitated through the connections of T-SET's staff, and by the leadership of the Port Authority, in particular Bill Miller and David Kramer.

Both the garage and the laboratory tests were simply experiments; they were an attempt to create the type of scenario we might encounter later on. The Pittsburgh Light Rail Line, known as "the T," offers an opportunity to continuously collect real data in an operational environment. While previously we had simulated temperature changes and damage scenarios, this project tries to account for natural temperature change and detect genuine damage. The downside of such a system is that we lose control over the bridges. While in the lab we can choose when to simulate damage, in this data collection process we have to wait until a bridge is damaged or until construction occurs on one of the bridges. Our data collection efforts began in September 2013. No work on this topic has been published. We will give a short summary of the work thus far and discuss the potential impact this project might have on the Structural Health Monitoring community.

Our first task was to design robust hardware capable of efficiently logging high frequency data. In September 2013 we placed our data collection system on a single train of the Light Rail System. The data acquisition system consists of an LGX AG150-V automotive computer (Figure 47a), a GPS module (UsGlobalSat BU-353), and two National Instruments Dynamic Signal Acquisition Modules (Figure 47b). There are two types of accelerometers; a triaxial accelerometer (PCB 354C03) placed on the middle truck of the train and two uniaxial accelerometers (5102 Vibra-metrics) placed near the data acquisition system. A schematic of the whole system as implemented in the train car.

(a) (b)

Figure 47 - Installed Equipment on the Train

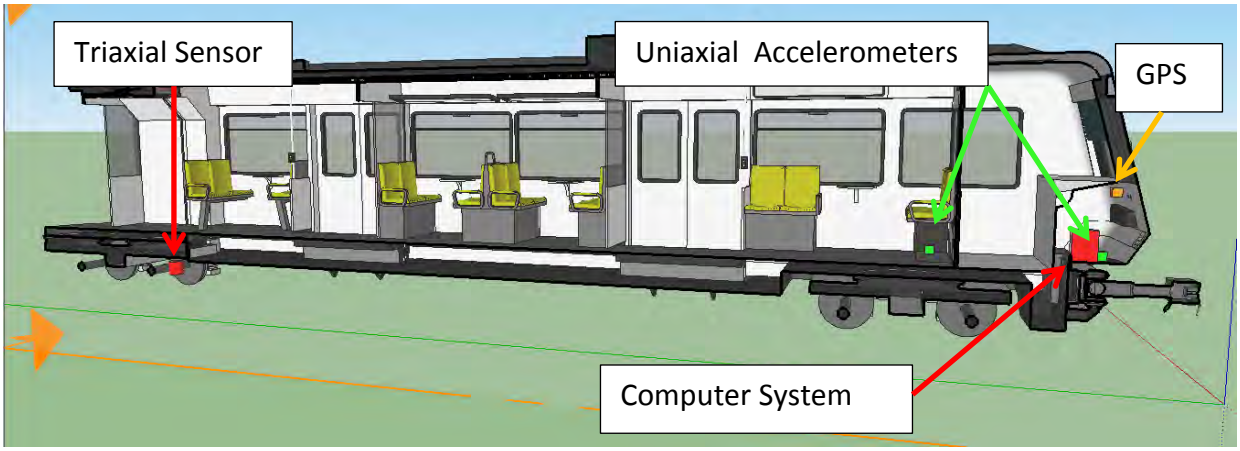

Figure 48 -Schematic of the Sensors Placed on the light rail car

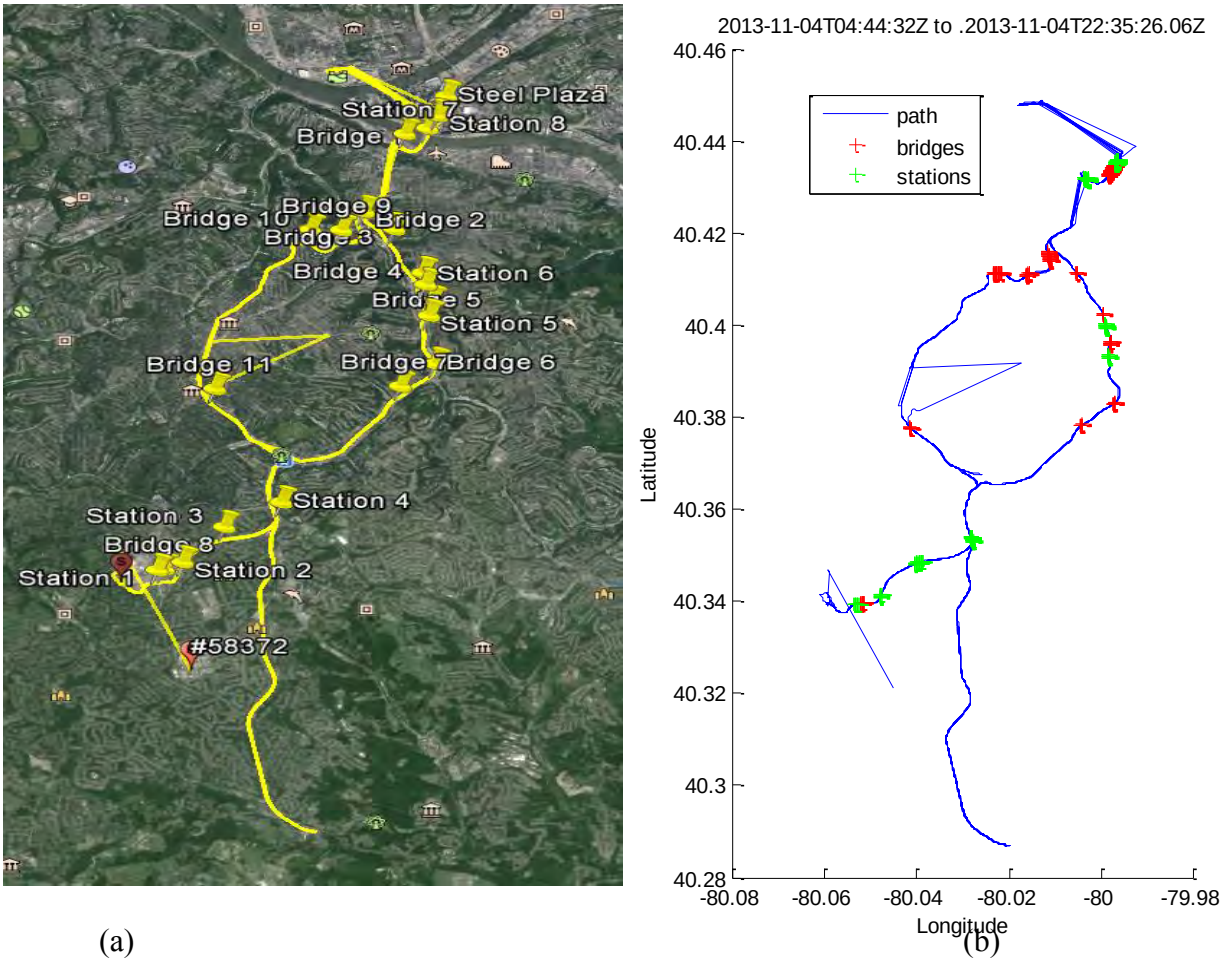

(a) (b)

Figure 49 - (a) GPS data overlaid onto Google Maps (b)GPS plotted with bridge locations identified. When the train enters into a tunnel, the GPS accuracy deteriorates rapidly. This is why there are some straight lines veering away from the track.

We collect accelerometer signals from the train continuously at 1651Hz. This generates lots of data—in the range of 10Gb per day. We store all this raw data, but then extract the portion of the signal when the train is actually on one of the 11 bridges, and place this data into a

database which is accessible via the internet. Details of the GPS output (overlaid onto Google Maps) and the portion of the signal we extract from the bridges are shown in figure 49.

We look at the signal of the train both as it crosses the bridges and while it remains in the station. We expect to see lower accelerations while the train is stopped in the station. This serves as important verification that we are extracting the correct portion of the signals. This is shown in Figure 50.

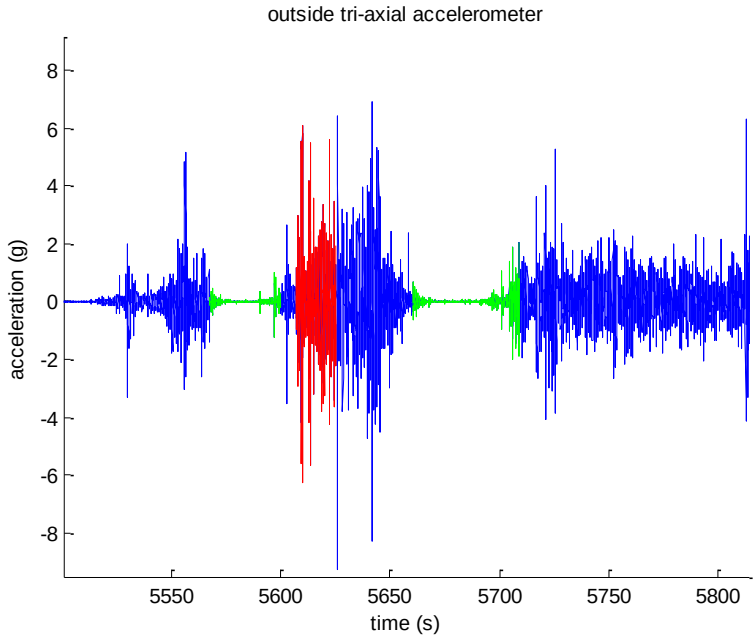

Figure 50 - This is an example of the raw signal collected from the train. Areas when the GPS shows the train to be in the station are shown in green. Areas where the GPS shows the train to be over a bridge are shown in red.

-

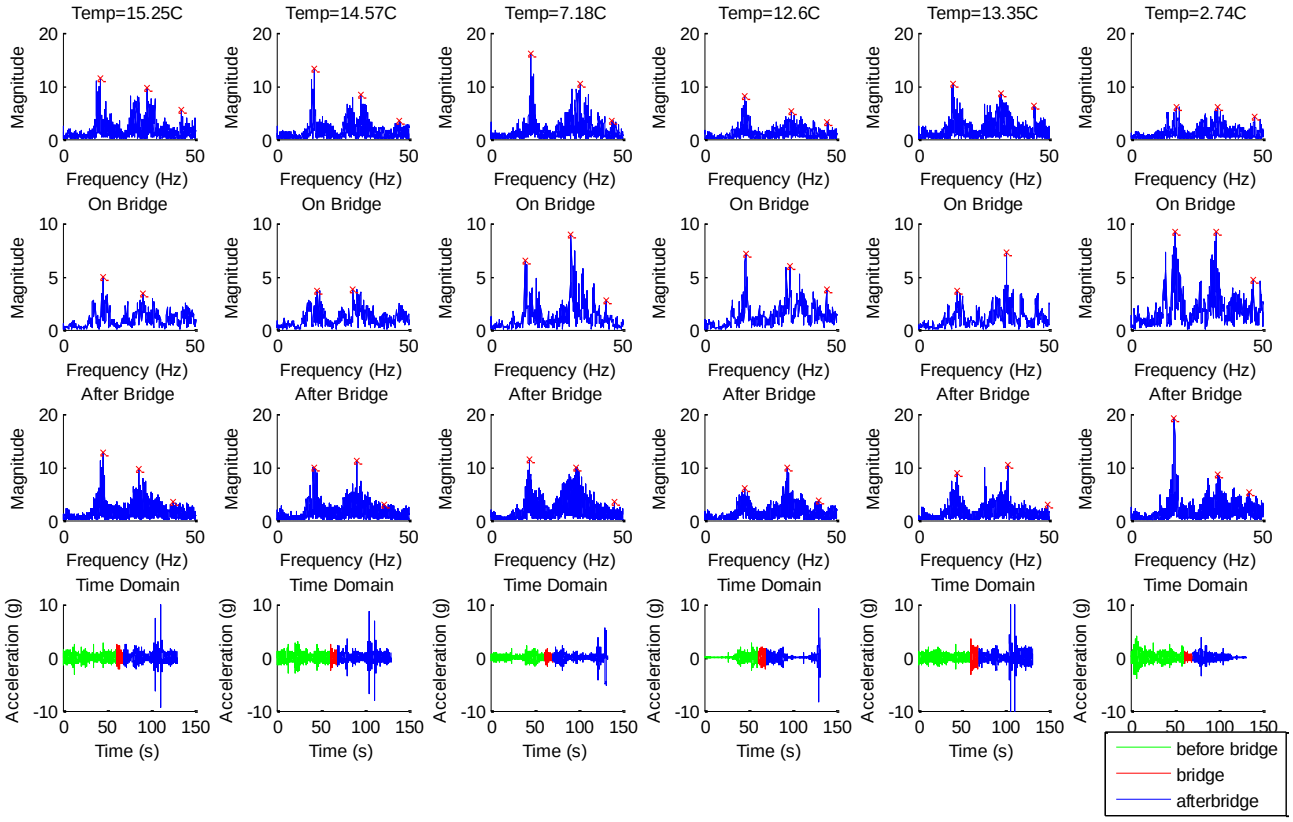

Figure 51 - Example Signals over a bridge at Different Temperatures

In addition to the accelerometer signals, we record environmental data in our database. This includes temperature, humidity, visibility and wind speed. These values are extrapolated from NOAA weather data, and triangulated to provide a "hyperlocal" forecasts. At the moment we have limited crossing for each bridge so it is difficult to know which changes are due to stochastic changes and which are due to temperature changes. Figure 51 shows six crossings over a single bridge. The frequency spectrum is shown before the train is on the bridge, while the train is on the bridge, and after the train crosses the bridge. The fourth plot at the bottom shows the time domain signal with different colors showing the different sections of the signal.

Recording the baseline dynamic response of the bridge could take a long time. As no experiments will be performed on the operational system, validation will occur when the algorithm is able to detect a fault of statistical significance that can be verified. Currently we are continuing to collect data about the bridges—to better understand how they behave at different temperatures.

In the coming months, we will apply existing signal processing and machine learning techniques to this dataset as well as try to develop new approaches.

Expected Significance: Ideally, we will find that we can use this technique to determine severity of damage, and location of damage. However, simply detecting a statistically significant change in the structure would be a huge breakthrough. This would allow the Port Authority to

prioritize when to send structural inspectors to a bridge. If nothing has changed in the bridge since the last inspection—this technique could save money. If something has happened before a scheduled inspection—this could save lives.

Conclusions

Over the course of this 23 month project we have made significant contributions in Indirect Structural Health Monitoring. This technology could provide an inexpensive way to continuously monitor the aging infrastructure both in the United States and abroad.

We have done fundamental work on understanding the vehicle bridge interaction, and the machine learning algorithms which will best detect changes in the bridge. And we have done practical work on efficiently collecting and storing data, and how we can work with agencies like the Port of Authority of Allegheny County to help them tackle the challenge of maintaining their assets.

Tasks 1 through 3 focused on a laboratory scale model and verified that indirect monitoring is in fact a viable method. We have shown we can detect small changes in a laboratory bridge in presence of environmental variables or in cases where there is uncertainty in the data. And we explored numerous types of damage including changes in rotational restrain, changes in damping ratio, and changes in the mass of the bridge. We have explored ways in which the fundamental features of the bridge change with damage so that the general trend can be understood. Without environmental variation, we can consistently achieve above 90 percent classification on the lab scale model. However adding some temperature variation reduces the classification accuracy down to roughly 70 percent. We have shown that the best algorithm thus far is a semi-supervised multiresolution classifier with label propagation and weighting. New methods using graph models may improve the accuracy while accounting for changes in temperature.

Task 4, collecting data from new sources, began only towards the end of this project, although we hope to continue this work with future support from the University Transportation Center. By applying our novel algorithms to the operational data we are collecting, we hope to be able to detect subtle structural changes, while avoiding false positives due to changes in the environmental conditions. We also hope to apply our methodology to damage detection of the railway track, in addition to our work on bridges. This problem is of considerable interest to the rail industry, not only in terms of maintenance, but also safety.

As many of the bridges in our national inventory approach their design life, ubiquitous computing and cheap sensors combined with novel signal processing algorithms may offer a solution. The work we have done with this funding, in particular, the papers we have published, bring us one step closer to harnessing this data to understand the state of our infrastructure. .

Perhaps the most exciting outcome of this research is the finding that algorithms developed for image processing can be applied to bridge monitoring. The semi-supervised approach described in the second half of Task 3 shows that even when only 10 percent of the data is labeled, we can still classify the state of the bridge. We hope we can continue to find such exciting connect when analyzing the data we gather from the Port Authority Trains. And we hope to develop new tools which help federal, state and local agencies maintain a safe and efficient infrastructure on tightly constrained budgets.

Bibliography

Abdi. H., Williams, and L.J. "Principal component analysis." Wiley Interdisciplinary Reviews: Computational Statistics, 2: 433–459, 2010.

Ahron, M. Elad, M. Bruckstein, A. "K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation", IEEE Transactions on Signal Processing, Vol. 54, No. 11. November 2006.

Alippi, C., and Galperti, C. 2008: "An Adaptive System for Optimal Solar Energy Harvesting in Wireless Sensor Network Nodes." *Circuits and Systems I: Regular Papers, IEEE Transactions On* 55, no. 6: 1742–1750.

Akinci, B., Hendrickson, C., and Karaesmen, I. 2003. "Exploiting Motor Vehicle Information and Communications Technology for Transportation Engineering." *Journal of Transportation Engineering* 129, no. 5: 469–474.

Beeby, S. P., Tudor, M. J. and White, N. M. 2006 "Energy Harvesting Vibration Sources for Microsystems Applications." *Measurement Science and Technology* 17, no. 12: R175–R195.

Bilello, C., and Bergman L. A. 2004. Vibration of damaged beams under a moving mass: theory and experimental validation. *Journal of Sound Vibration* 274 (Jul 1): 567-582.

Bishop, C. M. "Pattern Recognition and Machine Learning", Information Science and Statistics, Springer, 2006.

Brownjohn, J.M.W. 2007. Structural health monitoring of civil infrastructure. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 365, no. 1851 (February 15): 589-622.

Bu, J. Q., Law, S. S. and Zhu, X. Q. 2006. "Innovative Bridge Condition Assessment from Dynamic Response of a Passing Vehicle." *Journal of Engineering Mechanics* 132 (12) (December): 1372–1379.

Burges, Christopher J. C, "A Tutorial on Support Vector Machines for Pattern Recognition", Data Mining and Knowledge Discovery 2:121–167, 1998

Carden, E. P., and Fanning, P. 2004. "Vibration Based Condition Monitoring: A Review." Structural Health Monitoring 3 (4) (December 1): 355–377.

Casciati, F., and Giordano, M. 2010. *Structural Health Monitoring 2010: Proceedings of the Fifth European Workshop on Structural Health Monitoring Held at Sorrento, Naples, Itlay, June 28-july 4, 2010*.

Cerda, F., Garrett, J., Bielak, J., Bhagavatula, R., and Kovačević. J. 2010. "Exploring Indirect Vehicle-bridge Interaction for Bridge SHM." *Bridge Maintenance, Safety, Management and Life-Cycle Optimization - Proceedings of the 5th International Conference on Bridge Maintenance, Safety and Management:* 708–715.

Cerda, F., Chen, S, Bielak, J, Garrett, J, Rizzo, P, Kovačević, J. 2012 *"Indirect structural health monitoring of a simplified laboratory-scale bridge model*". Submitted to Smart Structures and Systems (Special Issue).

Chebira, A., and Kovačević, J. 2008. "Frames in Bioimaging." *Proc. CISS, Princeton, NJ, Mar. 2008.*

Chebira, A., Barbotin, Y., Jackson, C., Merryman, T., Srinivasa, G., Murphy, R., and Kovačević, J. 2007. "A Multiresolution Approach to Automated Classification of Protein Subcellular Location Images." *BMC Bioinformatics* 8 (1): 210.

Coifman, Y. Meyer, S. Quake, and M. V. Wickerhauser, "Signal processing and compression with wavelet packets," Tech. Rep., Yale Univ., 1991.

Cornwell, P., Farrar, C. R. Doebling, S. W. and Sohn, H. 1999. "Environmental Variability of Modal Properties." *Experimental Techniques* 23 (6): 45–48.

Daubechies, I. 1992. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics. http://portal.acm.org/citation.cfm?id=130655.

Deraemaeker, A., Reynders, E.,De Roeck, G., and Kullaa, J. 2008."Vibration-based Structural Health Monitoring Using Output-only Measurements Under Changing Environment." Mechanical Systems and Signal Processing 22, no. 1: 34–56.

Doebling, S. W., Farrar, C. R., and Prime, M. B. 1998. "A Summary Review of Vibration-based Damage Identification Methods." *Identification Methods," The Shock and Vibration Digest* 30: 91–105.

Domingos, Pedro and Michael Pazzani, "On the optimality of the simple Bayesian classifier under zero-one loss", Machine Learning, 29:103–137, 1997

Duda, R., Hart, P., and Stork, D. 2000. *Pattern Classification (2nd Edition)*. Wiley-Interscience.

FHWA 2001. Reliability of Visual Inspection for Highway Bridges - FHWA-RD-01-020 http://www.fhwa.dot.gov/publications/research/nde/01020.cfm

Farhey, D. 2005. "Bridge Instrumentation and Monitoring for Structural Diagnostics." *Structural Health Monitoring* 4 (4) (December 1): 301–318.

----2007. "Quantitative Assessment and Forecast for Structurally Deficient Bridge Diagnostics." *Structural Health Monitoring* 6 (1) (March 1): 39–48.

Farrar, C. R., and Jauregui, D.. 1998. "Comparative Study of Damage Identification Algorithms Applied to a Bridge: I. Experiment." *Smart Materials and Structures* 7 (5) (October 1): 704–719.

Farrar, C. R., and Worden, K. 2007. "An Introduction to Structural Health Monitoring." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365 (1851) (February 15): 303–315.

Frangopol, D. M., Strauss, A., and Kim, S. 2008. "Bridge Reliability Assessment Based on Monitoring." *Journal of Bridge Engineering* 13 (3) (May): 258–270.

Frangopol, D.M., Sause, R., and Kusko, C. eds. 2010. *Bridge Maintenance, Safety and Management - IABMAS'10: Proceedings of the Fifth International IABMAS Conference, Philadelphia, USA, 11-15 July 2010*. 1st ed. CRC Press.

Friswell, M. I. 2007. Damage identification using inverse methods. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 365, no. 1851 (February 15): 393-410.

Fugate, M. L., Sohn H., & Farrar C. R. 2000. Unsupervised learning methods for vibration-based damage detection. Presented at IMAC, San Antonio, Texas. February 7-10.

Haralick, R.M, Shanmugan, K., and Dinstein, I. (1973). "Textural Features for Image Classification." *IEEE Trans. on Systems, Man, and Cybernetics* 3 (6): 610–621.

Hilbe, Joseph M, "Logistic Regression Models", Chapman & Hall/CRC Press, ISBN 978-1-4200-7575-5, 2009

Ikenaga, S., Lewis, F.L., Campos, J., and Davis, L. 2000. "Active Suspension Control of Ground Vehicle Based on a Full-vehicle Model." *Proceedings American Control Conference, 2000. Proceedings of the 2000*. Vol 6:4019–4024 vol.6.

Isemoto, R., Kim,C.W., and Sugiura, K. 2010. "Abnormal Diagnosis of Bridges Using Traffic-induced Vibration Measurements." *The Twenty-Third KKCNN Symposium on Civil Engineering*, ,Taipei, Taiwan:189-192.

Kellogg, R.A., Chebira, A., Goyal, A., Cuadra, P.A. , Zappe, S.F., Minden, J.S., and Kovačević, J. 2007. "Towards an Image Analysis Toolbox for High-throughput Drosophila Embryo RNAi Screens." *Proc. IEEE Intl. Symp. Biomed. Imaging, Arlington, VA,: 288-291*.

Kim, C.W., and Kawatani, M. 2008. "Pseudo-static Approach for Damage Identification of Bridges Based on Coupling Vibration with a Moving Vehicle." *Structure and Infrastructure Engineering* 4 (5): 371–379.

Kim, C.W., Kawatani, M., and Fujimoto, T. 2010. "Identifying Bending Stiffness Change of a Beam Under a Moving Vehicle." *Bridge Maintenance, Safety, Management and Life-Cycle Optimization - Proceedings of the 5th International Conference on Bridge Maintenance, Safety and Management*, 180–180.

Kullback, S, and Leibler, R.A. 1951. "On Information and Sufficiency." *Ann. Math. Statist* 22: 79–86.

Law, S. S., and Zhu X. Q. 2004. Dynamic behavior of damaged concrete bridge structures under moving vehicular loads. *Engineering Structures* 26, no. 9 (July): 1279-1293.

Law, S. S, and Zhu. X. Q. 2005. Nonlinear characteristics of damaged concrete structures under vehicular load. *Journal of Structural Engineering-ASCE* 131, no. 8 (August): 1277-1285.

Li, H., Wekezer, J. and Kwasniewski, L. 2008. "Dynamic Response of a Highway Bridge Subjected to Moving Vehicles." *Journal of Bridge Engineering* 13 (5): 439–448.

Lin, C.W., and Yang, Y.B. 2005. "Use of a Passing Vehicle to Scan the Fundamental Bridge Frequencies: An Experimental Verification." Engineering Structures 27 (13) (November): 1865–1878.

Lynch, J. P. 2007. An overview of wireless structural health monitoring for civil structures. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 365, no. 1851 (February 15): 345-372.

Mahmoud, M. A., and Abou Zaid, M. A.. 2002. Dynamic response of a beam with a crack subject to a moving mass. *Journal of Sound and Vibration* 256, no. 4 (September 26): 591-603.

Mal, A., Ricci, F., Banerjee, S., and Shih, F. 2005. "A Conceptual Structural Health Monitoring System Based on Vibration and Wave Propagation." *Structural Health Monitoring* 4, no. 3: 283–293.

Mallat, S. 1999. *A Wavelet Tour of Signal Processing*. 2nd ed. Academic Press.

McGetrick, P.J., González, A., and O'Brien, E.J. 2009. "Theoretical Investigation of the Use of a Moving Vehicle to Identify Bridge Dynamic Parameters." *Insight: Non-Destructive Testing and Condition Monitoring* 51 (8): 433–438.

McGetrick, P.J., Kim, C.W., and O'Brien, E.J. 2010. "Experimental Investigation of the Detection of Bridge Dynamic Parameters Using a Moving Vehicle." *The Twenty-Third KKCNN Symposium on Civil Engineering,* Taipei, Taiwan:177-180.

Melhem, H., and Kim, H. 2003. "Damage Detection in Concrete by Fourier and Wavelet Analyses." *Journal of Engineering Mechanics* 129 (5) (May): 571–577.

Meruane, V., and Heylen, W. 2012. "Structural damage assessment under varying temperature conditions." *Structural Health Monitoring-an International Journal* 11 (3) (May): 345–357.

Peeters, B., and De Roeck, G. 2000 . “One Year Monitoring of the Z24-bridge : Environmental Influences Versus Damage Events.” In *SPIE Proceedings Series*, 1570–1576. Society of Photo-Optical Instrumentation Engineers.

Miyamoto, A., and Yabe, A. 2011. “Bridge Condition Assessment Based on Vibration Responses of Passenger Vehicle.” Journal of Physics: Conference Series 305 (July 19): 012103.

Reda Taha, M. M., Noureldin, A., Lucero, J. L. and Baca, T. J. 2006. “Wavelet Transform for Structural Health Monitoring: A Compendium of Uses and Features.” *Structural Health Monitoring* 5 (3) (September 1): 267–295.

Rytter, A. (1993). “Vibration Based Inspection of Civil Engineering Structures”. Denmmark: Department of Building Technology and Structural engineering, Aalborg University.

Serker Kamrujjaman, N. H. M., Wu, Z. and Li, S. 2010. “A Nonphysics-based Approach for Vibration-based Structural Health Monitoring Under Changing Environmental Conditions.” *Structural Health Monitoring* 9 (2) (March 1): 145–158.

Siringoringo, D.M., and Fujino, Y. 2012. “Estimating Bridge Fundamental Frequency from Vibration Response of Instrumented Passing Vehicle: Analytical and Experimental Study.” Advances in Structural Engineering 15 (3): 417–433.

Sohn, H. 2007. “Effects of Environmental and Operational Variability on Structural Health Monitoring.” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365 (1851) (February 15): 539–560.

Sugiyama, M. “Dimensionality reduction of multimodal labeled data by local Fisher discriminant analysis”, Journal of Machine Learning Research, vol.8 (May), pp.1027-1061, 2007.

Toshinami, T., Kawatani, M. and Kim, C.W. 2010. “Feasibility Investigation for Identifying Bridge’s Fundamental Frequencies from Vehicle Vibrations.” In *Bridge Maintenance, Safety, Management and Life-Cycle Optimization - Proceedings of the 5th International Conference on Bridge Maintenance, Safety and Management*, 329–334.

Van der Auweraer, H., and Peeters B. 2003. International Research Projects on Structural Health Monitoring: An Overview. Structural Health Monitoring 2, no. 4 (December 1): 341-358.

Vetterli, M and Kovačević, J. 1995. *Wavelets and Subband Coding*. Signal Processing. Prentice Hall PTR. http://waveletsandsubbandcoding.org.

Worden, K., &. Dulieu-Barton J. M. 2004. An Overview of Intelligent Fault Detection in Systems and Structures. Structural Health Monitoring 3, no. 1 (March 1): 85-98.

Worden, K., & Manson G. 2007. The application of machine learning to structural health monitoring. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 365, no. 1851 (Feb. 15): 515-537.

Wright, J. Yang, A. Ganesh, A. Sastry, S. and Ma, Y. “Robust Face Recognition via Sparse Representation” IEEE Transactions on Pattern Analysis and Machine intelligence, February 2009.

Wright, J., Yi Ma, Mairal, J., Sapiro, G., Huang, T.S., and Shuicheng Yan. "Sparse Representation for Computer Vision and Pattern Recognition." *Proceedings of the IEEE* 98, no. 6 (June 2010): 1031 –1044.

Yan, A.-M., G. Kerschen, P. de Boe, and J.-C. Golinval. 2005a "Structural Damage Diagnosis Under Varying Environmental conditions---Part I: A Linear Analysis." Mechanical Systems and Signal Processing 19: 847–864.

———2005b. "Structural Damage Diagnosis Under Varying Environmental Conditions---part II: Local PCA for Non-linear Cases." Mechanical Systems and Signal Processing 19: 865–880.

Yabe, A and Miyamoto, A. 2010. "Development of a bridge condition assessment system by using city bus." *Bridge Maintenance, Safety, Management and Life-Cycle Optimization - Proceedings of the 5th International Conference on Bridge Maintenance, Safety and Management*: 2056–2063.

Yang, Y.B., and Chang, K.C. 2009. "Extraction of Bridge Frequencies from the Dynamic Response of a Passing Vehicle Enhanced by the EMD Technique." Journal of Sound and Vibration 322 (4-5): 718–739.

Yang, Y. B., Lin, C.W., and Yau, J. D. 2004. "Extracting Bridge Frequencies from the Dynamic Response of a Passing Vehicle." *Journal of Sound and Vibration* 272 (3-5) (May 6): 471–493.

Yang, Y.B., and Lin, C.W. 2005. "Vehicle-bridge Interaction Dynamics and Potential Applications." *Journal of Sound and Vibration* 284 (1-2) (June 7): 205–226.

Yin, S.H., and Tang, C.Y. 2011. "Identifying Cable Tension Loss and Deck Damage in a Cable-Stayed Bridge Using a Moving Vehicle." Journal of Vibration and Acoustics 133 (2): 021007.

Zhu, X.Q., and Law, S.S. (2006). "Wavelet-based Crack Identification of Bridge Beam from Operational Deflection Time History." *International Journal of Solids and Structures* 43 (7-8) (April): 2299–2317.

Zhu, X., and Rizzo, P. (2011). "Guided Waves for the Health Monitoring of Sign Support Structures Under Varying Environmental Conditions," *Structural Control and Health Monitoring*, Available online: 21 AUG 2011

Zhu, X. "Semi-supervised learning literature survey", Tech. Rep. 1530, Univ. Wisconsin-Madison, 2005.

Zhu, X. and Ghahramani, Z. "Learning from labeled and unlabeled data with label propagation," Tech. Rep., Carnegie Mellon Univ., 2002.

Appendix

Indirect structural health monitoring of a simplified laboratory-scale bridge model

Fernando Cerda*[1], Siheng Chen[2], Jacobo Bielak[3], James H. Garrett[3], Piervincenzo Rizzo[5] and Jelena Kovačević[2,4]

[1] Universidad de Concepción, Concepción, Chile
Departments of [2]Electrical and Computer Engineering, [3]Civil and Environmental Engineering, and [4]Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[5]Department of Civil and Environmental Engineering, University of Pittsburgh, PA 15261, USA

Abstract. An indirect approach is explored for structural health bridge monitoring allowing for wide, yet cost-effective, bridge stock coverage. The detection capability of the approach is tested in a laboratory setting for three different reversible proxy types of damage scenarios: changes in the support conditions (rotational restraint), additional damping, and an added mass at the midspan. A set of frequency features is used in conjunction with a support vector machine classifier on data measured from a passing vehicle at the wheel and suspension levels, and directly from the bridge structure for comparison. For each type of damage, four levels of severity were explored. The results show that for each damage type, the classification accuracy based on data measured from the passing vehicle is, on average, as good as or better than the classification accuracy based on data measured from the bridge. Classification accuracy showed a steady trend for low (1-1.75 m/s) and high vehicle speeds (2-2.75 m/s), with a decrease of about 7% for the latter. These results show promise towards a highly mobile structural health bridge monitoring system for wide and cost-effective bridge stock coverage.

Keywords: Indirect SHM, laboratory experiment, damage detection, classification.

1. Introduction

The importance and need for bridge inspection and monitoring has increasingly become more apparent in the aftermath of catastrophic collapses, such as those that have occurred recently around the world (I-35W bridge over Mississippi River, USA, Aug 1 2007; Shershah Bridge, Pakistan, Sep 1 2007; Harp Road bridge, USA, Aug 15 2007; Loncomilla Bridge, Chile, Nov 18, 2004). In many countries, the assessment of bridges is done at fixed time intervals. For example, in the United States, bridges are visually inspected every two years, and then, if signs of deterioration are visible, more accurate evaluation is conducted using commercially available nondestructive evaluation techniques such as: acoustic emission, electromagnetic testing or liquid penetrant testing, among others. In the past two decades, the interest of both researchers and practitioners in structural health monitoring (SHM) methods has escalated. SHM allows for the evolution of the maintenance practice from "time-based" to "condition-based", which implies that a sensing system, integrated with the structure or the mechanical system, performs periodically-spaced measurements. The statistical analysis of damage-sensitive features extracted from these measurements enables one to determine the current state of system health and to notify in real time when degradation or damage occurs (Doebling et al 1998, Carden and Fanning 2004, Farrar and Worden 2007). This approach can be referred to as *direct bridge monitoring*. The objective of any direct bridge monitoring system is to establish the state of a bridge in terms of presence, location, severity, and type of damage (Rytter 1993, Farrar and Worden 2007). Owing to the size and the number of functionally obsolete or structurally deficient bridges in the United States alone (FHWA 2011), the price tag for instrumenting all of them is still prohibitive despite recent advances in the area of SHM (e.g., Chang 2011, Frangopol et al 2010, Casciati and Giordano 2010).

In this paper, the hypothesis experimentally investigated is that an array of sensors, mounted on moving vehicles that travel across the bridge of interest, can be helpful in identifying structural damage and thus serve as an indicator for more detailed analysis and evaluation that might include in-depth bridge inspection and specific instrumentation. This approach can be referred to as *indirect health monitoring*. The indirect approach can be used in multiple un-instrumented bridges at a low cost and without the need for on-site maintenance. It may be viewed as complementary and, if needed, as a substitute for a global direct health monitoring approach

* The work was performed while the corresponding author was with the Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Corresponding author, E-mail: facerda@udec.cl

if proved to be accurate and effective. Thus, this strategy might help fulfill the need for a practical and cost-effective solution for broad coverage of the bridge population, and help mitigate the costs associated with existing direct SHM practices (Farhey 2005, 2007, Frangopol et al. 2008).

The idea of indirect SHM is not new. It was first formulated by Yang (2004, 2005) who modeled the interaction of a sprung mass traveling on an Euler-Bernoulli beam to extract the beam's fundamental natural frequency. The model was validated experimentally using an instrumented two-wheeled cart attached to a vehicle traveling over a simply supported girder bridge (Lin and Yang 2005). A number of additional studies have been conducted since. In each case, the goal was to detect damage, develop more accurate techniques using limited data for identifying dynamic parameters such as the natural frequencies of the bridge and damping, and to validate these techniques by applying them to increasingly more realistic situations. Some of these efforts are reported in what follows.

Yang and Chang (2009) report results associated with field experiments where the first two natural frequencies of a bridge are extracted from the vehicle response by using empirical mode decomposition. Bu et al. (2006) develop a damage identification scheme based on optimizing a damage parameter vector. The dynamic response of a vehicle moving on top of a simply supported Euler–Bernoulli beam is simulated by a mathematical model. The damage is defined in terms of the reduction of flexural stiffness. The model incorporates noise measurements, road surface roughness, and errors such as underestimating vehicle parameters or bridge flexural stiffness.

Kim and Kawatani (2008) develop a pseudo-static damage detection method that makes use of the coupled vibration of a vehicle-bridge system. The method requires data collected from both the bridge and the vehicle to characterize the damage. A numerical model that includes the roadway roughness effect is used to test the approach. The pseudo-static approach is subsequently validated experimentally for different vehicle speeds and different reduction amounts of girder's moment of inertia (Kim et al. 2010). McGetrick et al. (2009) model a simplified quarter car-bridge interaction to extract the fundamental natural frequency and corresponding damping of the bridge from the spectra of the vehicle accelerations. They found that better accuracy is achieved at lower speeds and smoother road profiles. Moreover, the magnitudes of the acceleration power spectral density's peaks decrease with increasing bridge damping. This decrease is easier to detect with a smoother road profile. This work is validated experimentally by observing the effects of a vehicle moving across a steel girder that included a road surface profile different vehicle mass and speeds (McGetrick et al. 2010).

Isemoto et al. (2010) develop a hypothesis-testing scheme for damage detection based on the vehicle vertical acceleration data. An experimental vehicle-bridge model including roadway roughness, is used. However, only severe damage scenarios were identified. Miyamoto and Yabe (2011) exploit the vibration induced by a public bus for the indirect health monitoring of existing short- and medium-span reinforced/prestressed concrete bridges. The tests show a correlation between the vehicle vertical acceleration and the bridge vibration at midspan. Using a numerical 3D finite element model, they find the distribution of characteristic deflection for a particular driving speed and two severe damage scenarios.

Yin and Tang (2011) used the vertical displacement from the vehicle to identify tension loss and deck damage from numerical models. The relative displacement of a passing vehicle of a bridge with known damaged conditions is used to generate a vector basis. A proper orthogonal decomposition of the relative displacement of a vehicle passing a bridge with an unknown damage condition is optimized with a known basis, and parameters of the unknown damaged bridges are reconstructed. Sirigoringo and Fujino (2012) estimate the fundamental natural frequency of a bridge using the response of a passing instrumented vehicle. They conduct full-scale experiments on a simply supported short span bridge by using a light commercial vehicle. The frequency spectra from the vehicle's dynamic responses reveal the first natural frequency of the bridge. The experimental study considered traveling speed ranging from 10 to 30 m/s. More recently, González et al. (2012) proposed a methodology to identify the damping of a bridge from the vehicle response.

The initial results of an ongoing study are shown here. This work aims at bridging the gap from parameter identification to bridge health monitoring in an indirect fashion. A laboratory bridge model subjected to different vehicle speeds, two different types of reversible damage scenarios, and reversible changes of structural boundary conditions was used. In addition to the indirect measurement of the bridge motion obtained through the vehicle vibration, three sensors are installed directly on the bridge. Contrary to previous studies, the indirect and the direct data are used independently, and a new detection algorithm is introduced. This experimental setup allows for the comparison of the indirect and the direct strategies and evaluation of the effectiveness of the indirect damage detection algorithm.

The authors acknowledge that implementing the indirect monitoring approach assumes as the "ground truth" that the bridge is in an original sound condition. At this early research stage, only data from a specific pristine bridge is collected as the baseline and classified against data collected when the same bridge is modified. The baseline is taken from the response of the vehicle; therefore, there is no need to instrument the bridge.

2. Experimental Setup and Protocol

A laboratory experimental setup is built to collect data from a vehicle, which could be used to detect changes in the condition of the bridge. Using this model, acceleration data from a bridge structure and a vehicle passing over it is collected and later analyzed for characterizing vehicle-bridge interaction patterns. The complete experimental setup consists of mechanical structural components that make up the bridge and vehicle system, a vehicle motion control system, and data acquisition equipment. The different mechanical components resemble a simply supported bridge structure and a four-wheeled vehicle with an independent suspension system at each wheel. The motion control equipment is able to move the vehicle over an acceleration ramp, the simply supported bridge and deceleration ramp. Through this path, the vehicle accelerates until it reaches a target speed before the end of the acceleration ramp, and then maintains a constant speed over the bridge, followed by a deceleration of the vehicle so that it stops at the end of the deceleration ramp. The data acquisition system records accelerations at different locations on the vehicle and the bridge, as well as the position of the vehicle. This experimental setup was inspired by the work of Kim et al. (2010).

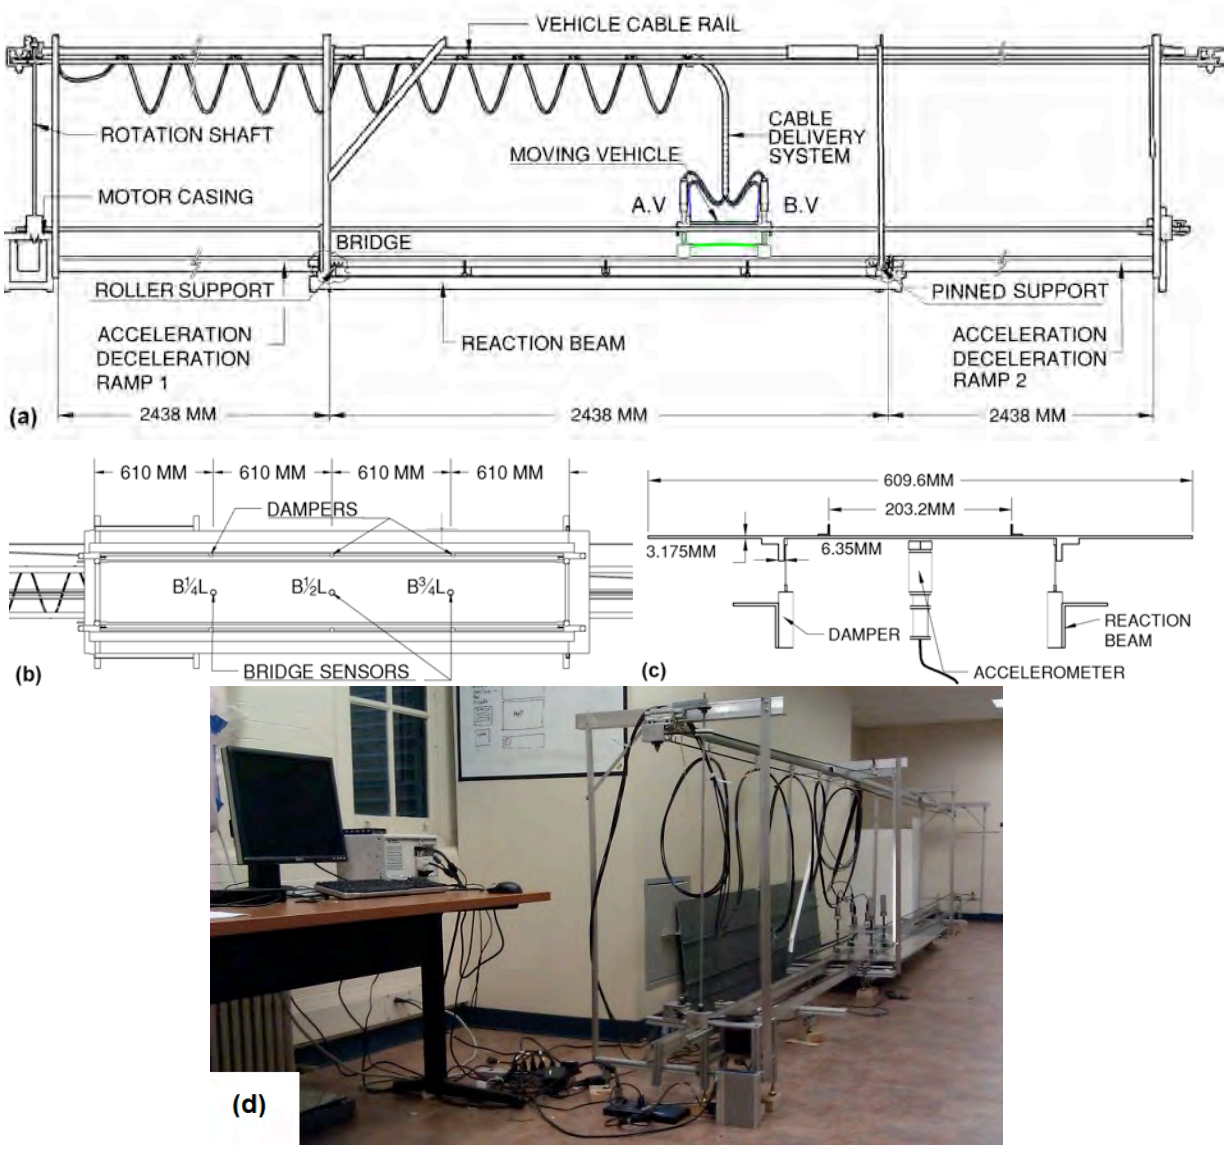

Fig. 1: Experimental setup: (a) Elevation of setup, (b) Bridge bottom view and (c) Bridge cross section (d) Setup overview picture.

An overview of the setup is shown in Fig. 1. The vehicle, approximately in the middle of Fig. 1, is pulled by a belt system. The belt is a neoprene belt, 1/4" width. The travelling path of the vehicle corresponds to the acceleration/deceleration ramps and the bridge as labeled in Fig. 1(a). The simply supported bridge structure is in the middle of the travelling path. Below the simply supported bridge there are two “reaction beams” that are connected to the bridge end supports and used to support added dampers.

The vehicle is instrumented with Vibra-metrics accelerometers (Model 5102) powered by cables supported by a cable delivery system that moves parallel to the vehicle. The cables and the vehicle are propelled by a motor at the leftmost part of the experimental setup.

2.1. Details of the Bridge and Vehicle Models

The bridge is simply supported by a roller support at the left and a pinned support at the right. The vehicle enters the bridge from the left. The whole system is constructed to act as a closed force loop system. The longitudinal forces generated by the motor to move the vehicle are transmitted between the two supports by two connecting beams labeled as "reaction beams" in Fig. 1(a). The bridge structure is instrumented with three 5102 Vibra-metrics accelerometers as in Fig. 1(b). The sensors are equally spaced along the longitudinal direction of the bridge and named accordingly as B1/4L, B1/2L and B3/4L where L=2438 mm. The reaction beams act as a support for localized dampers that connect to the bridge structure as in Fig. 1(c).The bridge deck consists of an aluminum plate, and two angle beams act as the bridge girders. On top of the plate two angle beams serve as rails for the travel path of the vehicle. Detailed dimensions of the bridge section are shown in Fig. 1(c) and a picture of the whole setup is shown in Fig. 1(d).
The bridge has a total weight of 18.3 kg, a fundamental natural frequency of 7.23 Hz, and fraction of critical damping of 3.6 percent. The corresponding aluminum modulus of elasticity and identified section second moment of area are $E=6.9\ E^{10} N/m^2$ and $I \cong 8.15E^{-8}\ m^4$ respectively. These are the properties of the bridge in the pristine condition, later referred to as Scenario 1. The bridge and vehicle models do not resemble a particular full-scale structure. They are treated as a vehicle/bridge system in itself. However, the dimensions follow an approximate scaling factor of S=8 to those of a real structure. Following the scaling laws for elastic vibration analysis, this is proxy for a 2.6 Hz simply supported girder structure. The speed has a scaling factor of $S^{1/2}$.

Fig. 2(a) shows a 3D view of the vehicle constructed for the experimental setup with the main components labeled. The vehicle is instrumented with two accelerometers connected to the suspension shafts in order to record the acceleration at the wheel level and two accelerometers placed on the suspension to acquire data filtered by the suspension system. To keep the symmetry of the vehicle, two calibrated weights are placed on top of the un-sensed wheel shafts. Similarly to the bridge structure, the vehicle is built mainly with aluminum parts.

Two reference points are labeled on the longitudinal direction of the vehicle as points A Vehicle (A.V) and B Vehicle (B.V). A top view of the vehicle is shown in Fig. 2(b). The length and width of the vehicle as well as the labels assigned to each of the four sensors are also indicated. The sensor labels are defined by their position and location to the reference point. Suspension A.V and Suspension B.V are labeled S.A.V and S.B.V respectively, and the two Wheel level sensor locations are labeled W.A.V and W.B.V. A picture of the model vehicle is shown in Fig. 2(c).

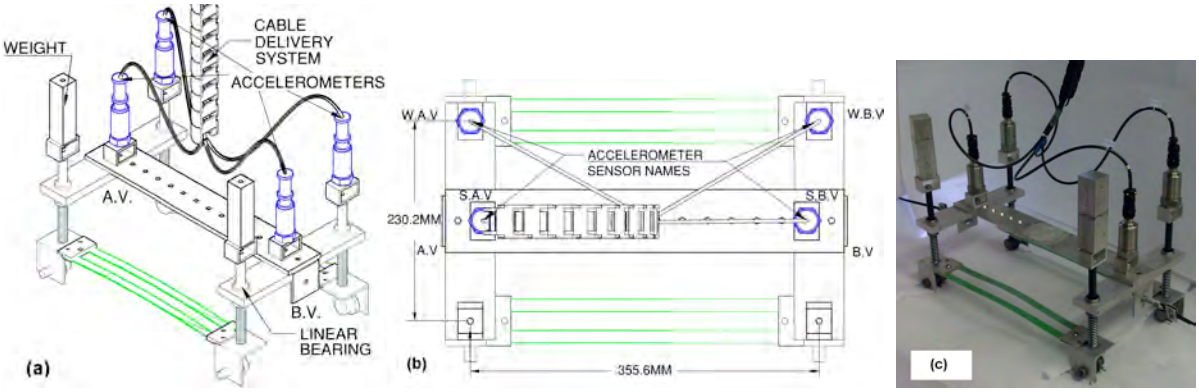

Fig. 2: Experimental vehicle CAD drawings. (a) 3D view. (b) Top view. (c) Vehicle picture.

The frequency of the A.V and B.V axles of the vehicle are examined through free vibration experiments on the suspension. As the vehicle moves along the traveling path, the lengths of the belts at both sides of the vehicle vary. This variation affects the vehicle's dynamics as a whole vehicle-belt system. Free vibration experiments are performed on the vehicle-belt system with the vehicle at different locations. A frequency of about 6 Hz can be identified for both axles by averaging the frequency power spectrum obtained from five free vibration signals. However, depending on the location of the vehicle, other frequencies are also present in the system. The varying frequencies introduced by the belt can be regarded as system noise and make the damage classification task more difficult; good results in such conditions make therefore a stronger point for the classification results shown later in this paper.

Another set of free-vibration experiments with the belt disconnected from the vehicle are run for both axles. The vehicle frequencies are summarized in Table **1**. Damping in both axles is observed to be similar to the critical damping; that is, at the limit between vibration and non-vibration. The vehicle/bridge ratio of the model is about 25%, which is higher than what would be expected in a full-scale scenario.

Table 1: Vehicle properties.

Vehicle weight [kg]	4.8
A.V axle frequency [Hz]	5.0
B.V axle frequency [Hz]	5.5

2.2. Motion Control and Data Acquisition Equipment

A National Instruments® PXI system running in LabView® is assembled to operate the instrumented vehicle and to allow for data acquisition and storage. The system consists of a PXI Chassis (NI PXI 1031) with a motion control card (NI PXI 7342), a motion interface (UMI 7772), a stepper drive (P70360) and a dual shaft stepper motor (NEMA 34). A feedback loop for position is achieved with an encoder. The acceleration data are digitized and stored for post-processing using two digitizers (NI 9234)) in an NI CompactDAQ module. Both digitizers feed the data into the same data file. The effective sampling rate is about 1650 Hz, the minimum for the equipment.

A physical mark is made at the beginning and at the end of the travelling path of the vehicle to produce an acceleration spike that is used to align all the acceleration runs of a specific scenario.

2.3. Protocol

Three different types of "damage" scenarios are designed: (1) variations on the support condition by imposing rotational restraints, 2) increase of damping at different locations, and (3) a mass increase at the midspan. For each type, four levels of severity are devised in order to obtain a total of 12 different damage scenarios. Table 2 shows the twelve conditions of damage. For each case the resonance frequency and the critical damping are reported and compared to the baseline, i.e. Scenario 1. For all cases, the fundamental natural frequency of vibration and the damping coefficient are determined by means of conventional free-vibration experiments.

The rotational restraint mechanism is built into each of the four beam supports of the bridge model. As shown in Fig. 3(a), an aluminum bar is attached to the main girder of the bridge at one extreme and connected to the support at the other extreme to provide vertical restraint. The plate is drilled down to 6.35 mm with 12.7 mm radius to provide only a partial restraint.

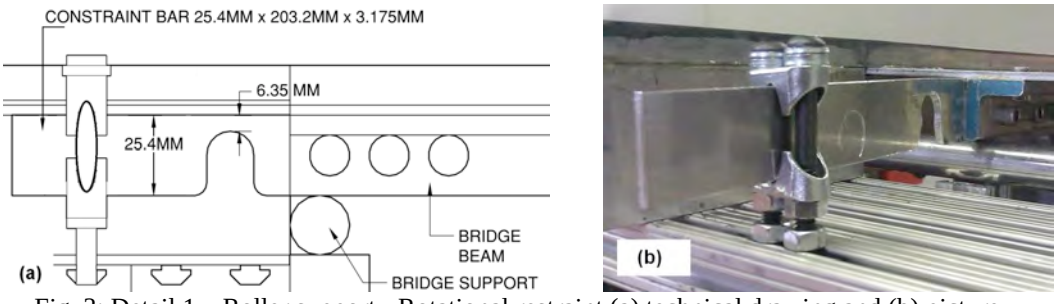

Fig. 3: Detail 1 – Roller support - Rotational restraint (a) technical drawing and (b) picture.

Variation of the rotational restraints simulates the case of rubber bearings becoming stiffer in time or steel corrosion occurring on rocker supports (Chajes et al. 1997, Kim et al. 2009). This condition is a common cause of undesired stress in the structure, and therefore a reduction in the load capacity. In Table 2, the variations of the rotational restraints are described as scenarios SC2 to SC5. In Scenarios 2-5, one, two, three and all four supports are restrained, respectively. As expected, the greater the number of rotational restraints, the higher the fundamental frequency, which provides an indicator of change to the bridge structure.

The variation of localized damping is achieved by adding dampers to the bridge at those locations schematized in Table 2 (Scenarios 6 to 9) according to the scheme presented in Fig. 1(c). A set of AIRPOT adjustable dampers is used. The dampers are calibrated to provide a damping coefficient c≅1.3 N • s/m. In Scenarios 6-9, one, two, four and six dampers, respectively are attached to the bridge structure as depicted in the schematics of Table 2.

Finally, the concentrated mass at the midspan of the structure consisted of weights equal to 50 g, 100 g , 200g and 300 g, and referred to as Scenarios 10 to 13, respectively, in 错误！未找到引用源。. As expected, the presence of the mass, decreased the fundamental frequency of vibration of the structure.

Eight different vehicle speeds, varying from 1 m/s to 2.75 m/s, are considered for each damage scenario. The speed range of the experiments is close to 10-30 km/hr range in a full-case structure.

Table 2: Damage scenarios.

Schematics	SC	f [Hz]	% f shift	$\%\zeta_{crit}$	$\%\ \zeta_{crit}$ shift
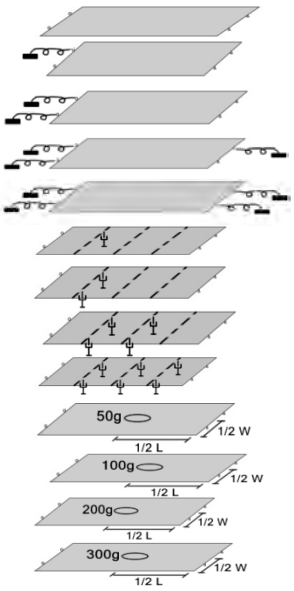	01	7.23	0.00	3.63	0.00
	02	7.46	3.17	6.34	74.9
	03	7.66	6.00	6.45	77.8
	04	8.11	12.2	7.97	120
	05	8.56	18.4	9.37	158
	06	7.24	0.17	8.52	135
	07	7.25	0.28	11.3	212
	08	7.28	0.73	26.4	629
	09	7.30	0.98	31.4	767
	10	7.19	-0.56	4.44	22.5
	11	7.18	-0.66	4.34	19.5
	12	7.14	-1.29	4.07	12.3
	13	7.09	-1.85	4.37	20.5

Because of the design of the system, when applying additional damping the system does not behave as a single degree of freedom damped system. In Scenarios 06-09 from Table 2 with additional damping, there is a coupling effect between the stiffness of the reaction beams and the bridge structure causing a slight increase in the measured fundamental frequency. This effect was verified with an analytical model.

3. Signal Analysis and Classification

The task of distinguishing various bridge conditions is a signal-processing task of classification. The classification process is described first in general and then is explained how it is used in the classification experiments. Assume a real signal *x* of length *N*, i.e., $X \subset R^N$ (see Original signal in Fig. 4). The problem, then, can be formulated as that of designing a map from the signal space of vibrational signals $X \subset R^N$ to a response space of class labels $Y \subset \{1,2,...,C\}$ (in Fig. 4 these are Damaged and Pristine labels). That is, the decision $d : X \to Y$ is the map that associates an input signal with a class label.

A general classification system consists of a feature extractor and a classifier (see Fig. 4). Since the dimensionality of the input space is typically large, the feature extractor is introduced to reduce this dimensionality by setting up a feature space $F \subset R^k$ where $k \le N$ between the input space and the response space. The feature extractor is the map defined as $f : X \to F$ and the classifier as a map $g : F \to Y$.

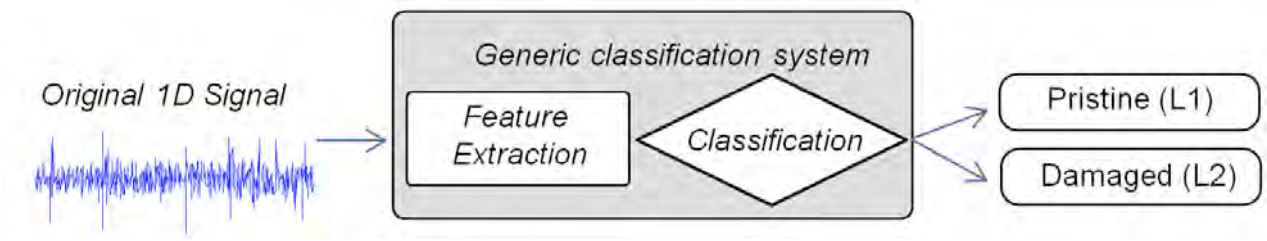

Fig. 4: General classification system.

3.1. Preprocessing

Fig. 5 shows the signal obtained from the vibration of the vehicle, from the time when it starts moving, through its motion across the bridge, until it has come to a stop. The only relevant information for the bridge characterization, however, is that of the vehicle moving across the bridge. A reference start time is chosen as the moment when the rear wheels enter the bridge and reference end time the moment when the front wheels exit the bridge; that portion of the signal is highlighted in Fig. 5(a)-(d) by the two vertical lines. The extracted portions of the signals are then normalized to have zero mean and unit variance.

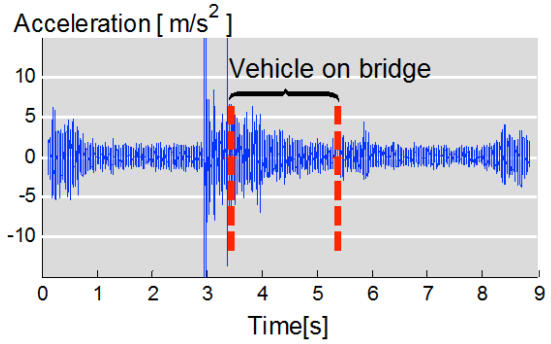

(a) Sensor B ½ L, SC 01, Pristine.

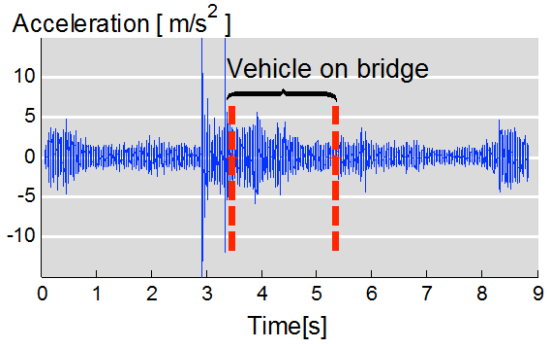

(b) Sensor B ½ L, SC 02, Damaged.

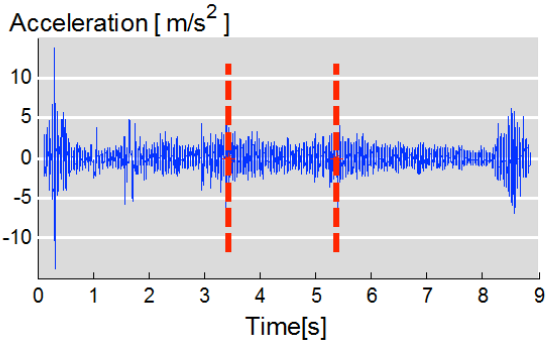

(c) Sensor WBV, SC 01, Pristine.

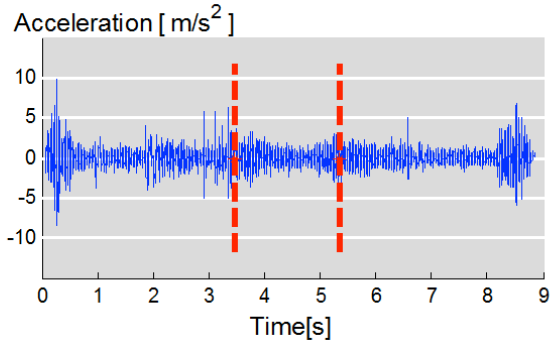

(d) Sensor WBV, SC 02, Damaged.

Fig. 5: Time-domain signals.

3.2. Feature Extraction

A linear structural system can be characterized in the frequency domain by its predominant natural frequencies, their corresponding mode shapes and damping values. The use of frequency spectra characteristics for damage detection is explored in this work.

The fundamental frequency of the bridge, as shown in Table 2 ranges from 7.2 to 8.6 Hz approximately. The n^{th} frequency of a simple supported beam is $n^2 \bullet f_1$, where f_1 is the first mode. The second frequency range goes approximately from 28 to 35 Hz. However, this range is reduced because of the presence of the vehicle mass acting on the bridge. Considering the fundamental natural frequencies of the damage scenarios and the vehicle main bouncing frequency, the analysis is limited arbitrarily to a frequency spectrum of up to 33 Hz allowing for a second bridge mode to influence the response.

For example, looking at Fig. 6 and the spectra (discrete Fourier transform, DFT) of the signals from

two different scenarios, a pristine and a damaged one, it can be seen that, potentially, the responses can be told apart from separate sensors by looking at magnitudes of certain characteristic frequencies. Thus, the authors decided to use frequencies as features, hoping to distinguish among different scenarios. The use of additional features might even improve the results shown later in this study. The task is then to find a set of such features to maximize differentiation between classes.

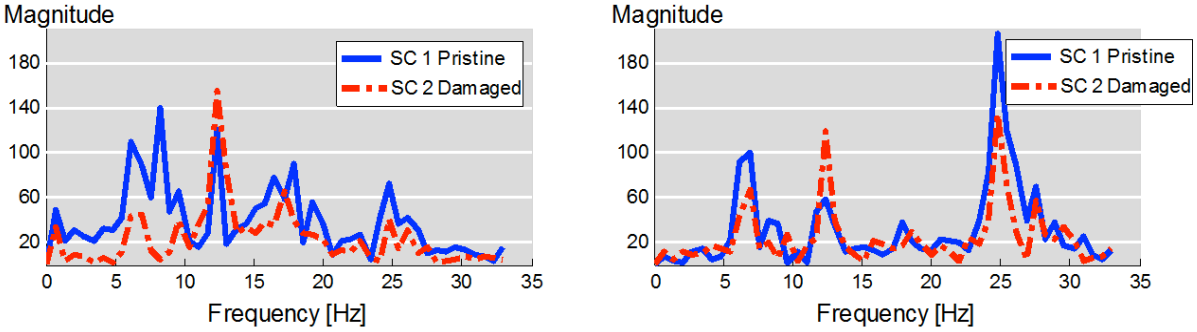

(a) Two scenarios, sensor B ½ L. (b) Two scenarios, sensor W.B.V.

Fig. 6: Discrete Fourier transform of the signal.

Since the acceleration signal contains a large number of spikes and other transient signals, the spectra are noisy with little consistency between runs. To reduce noise and keep non-transient frequencies of interest, a typical approach is to average the spectra across frequency. Averaging the frequency spectra is a well-known technique used in noisy signal processing. For example, in radar signal analysis, a redundant number of antennas capture noisy signals from the same source and average them to increase the signal to noise ratio (Keeler and Passarelli, 1989).

After averaging, the frequency-domain energy distribution for each scenario is calculated. This technique relies on the assumption that each scenario has its unique energy distribution in the frequency domain. Since the goal is to tell classes apart and not individual runs, all the energy distributions from the same class are averaged. The mean energy distribution is used as the representative member of that class.

Fig. 7 shows the distinctive presence of some frequencies that are neither from the bridge nor the vehicle. For example, the most significant peaks are around 13 Hz and 25 Hz in Fig. 7(b). Considering the nature of the problem, the analyzed signals represent the portion of time where the vehicle is traversing the bridge, therefore, these signals do not represent a single structure, rather, a structure made up of the vehicle-bridge coupled effect with vehicle at different locations along the bridge.

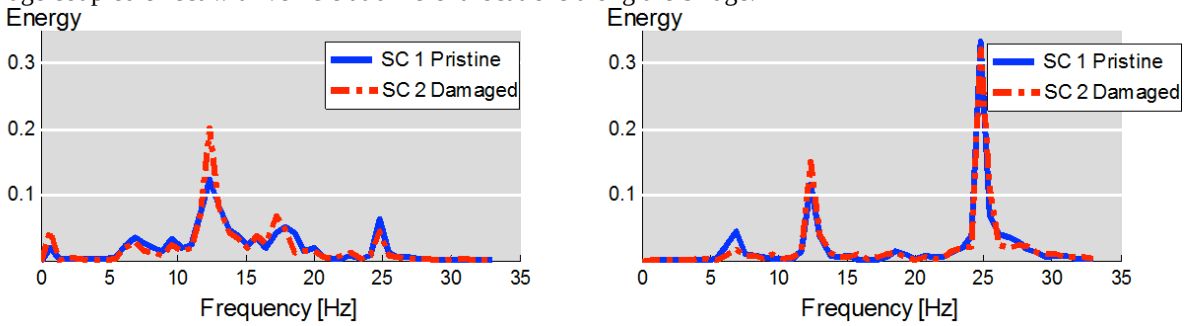

(a) 2 scenarios, sensor B ½ L. (b) 2 scenarios, sensor W.B.V.

Fig. 7: Mean energy distribution (normalized to unit energy).

Let $\{x_i^{(c)}\}_{i=1}^{N_c}$ be a set of signals with N_c samples belonging to Class c. The Fourier energy map is

$$\Gamma_c(j) = \frac{\sum_{i=1}^{N_c} \| w_j^T x_i^{(c)} \|^2}{\sum_{i=1}^{N_c} \| x_i^{(c)} \|^2} \quad (1)$$

where w denotes Fourier basis vector and j denotes the frequency band. To evaluate the power of discrimination of every Fourier basis vector, we need a discriminant measure D to evaluate the power of discrimination. The higher discriminant power is assumed to provide higher discrimination between classes.

For the j th Fourier basis vector, the power of discrimination Δ is denoted by

$$\Delta_j = D(\{\Gamma_c(j)\}_{c=1}^{C}). \quad (2)$$

There exist numerous choices for the discriminant measure; the J-divergence was used in this work (Kullback and Leibler 1951) because it is one of the most famous measurements to quantify the difference or discrepancy of probability density functions (PDFs) in information theory. Moreover, by definition it is additive, which helps in multiclass situations. Let $p = \{p_i\}_{i=1}^n$, $q = \{q_i\}_{i=1}^n$ be two nonnegative sequences with $\sum p_i = \sum q_i = 1$, J-divergence between p and q is:

$$J(p,q) = \sum_{i=1}^{n} p_i \log \frac{p_i}{q_i} + \sum_{i=1}^{n} q_i \log \frac{q_i}{p_i} \tag{3}$$

Fig. 8 shows a graph of the discriminant power between the frequency signals previously depicted in Fig. 7.

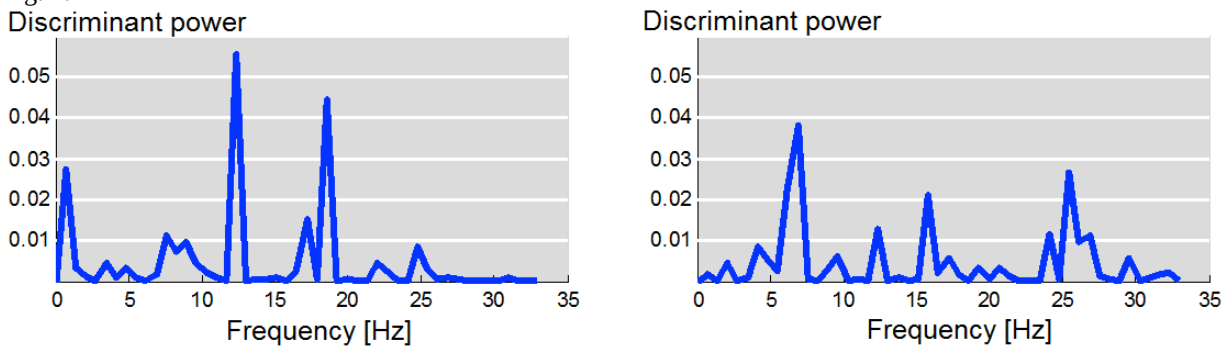

(a) Sensor B ½ L. (b) Sensor W.B.V.

Fig. 8: Discriminant power (normalized to unit discriminant power).

To help understand the feature selection method, a summary of the assumptions and conclusions thus far is presented: 1) To differentiate signals from different scenarios, frequencies are used as features. 2) If the discriminant power is higher, it is easier to discriminate between classes. The discriminant power will thus predict how well a feature will perform during classification. 3) A small number of frequencies provide most of the discriminative power; in other words, the frequency feature set is sparse. Just those frequencies that have large discriminative power are selected (this is called nonlinear approximation); see Fig. 9.

Algorithm (Fourier Discriminant Basis Vectors Selection)
Task: Find $k(\leq n)$ most discriminant Fourier basis vectors
Given a dataset consisting of C classes of signals $\{\{x_i^{(c)}\}_{i=1}^{N_c}\}_{c=1}^{C}$
Step 1: Take the DFT of *x*.
Step 2: Construct Fourier energy map Γ_c for $c = 1,...,C$
Step 3: Determine the power of discrimination $\Delta_j = D(\{\Gamma_c(j)\}_{c=1}^{C})$ for every Fourier basis vector w_j
Step 4: Order Fourier basis vectors by their power of discrimination.
Step 5: Use $k(\leq n)$ most discriminant Fourier basis vectors for constructing classifier

Fig. 9: Summary of feature extraction process.

This selection method performs nonlinear approximation in the Fourier domain and is data adaptive. Different data may give different frequency information and different discriminant powers. Since this method learns from the data and always chooses the frequencies with large discriminant power, it is more robust than traditional linear approximation.

Fig. 10 shows the feature space of the first three features. Blue circles denote the pristine scenario and red asterisks the damaged scenario. Using just three Fourier discriminant basis vectors, it is easy to separate the two classes.

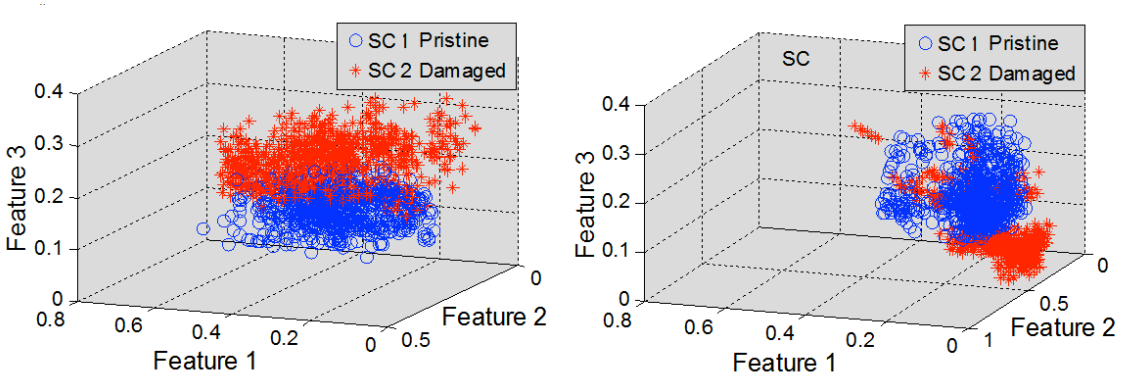

(a) 2 scenarios, sensor B ½ L. (b) 2 scenarios, sensor W.B.V.

Fig. 10: Clustering of two scenarios.

4. Classification

The second part of a classification system is the classifier itself. It takes as input a feature vector and outputs a class label. The classification problem here is called supervised learning, as a labeled training set is given. Many different classifiers are available, such as naïve Bayes, neural networks and many others (Duda et al 2000). In this work, the support vector machine (SVM) classifier is chosen, which is now briefly described.

When looking for the best boundary to separate classes, two things are desired: 1) the boundary should give high classification accuracy; 2) avoid overfitting. To satisfy these two requirements, SVM maximizes the margin, that is, the distance between a decision boundary and a data point, and expresses it as a function of the weight vector and bias of the separating hyperplane, which is used to separate the space in two.

There are 30 samples in each class. If 3 out of 30 samples are randomly chosen and averaged, then there will be 4060 potential different choices (30 choose 3 binomial coefficients). As the study dataset, 1000 out of 4060 samples are randomly chosen for each scenario. A 20-fold cross validation is performed. Each time, 2000 data samples, consisting of 1900 training samples and 100 testing samples, are used to create and test the SVM-based classifier. We then used our Fourier discriminant basis search algorithm and selected the top 5 frequencies that provide the largest discriminant power as features. Then kernel SVM is used as the classifier.

5. Results and Discussion

The results of the classification experiments are presented and discussed in terms of the classification accuracy, which is defined as the number of test samples correctly classified divided by the total number of test samples.

In the two classes defined, pristine and damaged, scenarios 2-13 belonged to the latter class. The data collected from all seven accelerometers are used.

Fig. 11 shows the variation of the average classification accuracy for different variables. Fig. 11(a) shows the average across the different severities, speeds and sensor locations for each damage type. The three bridge sensors, B1/4L, B1/2L and B3/4L, are averaged and referred to as “Bridge”, the two sensors at the wheel level, (W.A.V and W.B.V) are averaged and referred to as “Wheel” and the two sensors at the suspension level, (S.A.V and S.B.V) are averaged and referred to as “Suspension”. The standard deviation across the averaged variables is shown at the top of each bar. An average classification accuracy for all the sensors for each damage type is depicted with a black line and corresponding percentage. The baseline in Fig. 11(a) and (b) is 50%, which is the expected probability of randomly choosing between two labels (pristine or damaged). Classification accuracy values of over 90% are obtained despite the subtle changes introduced in the bridge structure. The amount of change inflicted is deliberately small to test the detection capability of the combined indirect approach using the signal processing techniques described in Section 错误！未找到引用源。. The signals from the sensors located at the wheel level are classified consistently across the different damage types, and more accurately than those from the sensors located on the bridge or on the vehicle at the suspension level.

The classification results in Fig. 11(b) show how the average classification accuracy for all damage scenarios varies for different vehicle speeds. Similarly to Fig. 11(a), each bar represents the mean accuracy classification across the different damage scenarios. At the top of each bar, the corresponding standard deviation is shown. Looking at Fig. 11(b), one can see that there is a jump between the first four speeds, between 1 and 1.75 m/s and the four higher speeds from 2 to 2.75 m/s. The average across the two groups of speeds is shown

by a black line and corresponding percentage. There is about a 7% difference in classification accuracy between these two speed ranges. This classification accuracy difference is consistent for the average classification accuracy of the sensors at the different locations (i.e., Bridge, Wheel and Suspension).

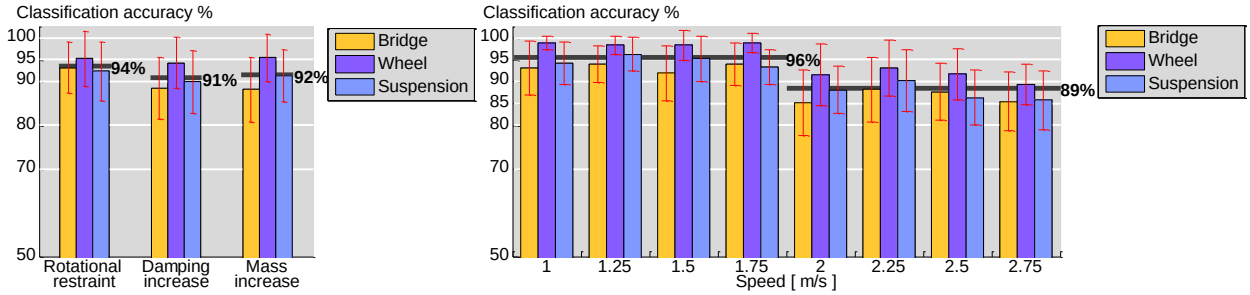

(a) Average by damage types. (b) Average by speeds.

Fig. 11: Classification accuracy results.

Fig. 12 illustrates the sensitivity of the classification method to different levels of severity of the different damage types inflicted in the damage scenarios. Fig. 12(a), (b) and (c) show the average classification for the different damage severity levels for the rotational restraint damage type for different vehicle speeds. For all of these graphs, the thickness of the line depicts the level of damage. The thinnest line indicates the least amount of damage inflicted (e.g., only one of four rotational restraints invoked in SC2) and the thickest linethe maximum amount of damage inflicted (e.g., all four rotational restraints invoked in SC5).

Fig. 12(a), (b) and (c) show the average classification accuracy for each rotational restraint damage severity level for the signals from all the sensors on the bridge, all the sensors on the suspension, and all the sensors on the wheel, respectively. Fig. 12(d), (e), and (f) break the results down for each sensor and show the average classification accuracy for each rotational restraint damage severity level for each signal from the three sensors on the bridge, B1/4L, B1/2L and B3/4L; from the two sensors on the suspension, S.A.V and S.B.V; and from the two sensors on the wheel, W.A.V and W.B.V, respectively.

Fig. 12 demonstrates that the classification accuracy for SC2 is lower than for the other rotational restraint scenarios (SC3, SC4 and SC5). For SC2, there is a variation in the classification accuracy with respect to speed. However, more severe rotational restraint scenarios seem to be less dependent on speed, with high classification accuracy for low speeds and a slight parabolic decrease for higher speeds. No significant difference in the classification accuracy is apparent in Fig. 12(a), (b) and (c) regarding the sensor location. This shows that, in terms of classification accuracy, the signal processing approach performs well with sensor data from the vehicle (sensor or wheel) as with sensor data directly measured on the bridge. The wheel sensors perform slightly better than the suspension or bridge sensors. In other words, the results indicate that in this particular set of experiments, and for the signal processing scheme used, the indirect approach has a classification accuracy that is as good as that of the direct approach. Fig. 12(d), (e), and (f) show that the lowest classification accuracies are those of the least severe damage scenarios. The most severe scenarios are all grouped with high classification accuracies.

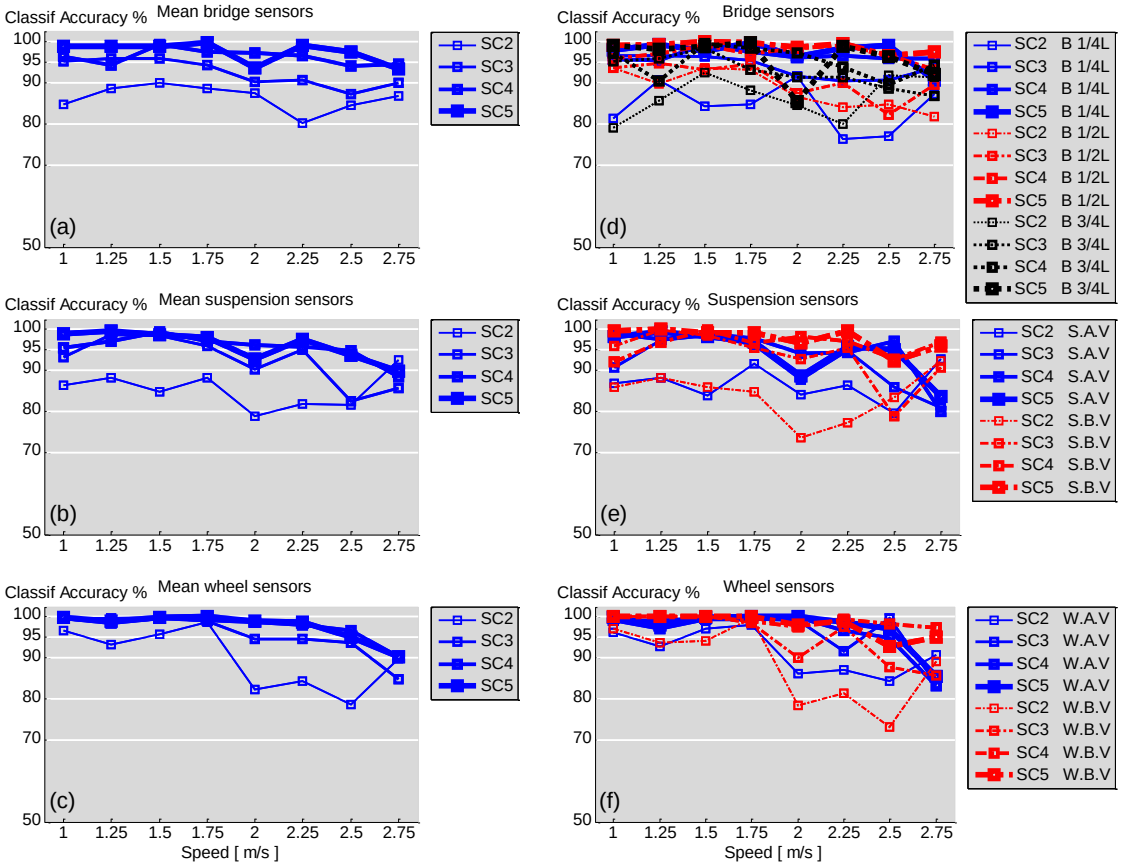

Fig. 12: Classification results for rotational restraint scenarios.

Fig. 13 and Fig. 14 are similar in nature to Fig. 12, but display the classification accuracies for the two other damage types explored: the scenarios with increasing amounts of damping and with increasing amounts of mass. Fig. 13(a), (b) and (c) show Scenario 6, the scenario with the single damper, as the one with the least classification accuracy across all speeds. In terms of vehicle speed, a decrease in the classification accuracy appears to occur when the vehicle speed is 2 m/s as shown in Fig. 13(a). The same observations from Fig. 12 apply to Fig. 13 as well. There is a slight decrease of the classification accuracy at higher speeds, and the classification accuracy seems to be independent of the sensor location; that is, there is little difference in the classification capability between the direct and the indirect approaches.

Fig. 14 shows the classification results for the scenarios with a mass increase at the midspan. Even though the inflicted change in the bridge structure is quite subtle, the classification accuracy is high, especially at lower speeds. The same observations made for Fig. 12 and Fig. 13 are valid for Fig. 14 as well, that is, classification accuracies exhibit small variations with respect to the vehicle speed and the sensor locations.

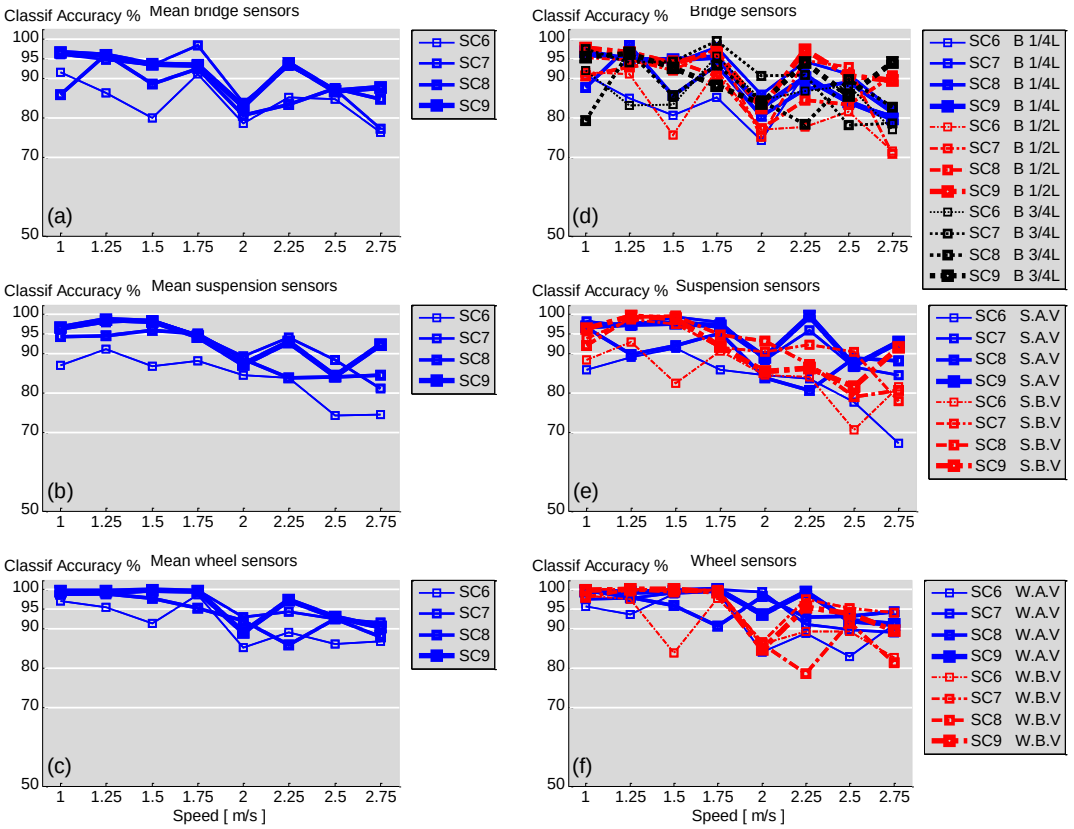

Fig. 13: Classification results for damping increase scenarios.

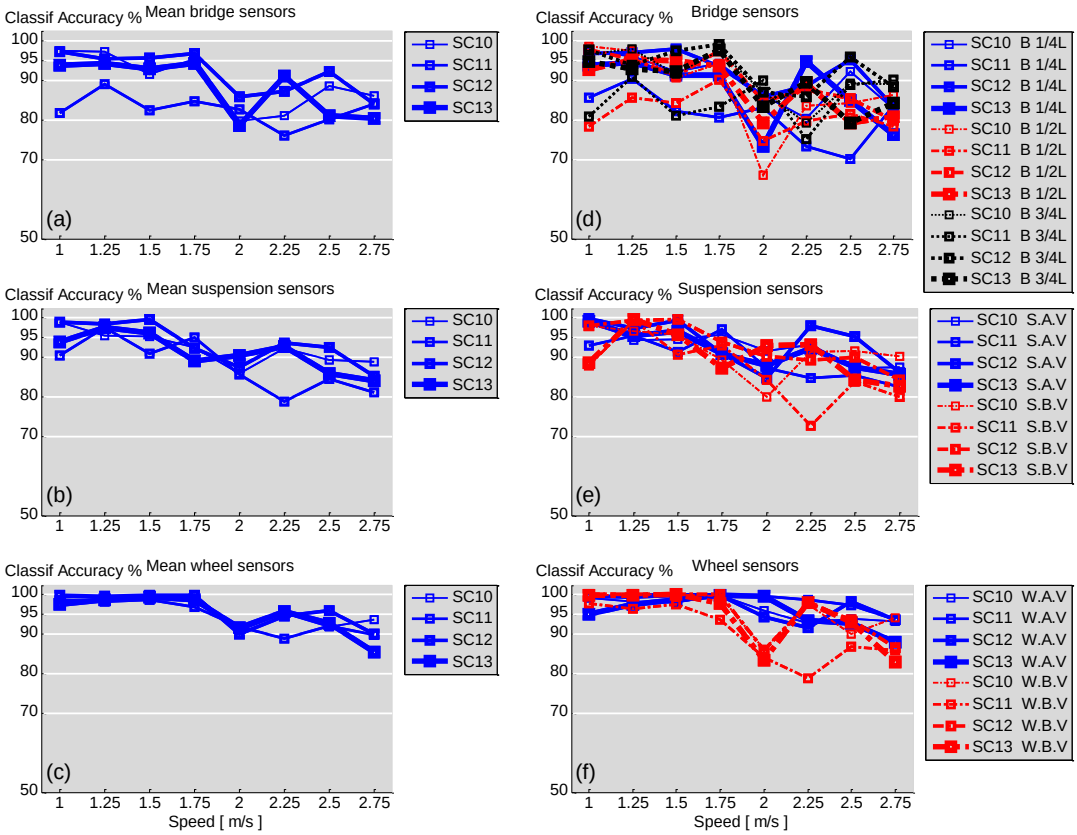

Fig. 14: Classification results for concentrated mass scenarios.

6. Conclusions

The damage detection capability of an indirect bridge monitoring approach is studied by means of data collected from moving vehicles and applying signal processing techniques to detect damage of a bridge. The study uses data from a laboratory vehicle-bridge physical model. The indirect monitoring approach requires no pre-condition on the bridge and is pursued by the authors as an economical and effective bridge SHM approach for a large bridge stock. However, in a full-scale deployment, some bridge parameters such as their geometric configuration and location will be necessary for data pre-processing.

To perform numerous test repetitions, a fully automated vehicle-bridge model needed to be built. In this study, the amount of experimental data samples is significantly greater than that in previous experiments. Each scenario is run 30 times at eight different velocities. Three different reversible damage proxy types were built into the experimental setting, and each damage type had four different severity scenarios.

The synchronized acceleration data from the bridge and the vehicle, and the vehicle position data, allowed for the comparison of the direct and indirect approaches in terms of the accuracy with which each could classify the existence of damage for different extents of damage.

A feature extraction technique based on averaging the power spectrum from a set of data is used to achieve high noise reduction. Then, features extracted from the Fourier domain are automatically chosen from the denoised data samples based on their significance and classified using an SVM classifier. The use of other features and signal processing techniques will be pursued in future research. However, using frequencies and the specific signal processing scheme presented, allowed comparing the direct and indirect monitoring approaches damage detection capability and trends regarding different types of damage, severities of damage and vehicle speeds.

High classification accuracy is achieved across three distinct types of changes inflicted into the bridge structure: 1) a change in the support conditions obtained by introducing rotational restraints at the supports; 2) an increase in the damping of the bridge structure; and 3) a localized mass increase at the midspan of the bridge.

The severity of the changes inflicted in the bridge structure is consistent with higher classification accuracy. For example, SC3, SC4 and SC5 imposed more significant changes into the bridge structure than SC2, and consistently higher classification accuracy is obtained. Nonetheless, the classification accuracy achieved for the subtle change inflicted on SC2 is on average above 85%.

The detection of the various changes in the bridge structure is quite insensitive to the vehicle speed. This effect can be important for practical applications where vehicle speeds cannot be readily controlled. However, a small jump is observed between the lower and higher speeds, where the classification accuracy decreases by about 7 percent at the higher speeds. At this early research stage, this sudden decrease can only be attributed to the non-linear nature of the dynamic interaction problem.

Independent of the sensor location, high classification accuracy is achieved across all the sensors. The indirect and direct approaches seem to be equally effective for damage detection when applying the proposed signal processing techniques. Of the two sensor locations considered in the indirect approach, the wheel level and the suspension level, the sensors at the wheel level performed better than the sensors at the suspension level.

Given the simplicity of the model considered, the results presented are strictly applicable only to the particular experimental setup and cannot be generalized for full-scale structures at this time. The authors have not yet tested whether the indirect approach can be generalized to different bridge structural configurations and/or with a specific range of dynamic properties. The effect of the vehicle/bridge ratio is not addressed in this work. This parameter will be subject of future investigations. On the other hand, a high degree of consistency is observed in the classification accuracies across the very different types and severity of damage and for different vehicle speeds. This gives us hope that our approach might be applicable to more general systems.

In reality, data collected about actual damage conditions from the actual structure being monitored will not always be available for use in training. In this case, different detection schemes would have to be used, such as outlier detection, which would help to trigger more in-depth inspection. In the long run numerical and experimental models that proxy real case scenarios would allow learning the evolution of damage sensitive features obtained either in the direct or indirect fashion from different damaged scenarios. The early detection of those features will allow making more accurate diagnostics and prognosis about the overall structural health.

Clearly, further research is needed to validate the robustness of these results for more realistic systems and conditions, including different roadway roughness profiles, atmospheric conditions and other bridge interaction variables such as different vehicle/bridge mass ratios, the effect of ongoing traffic and torsional effects on the bridge by non-symmetric loading from the vehicle path.

7. Acknowledgements

We gratefully acknowledge the support from the National Science Foundation (Award No. CMMI1130616); the Traffic 21 initiative and the T-SET University Transportation Center sponsored by the US Department of Transportation under Grant No. DTRT12-G-UTC11 at Carnegie Mellon University; and the Fulbright-MECESUP Fellowship awarded to the corresponding author. We thank Eduard Romero, Ranny Zhao, Zhe Zhuang and Timothy Pianka, who participated in the construction of the experimental setup and running the experiments. We also thank Michael McCann for his input in the signal processing section.

References

Bu, J. Q., Law, S. S., and X. Q. Zhu. 2006. "Innovative Bridge Condition Assessment from Dynamic Response of a Passing Vehicle." *J. Eng. Mech* 132 (12) (December): 1372–1379. doi:10.1061/(ASCE)0733-9399(2006)132:12(1372).

Carden, E. P., and Fanning, P. (2004). "Vibration Based Condition Monitoring: A Review." *Struct. Health Monit.* 3 (4) (December 1): 355–377. doi:10.1177/1475921704047500.

Casciati, F., and Giordano, M. (2010). *Structural Health Monitoring 2010: Proceedings of the Fifth European Workshop on Structural Health Monitoring Held at Sorrento, Naples, Itlay, June 28-july 4, 2010*. DEStech Publications, Inc.

Chang, Fu-Kuo. 2011. *Structural Health Monitoring 2011: Condition-Based Maintenance and Intelligent Structures*. Destech Publications.

Chajes, J., Mertz, D. and Commander, B. (1997). "Experimental Load Rating of a Posted Bridge." *J. of Bridge Eng.*2, no. 1: 1–10.

Doebling, S.W., Farrar, Ch.R., and Prime, M.B. (1998). "A Summary Review of Vibration-based Damage Identification Methods." *The Shock and Vibration Digest* 30: 91–105. doi:10.1.1.57.9721.

Duda, R. O., Hart, P.E. and Stork, D. G. (2000). *Pattern Classification (2nd Edition)*. Wiley-Interscience.

Farhey, D. N. (2005). "Bridge Instrumentation and Monitoring for Structural Diagnostics." *Struct Health Monit* 4 (4) (December 1): 301–318. doi:10.1177/1475921705057966.

Farhey, D. N (2007). "Quantitative Assessment and Forecast for Structurally Deficient Bridge Diagnostics." *Struct Health Monit* 6 (1) (March 1): 39–48. doi:10.1177/1475921707072061.

Farrar, C.R., and Worden, K. 2007. "An Introduction to Structural Health Monitoring." *Phil. Trans. R. Soc.* A 365 (1851) (February 15): 303–315. doi:10.1098/rsta.2006.1928.

FHWA (2011). "Deficient Bridges by State and Highway System." http://www.fhwa.dot.gov/bridge/deficient.cfm.

Frangopol, D. M., Strauss, A, and Kim, S. 2008. "Bridge Reliability Assessment Based on Monitoring." *J. Bridge Eng.* 13 (3) (May): 258–270. doi:10.1061/(ASCE)1084-0702(2008)13:3(258).

Frangopol, D. M., Sause, R., and Kusko, Ch. eds. 2010. *Bridge Maintenance, Safety and Management - IABMAS'10: Proceedings of the Fifth International IABMAS Conference, Philadelphia, USA, 11-15 July 2010*. 1st ed. CRC Press.

Isemoto, R., Kim,C.W., and Sugiura, K. 2010. "Abnormal Diagnosis of Bridges Using Traffic-induced Vibration Measurements." *The Twenty-Third KKCNN Symposium on Civil Engineering,*Taipei, Taiwan.

González, A., E.J. OBrien, and McGetrick, P.J.. "Identification of Damping in a Bridge Using a Moving Instrumented Vehicle." (2012) *J. Sound Vib.* 331, no. 18 (August 27): 4115–4131.

Keeler, R.J., Passarelli, R.E. (1989). "Signal processing for atmospheric radars". NCAR/TN-331+STR NCAR TECHNICAL NOTE

Kim, C.W., and Kawatani, M. (2008). "Pseudo-static Approach for Damage Identification of Bridges Based on Coupling Vibration with a Moving Vehicle." *Struct. Infrastruct E.* 4 (5): 371–379.

Kullback, S, and Leibler, R.A. (1951). "On Information and Sufficiency." *Ann. Math. Statist* 22: 79–86.

Kim, C. W., Kawatani, M, and Fujimoto, T. (2010). "Identifying Bending Stiffness Change of a Beam Under a Moving Vehicle." In *Bridge Maintenance, Safety and Management - IABMAS'10,* 180–180. Bridge Maintenance, Safety and Management. CRC Press. http://dx.doi.org/10.1201/b10430-108.

Kim, Y. J., Tanovic, R., and Wight, R. G.. (2009) "Recent Advances in Performance Evaluation and Flexural Response of Existing Bridges." *J. Perform. Constr. Facil.* 23, no. 3 (June 2009): 190–200. doi: 10.1061/(ASCE)CF.1943-5509.0000007

Lin, C.W., and Yang, Y.B. (2005). "Use of a Passing Vehicle to Scan the Fundamental Bridge Frequencies: An Experimental Verification." *Engineering Structures* 27 (13) (November): 1865–1878. doi:10.1016/j.engstruct.2005.06.016.

McGetrick, P.J., González, A., and O'Brien, E.J. (2009). "Theoretical Investigation of the Use of a Moving Vehicle to Identify Bridge Dynamic Parameters." *Insight: Non-Destructive Testing and Condition Monitoring* 51 (8): 433–438.

McGetrick, P.J., Kim, C.W., and Obrien, E.J. (2010). "Experimental Investigation of the Detection of Bridge Dynamic Parameters Using a Moving Vehicle." *The Twenty-Third KKCNN Symposium on Civil Engineering*, Taipei, Taiwan.

Miyamoto, A., and Yabe, A. (2011). "Bridge Condition Assessment Based on Vibration Responses of Passenger Vehicle." *J. Phys. Conf. Ser.* 305 (July 19): 012103. doi:10.1088/1742-6596/305/1/012103.

Rytter, A. (1993) Vibration based inspection of civil engineering structures. Ph.D. Dissertation, Department of Building Technology and Structural Engineering, Aalborg University, Denmark.

Siringoringo, D.M., and Y. Fujino. (2012). "Estimating Bridge Fundamental Frequency from Vibration Response of Instrumented Passing Vehicle: Analytical and Experimental Study." *Adv. Struct Eng.* 15 (3): 417–433.

Yang, Y. B., Lin, C. W. and Yau, J. D. (2004). "Extracting Bridge Frequencies from the Dynamic Response of a Passing Vehicle." *J. Sound Vib.* 272 (3-5) (May 6): 471–493. doi:10.1016/S0022-460X(03)00378-X.

Yang, Y.B., and Chang, K.C. (2009). "Extraction of Bridge Frequencies from the Dynamic Response of a Passing Vehicle Enhanced by the EMD Technique." *J. Sound Vib.* 322 (4-5): 718–739.

Yang, Y.B., and Lin, C.W. (2005). "Vehicle-bridge Interaction Dynamics and Potential Applications." *J. Sound Vib* 284 (1-2) (June 7): 205–226. doi:10.1016/j.jsv.2004.06.032.

Yin, S-H, and Tang, C-Y. (2011). "Identifying Cable Tension Loss and Deck Damage in a Cable-Stayed Bridge Using a Moving Vehicle." *J. Vib. Acoust.* 133 (2): 021007. doi:10.1115/1.4002128.

Semi-Supervised Multiresolution Classification Using Adaptive Graph Filtering with Application to Indirect Bridge Structural Health Monitoring

Siheng Chen, *Student Member, IEEE*, Fernando Cerda, Piervincenzo Rizzo, Jacobo Bielak, James H. Garrett and Jelena Kovačević, *Fellow, IEEE*

***Abstract*—We present a multiresolution classification framework with semi-supervised learning on graphs with application to the indirect bridge structural health monitoring. Classification in real-world applications faces two main challenges: reliable features can be hard to extract and few labeled signals are available for training. We propose a novel classification framework to address these problems: we use a multiresolution framework to reliably extract features and semi-supervised learning to train both labeled and unlabeled signals. We further propose an adaptive graph filter for semi-supervised classification that allows for classifying unlabeled as well as unseen signals and for correcting mislabeled signals. We validate the proposed framework on indirect bridge structural health monitoring and show that it performs significantly better than previous approaches.**

***Index Terms*—multiresolution classification, semi-supervised learning, discrete signal processing on graphs, adaptive graph filter, indirect bridge structural health monitoring**

I. Introduction

CLASSIFICATION is a signal processing task whose goal is to design a map that associates each input signal with a predefined class label. It is widely used in a number of real-world applications, such as geophysical waveform classification [2], radar signal classification [3], structural health monitoring [4], computer-aided diagnosis of medical images and classification of biological images [5]–[7]. A generic classification system consists of a feature extractor and a classifier: a feature extractor reduces the dimensionality of the problem, while a classifier labels the features. In many real-world problems, however, reliable features can be hard to extract; for example, different lighting conditions can derail robust face recognition [8]. Moreover, few labeled signals could be available for training; for example, in bridge structural health monitoring, one of the practical problems is that it is financially impractical to label a large dataset, and thus, few labeled signals are available.

S. Chen is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 15213 USA. Email: sihengc@andrew.cmu.edu. F. Cerda is with Universidad de Concepción, Concepción, Chile. Email: fernando.cerda@gmail.com. P. Rizzo is with the Department of Civil and Environmental Engineering, University of Pittsburgh, Pittsburgh, PA, 15261, USA. Email: pir3@pitt.edu. J. Bielak and J. H. Garrett are with the Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA, 15213 USA. Emails: {jbielak, garrett}@cmu.edu}. J. Kovačević is with the Departments of Biomedical Engineering and Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA. Email: jelenak@cmu.edu.

The authors gratefully acknowledge support from the NSF through awards 1130616 and 1017278, as well as CMU Carnegie Institute of Technology Infrastructure Award. Parts of this work were presented at ICASSP 2013 [1].

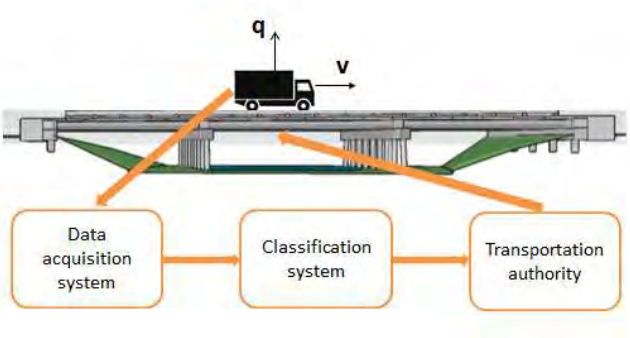

Fig. 1: Indirect bridge structural health monitoring system. Acceleration signals are collected from a moving vehicle and sent to a classification system, which identifies the bridge status and reports it to a transportation authority.

To address these challenges, we propose a novel classification framework that takes advantage of supervised multiresolution classification [5], which extracts hidden features in localized time-frequency regions (subbands), and semi-supervised learning [9], which uses both labeled and unlabeled signals for classification. We follow this by a novel weighting algorithm that combines information from all the subbands of all the signals to make a global decision in a semi-supervised fashion. In the new framework, (1) each localized subband contributes to the classification by its discriminative power; and (2) both labeled and unlabeled signals provide information.

We further propose an adaptive graph filter for semi-supervised classification based on discrete signal processing on graphs [10]; this adaptive graph filter extends the applications of signal processing on graphs to classification [11]. A graph signal is built by defining each node to be a signal in a given dataset and each edge to be the similarity between each pair of signals. The adaptive graph filter classifies signals by filtering the graph signal and producing labels. This adaptive graph filter allows for classifying unlabeled as well as unseen signals (together with regression) and for correcting mislabeled signals. We further establish the connection to the theory of diffusion maps [12] as well as that of diffusion wavelets [13].

We validate the proposed framework on a real-world classification problem, indirect bridge structural health monitoring. Assessing and monitoring bridge health has been an intense

Generic classification		
$\mathcal{X} = \{x^{(i)}\}$	input dataset	$i = 1, \ldots, N$
$\mathcal{L} = \{x^{(i)}\}$	labeled dataset	$i = 1, \ldots, L$
$\mathcal{U} = \{x^{(i)}\}$	unlabeled dataset	$i = L+1, \ldots, N$
$\mathcal{Y} = \{y^{(i)}\}$	ground-truth labels for $\mathcal{L}$	$i = 1, \ldots, L$
$\hat{\mathcal{Y}} = \{\hat{y}^{(i)}\}$	estimated labels for $\mathcal{U}$	$i = L+1, \ldots, N$
F	feature extraction function	
$f^{(i)}$	feature vector	$i = 1, \ldots, N$
$q^{(i)}$	ground-truth vector	$i = 1, \ldots, L$
Q	ground-truth matrix	$L \times C$
$\hat{q}^{(i)}$	confidence vector	$i = L+1, \ldots, N$

TABLE I: Parameters used in a generic classification system.

area of interest for some time, especially in the aftermath of several bridge collapses. One of the practical problems in bridge structural health monitoring is the scarcity of labeled signals; these are collected based on visual inspection, an expensive process. Moreover, visual inspection is not always accurate, producing unreliable labels. Sensor-based structural health monitoring systems have been proposed to automate and improve on the visual inspection process. One approach is to install strain gauges, accelerometers, or other sensors directly on the bridge. The drawback is that such sensors still require a sophisticated and expensive electronic infrastructure with installation, maintenance and power support. Recently, indirect approaches have been proposed [14]–[16], based on using moving vehicles to collect signals from accelerometers inside the vehicles, a more efficient solution that is expected to be economically feasible by design (see Figure 1).

Because in indirect approaches less data is collected and the data is noisier (as it is farther from the source), data analysis plays a crucial role. Moreover, although it is easy to get a large number of signals, it is expensive to label them because the process involves physically inspecting the bridge and determining its health; thus, very few signals are actually labeled. We show that our proposed framework and algorithm perform remarkably well for the case of a lab-scale bridge-vehicle dynamic system.

Previous Work. Multiresolution classification was originally proposed for bioimaging applications with excellent performance on classifying images of protein subcellular locations [5], developmental stages of Drosophila embryos [6], germ layer components in teratomas [7], and even fingerprint recognition [17]. Previous work on semi-supervised learning includes generative mixture models with expectation maximization, co-training, transductive support vector machine and graph-based approaches [9], each of which makes specific assumptions on how to use unlabeled signals to help classification. Signal processing on graphs has been proposed as a framework to build tools to analyze structured signals and is a rather recent development [10], [11]. Indirect bridge structural health monitoring determines the state of the bridge by using advanced signal processing techniques to analyze vibrational signals collected from the dynamic responses of vehicles traversing a bridge [15], [16], [18], [19].

Contributions. Our contributions are as follows: We propose a novel

- classification framework that combines multiresolution classification with semi-supervised learning;
- adaptive graph filter for semi-supervised classification that allows for classifying unlabeled as well as unseen signals and for correcting mislabeled signals; and
- a tentative solution to indirect bridge structural health monitoring.

Algorithm 1 Generic classification

Input	$\mathcal{X}$	input dataset
	$\mathcal{Y}$	ground-truth labels for $\mathcal{L}$
Output	$\hat{\mathcal{Y}}$	estimated labels for $\mathcal{U}$
Function	$\mathbf{C}(\mathcal{X})$	
	$f^{(i)} = \mathrm{F}(x^{(i)})$	feature extraction
	$\hat{q}^{(i)} = \mathrm{C}(f^{(i)}, Q)$	classification
	$\hat{y}^{(i)} = \arg\max_c \hat{q}_c^{(i)}$	
	return $\hat{\mathcal{Y}}$	

Outline of the Paper. Section II states the problem and briefly reviews multiresolution classification, semi-supervised learning and signal processing on graphs; Section III describes our proposed framework for semi-supervised multiresolution classification, while Section IV describes our proposed adaptive graph filter for semi-supervised classification. The algorithms are validated in Section V on acceleration signals collected from a lab-scale bridge-vehicle dynamic system. Section VI concludes with discussion and pointers to future directions.

II. Background and Problem Formulation

In this section, we cover the background material necessary for the rest of the paper. We start with the classification problem and then a supervised classification framework, multiresolution classification system. Next, we introduce signal processing on graphs, which lays a foundation for our proposed semi-supervised classifier. Finally, we overview semi-supervised learning, which we will use in Section III.

A. Classification

The goal of classification is to label signals as belonging to one of a number of given classes [20]. Let $\mathcal{X} = \{x^{(i)} \in \mathbb{R}^D\}_{i=1}^N$ be the given dataset with $N = L+U$ signals, the first L belonging to the *labeled dataset* $\mathcal{L} = \{x^{(i)} \in \mathcal{X}\}_{i=1}^L$, and the last U belonging to the *unlabeled dataset* $\mathcal{U} = \{x^{(i)} \in \mathcal{X}\}_{i=L+1}^N$. Thus, the inputs to the classifier are the dataset $\mathcal{X}$ and the *ground-truth labels* $\mathcal{Y} = \{y^{(i)} \in \{1, 2, \ldots, C\}\}_{i=1}^L$ for the labeled dataset $\mathcal{L}$, while the outputs are the *estimated labels* $\hat{\mathcal{Y}} = \{\hat{y}^{(i)} \in \{1, 2, \ldots, C\}\}_{i=L+1}^N$ for the unlabeled dataset $\mathcal{U}$ (see Table I and Algorithm 1).

We formulate the problem as designing a map that associates an input signal to a class label with a certain probability. That is, we view the label as a posterior probability vector $\hat{q} \in \mathbb{R}^C$, where the cth component, $\hat{q}_c$, is the probability that a signal belongs to the cth class. Since $\hat{q}$ expresses a confidence of an assigned label, we name it a *confidence vector*. The confidence vector for a labeled signal $i = 1, 2, \ldots, L$, is the *ground-truth vector*, $q \in \mathbb{R}^C$, with 0s everywhere except 1 in

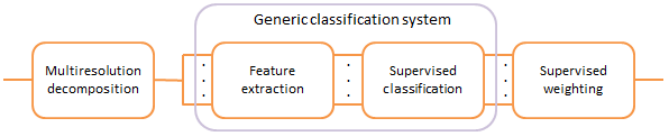

Fig. 2: Supervised multiresolution classification decomposes a signal into localized space-frequency subbands using a given multiresolution transform, followed by feature extraction and supervised classification in each subband, yielding a local classification decision. A supervised weighting algorithm combines all local decisions into a global decision.

position c indicating membership in class c. The *ground-truth matrix* $Q = \left[q^{(1)} \quad q^{(2)} \quad \dots \quad q^{(L)}\right]^T$ of size $L \times C$ collects all L ground-truth vectors as its rows.

Typically, a generic classification system will have an intermediate block between the the input and the output, a *feature extractor* F, aimed at reducing the dimensionality of the problem; this is followed by a *classifier* C. If the classifier is supervised, we denote it by SC (see Figure 2 and Algorithm 1).

B. Multiresolution Classification

Multiresolution classification is a supervised classification framework (see Figure 2). It decomposes a signal into S localized space-frequency subbands using multiresolution transforms, both bases and frames [21]–[26]. In each subband, multiresolution classification extracts features, classifies them, and produces a local classification decision. A supervised weighting algorithm combines all local decisions into a global decision. This process implicitly mimics the use of wavelet packets, a data-adaptive multiresolution technique [27], and is summarized in Algorithm 2.

1) Multiresolution Decomposition: Multiresolution classification starts with decomposing signals using a given multiresolution transform into several localized space-frequency subbands. For the ith signal, the multiresolution coefficients in the sth subband are

$$a_s^{(i)} \;=\; \mathrm{D}_s(x^{(i)}),$$

where D_s is the multiresolution transform function in that subband. It is implemented using a signal processing device called a *filter bank*; we have a choice of a number of multiresolution transforms D available, both nonredundant (bases) as well as redundant ones (frames) [21]–[26].

2) Feature Extraction: In each subband, features are extracted depending on the application at hand. These features can be generic features, such as texture, Gabor, etc [28]–[35], or can be designed using expert knowledge [36]–[39]. For the ith signal, the feature vector in subband s is

$$f_s^{(i)} = \mathrm{F}_s(a_s^{(i)}),$$

where F_s is the feature extraction function in that subband; different subbands can use different feature extraction functions.

Algorithm 2 Multiresolution classification

Input	$\mathcal{X}$	input dataset
	$\mathcal{Y}$	ground-truth labels for $\mathcal{L}$
Output	$\hat{\mathcal{Y}}$	estimated labels for $\mathcal{U}$
Parameters	**per subband** s	$s = 1, \dots, S$
	D_s	multiresolution function
	$a_s^{(i)}$	multiresolution coefficients
	F_s	feature extraction function
	$f_s^{(i)}$	feature vector
	SC_s	supervised classification function
	$\hat{q}_s^{(i)}$	confidence vector
	for all subbands	
	SW	supervised weighting function
	$\hat{Q}^{(i)}$	confidence matrix
	w	weight vector, $\|w\|_1 = 1$
	$\hat{q}^{(i)}$	final confidence vector
Function	**MRC**$(\mathcal{X})$	
	$a_s^{(i)} = \mathrm{D}_s(x^{(i)})$	multiresolution decomposition
	$f_s^{(i)} = \mathrm{F}_s(a_s^{(i)})$	feature extraction
	$\hat{q}_s^{(i)} = \mathrm{SC}_s(f_s^{(i)}, Q)$	supervised classification
	$\hat{q}^{(i)} = \mathrm{SW}(\hat{Q}^{(i)}, Q)$	supervised weighting
	$\hat{y}^{(i)} = \arg\max_c \hat{q}_c^{(i)}$	
	return $\hat{\mathcal{Y}}$	

3) Supervised Classification: In each subband, the features extracted in the previous stage are fed into a supervised classifier. The classifier can be any state-of-the-art classifier, such as logistic regression or support vector machine [20]. For the ith signal, the confidence vector in the sth subband is

$$\hat{q}_s^{(i)} \;=\; \mathrm{SC}_s(f_s^{(i)}, Q),$$

where SC_s is the supervised classification function in that subband; different subbands can use different classification functions.

4) Supervised Weighting: To combine the subbands' classification decisions, we collect subbands' individual confidence vectors $\hat{q}_s^{(i)}$ into a $C \times S$ confidence matrix $\hat{Q}^{(i)}$, and define the weighting function SW as that taking weighted subbands' individual confidence vectors and producing a single confidence vector,

$$\hat{q}^{(i)} \;=\; \mathrm{SW}(\hat{Q}^{(i)}, Q) \;=\; \hat{Q}^{(i)} w.$$

Here, the $S \times 1$ weight vector w assigns a weight to each subband according to its discriminative power; w_s thus tells us how reliable subband s is. The weight vector is found by optimizing a supervised weighting objective function

$$w \;=\; \arg\min_{\omega} \{ \sum_{i=1}^{L} \|q^{(i)} - \hat{Q}^{(i)}\omega\| \}, \qquad (1)$$

with the constraint $\|\omega\|_1 = 1$. In other words, the optimal weight vector is the one found to be the most reliable over labeled signals only. After weighting, we compute the global decision as

$$\hat{y}^{(i)} \;=\; \arg\max_c \hat{q}_c^{(i)}. \qquad (2)$$

Note that in this section we overloaded the symbol $\hat{q}$, hopefully without confusion; $\hat{q}_c$ is the cth component of the confidence vector $\hat{q}$, while $\hat{q}_s$ is the confidence vector of subband s.

C. Signal Processing on Graphs

With the development of social, biological, and physical networks, signals with complex structure are arising. Traditional discrete signal processing is mainly suited to processing regularly sampled low-dimensional signals, such as discrete time and space signals. To mitigate the problem, signal processing on graphs is emerging as a tool to analyze high-dimensional signals with irregular structure [11], [40], [41], defined on a more general domain.

We focus here on one of the recent developments, discrete signal processing on graphs [10]. The dataset is represented by a *graph* $\mathcal{G} = (\mathcal{V}, P)$, where $\mathcal{V} = \{v_i\}_{i=1}^N$ is the set of *nodes* representing signals and $P \in \mathrm{C}^{N\times N}$ is a *graph shift* describing the relational dependencies among the nodes. The graph shift P is not necessarily an adjacency matrix nor does it necessarily have a probabilistic meaning. A *graph signal* s is then defined as the following map:

$$s \ : \ \mathcal{V} \to \mathrm{C}^{N\times D},$$

where D is the dimension of the graph signal on each node. A linear shift-invariant system, or, a *graph filter*, is defined as

$$H \ = \ h(P) \ = \ \sum_{k=0}^{K} h_k P^k, \tag{3}$$

with $h_k \in \mathrm{C}$, $k = 0, 1, \ldots, K$.

Then, a graph filter $H \in \mathrm{C}^{N\times N}$ applied to a graph signal $s \in \mathrm{C}^{N\times D}$ produces an output, which is again a graph signal,

$$Hs \ = \ h(P)s.$$

Discrete signal processing on graphs then defines a series of standard signal processing concepts including the graph Fourier transform, frequency, spectrum, spectral decomposition, and impulse and frequency responses [10].

D. Semi-Supervised Learning

Traditional classifiers typically fall under supervised learning, with only labeled signals to train. In many real-world applications, however, a large number of labeled signals is not available, which can cause overfitting. Semi-supervised learning is a technique for training classifiers with both labeled and unlabeled signals, which assumes that unlabeled signals can provide distribution information to build a stronger classifier. Some well-known semi-supervised learning algorithms include generative mixture models with expectation maximization, co-training and graph-based approaches [9]. Generative mixture models with expectation maximization assume that classes produce well clustered signals, and that with large number of unlabeled signals, the mixture components can be identified [42]. Co-training assumes that the features are discriminative enough so they can be split into two sets, with each set being able to build a good classifier [43]. Graph-based approaches assume that while the measured signals are defined in a high-dimensional space, they exist in a low-dimensional manifold; a graph is then constructed by measuring the similarity of each pair of signals, and those deemed similar are labeled as belonging to the same class [44].

We focus here on label propagation, one of graph-based approaches. Label propagation classifies signals by understanding how labels propagate on a graph; two methods are in use, diffusion functions [45] and harmonic functions [46], [47]. Both methods work based on propagating the known labels on the transition matrices. Diffusion functions propagate those labels a finite number of times without any intervention. Harmonic functions, on the other hand, correct the known labels to the initial values after each propagation and propagate an infinite number of times. The advantage of harmonic functions is that the known labels keep pushing the decision boundaries to low-density gaps. The drawback is that if the known labels are not reliable, harmonic functions may keep diffusing wrong information.

III. Semi-Supervised Multiresolution Classification

Multiresolution classification analyzes data to uncover hidden information; in its original form, it uses supervised classification, and can thus train on labeled signals only. When the labeled dataset is small or contains improperly labeled signals, the classification boundary and the weights assigned to subbands can be unreliable. Semi-supervised learning, on the other hand, uses the entire dataset to help classification, but works on one resolution level only. We thus propose to merge these two concepts and gain the best of both worlds: a *semi-supervised multiresolution classifier*. We stress here that this is not a simple combination of known techniques, as there is no known way to weigh subband decisions for unlabeled signals; this is one of our contributions.

A. Semi-Supervised Classification

Figure 3 summarizes our proposed framework; the multiresolution decomposition and feature extraction blocks from Figure 2 work as before. The first change is that the supervised classification block is replaced by a semi-supervised one so we can use both labeled and unlabeled signals to make a labeling decision in each subband. For the ith signal, the confidence vector in the sth subband is now

$$\hat{q}_s^{(i)} \ = \ \mathrm{SSC}_s(f_s^{(i)}, Q), \tag{4}$$

where SSC_s is the semi-supervised classification function in subband s. As for supervised classification, SSC can be chosen from a variety of approaches; we propose a new one, adaptive graph filter, described in Section IV.

B. Semi-Supervised Weighting

We now explain how to build a semi-supervised weighting block, that is, how to weigh decisions from all the subbands to get a global decision in a semi-supervised manner.

Labeled signals contribute to weighting directly by fitting their confidence vectors from all the subbands to the ground truth; unlabeled signals cannot do the same as they do not have the ground truth. We could use Shannon entropy to measure the confidence of labeling an unlabeled signal; if the entropy is small (less uncertainty, high confidence), it is

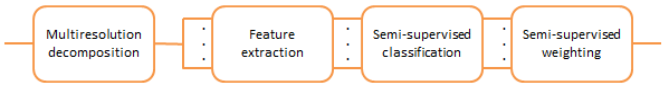

Fig. 3: Semi-supervised multiresolution classification. Supervised classification and weighting algorithm in Figure 2 are replaced with their semi-supervised counterparts so that unlabeled data can contribute to classification.

easy to assign a label to the signal, and vice versa. We could thus normalize each confidence vector to sum to 1 to measure its entropy. We encounter a problem, however; for example, let $\hat{q}^{(1)} = \begin{bmatrix}0.5 & 0.5 & 0\end{bmatrix}^T$ and $\hat{q}^{(2)} = \begin{bmatrix}0.5 & 0.25, & 0.25\end{bmatrix}^T$ be confidence vectors for two signals. While we can label $\hat{q}^{(2)}$ as Class 1, we clearly cannot make a decision for $\hat{q}^{(1)}$. The entropy measure, however, tells us that we can label $\hat{q}^{(1)}$ with higher confidence (less uncertainty) because its entropy, $\mathrm{H}(\hat{q}^{(1)}) = 1$, is lower than that for $\hat{q}^{(2)}$, $\mathrm{H}(\hat{q}^{(2)}) = 3/2$. To resolve this issue, we define a new uncertainty measure,

$$\mathrm{M}(\hat{q}) \;=\; \mathrm{H}(\hat{q})\,(\chi_{d>T} + \lambda(d)\,\chi_{d\leq T}), \tag{5}$$

where χ_I is the indicator function of an interval I, $d = |\hat{q}_{(1)} - \hat{q}_{(2)}|$ with $\hat{q}_{(1)}, \hat{q}_{(2)}$ the first and second largest element in $\hat{q}$, respectively, T is the threshold, and $\lambda(d)$ is a penalty function that is large when the first and second largest elements are close. In other words, when the first and second largest elements in $\hat{q}$ are far apart, d is large, the first term in (5) takes over and $\mathrm{M}(\hat{q}) = \mathrm{H}(\hat{q})$; when, on the other hand, the first and second largest elements in $\hat{q}$ are close, d is small, the second term in (5) takes over and the uncertainty $\mathrm{M}(\hat{q})$ is large.

We can now use this new uncertainty measure to say that the uncertainty of the sth subband in labeling the ith signal is $\mathrm{M}(\hat{q}_s^{(i)})$. Since entropy is additive, the total uncertainty of a subband when classifying unlabeled signals is the mean uncertainty over all the unlabeled signals for this subband,

$$\mathrm{M}_s \;=\; \frac{1}{U}\sum_{i=L+1}^{N} \mathrm{M}(\hat{q}_s^{(i)}).$$

We now define the discriminative power of the sth subband to be the confidence

$$g_s \;=\; \frac{e^{-\beta \mathrm{M}_s}}{\sum_{j=1}^{S} e^{-\beta \mathrm{M}_j}}, \tag{6}$$

where β is the decay coefficient that controls the distribution of the discriminant power from all the subbands. When the uncertainty of a subband is large, the confidence is small and the subband gets assigned a low weight, and vice versa. Confidences from all the subbands are collected into a vector g; note that $\|g\|_1 = 1$.

We now find the weight vector by optimizing a *semi-supervised weighting objective function*,

$$\begin{aligned} w \;&=\; \arg\min_{\omega}\; \{\alpha \sum_{i=1}^{L} \|q^{(i)} - \hat{Q}^{(i)}\omega\| \\ &+\; (1-\alpha)\|\omega - g\|\}, \end{aligned} \tag{7}$$

Algorithm 3 Semi-supervised multiresolution classification

Input	$\mathcal{X}$	input dataset
	$\mathcal{Y}$	ground-truth labels for $\mathcal{L}$
Output	$\hat{\mathcal{Y}}$	estimated labels for $\mathcal{U}$
Parameters	**per subband** s	$s = 1, \ldots, S$
	D_s	multiresolution function
	$a_s^{(i)}$	multiresolution coefficients
	F_s	feature extraction function
	$f_s^{(i)}$	feature vector
	SSC_s	semi-supervised classification function
	$\hat{q}_s^{(i)}$	confidence vector
	for all subbands	
	SSW	semi-supervised weighting function
	$\hat{Q}^{(i)}$	confidence matrix
	w	weight vector, $\|w\|_1 = 1$
	$\hat{q}^{(i)}$	final confidence vector
Function	**SSMRC**$(\mathcal{X})$	
	$a_s^{(i)} = \mathrm{D}_s(x^{(i)})$	multiresolution decomposition
	$f_s^{(i)} = \mathrm{F}_s(a_s^{(i)})$	feature extraction
	$\hat{q}_s^{(i)} = \mathrm{SSC}_s(f_s^{(i)}, Q)$	semi-supervised classification
	$\hat{q}^{(i)} = \mathrm{SSW}(\hat{Q}^{(i)}, Q)$	semi-supervised weighting
	$\hat{y}^{(i)} = \arg\max_c \hat{q}_c^{(i)}$	
	return $\hat{\mathcal{Y}}$	

with the constraint $\|\omega\|_1 = 1$, and where $\alpha = L/(L+U)$ is the *labeling ratio*. The first term in (7) represents the contribution from all labeled signals and is a scaled version of (1). The second term in (7) represents the contribution from all unlabeled signals; to obtain it, we fit weights to subbands' confidences. We use the labeling ratio to balance these two terms. Since this is a convex optimization problem, it is numerically efficient to solve. As in (2), after weighting, we compute the global decision as

$$\hat{y}^{(i)} \;=\; \arg\max_c \hat{q}_c^{(i)}, \tag{8}$$

where $\hat{q}^{(i)} = \mathrm{SW}(\hat{Q}^{(i)}, Q) = \hat{Q}^{(i)}w$ (see Algorithm 3).

IV. Adaptive Graph Filter

The idea of using a graph filter as a binary classifier was first proposed in [10]. For practical applications, however, its use is limited: First, it can only perform binary classification; then, it trains the filter coefficients by choosing a local optimum; and finally, it cannot classify unseen signals. Here, we propose an *adaptive graph filter* as a semi-supervised classifier in (4) that resolves these problems; we also connect this new adaptive graph filter to diffusion maps and diffusion wavelets.

In (4), for the signal $x^{(i)}$ and subband s, the inputs to the semi-supervised classifier are the feature vector $f_s^{(i)}$ and the ground-truth matrix Q, and the output is the confidence vector $\hat{q}_s^{(i)}$. For simplicity, in this section we omit the subband index s and write $f^{(i)}$, $\hat{q}^{(i)}$; it should be understood, however, that an adaptive graph filter is applied in each subband.

A. Graph Filtering as Semi-Supervised Classification

We start by outlining the basic idea, followed by detailed developments. Let the input graph signal be a prior confidence matrix formed from the ground-truth matrix and the graph

shift be the Hermitian transpose of the transition matrix. An adaptive graph filter is then built by combining a series of graph shifts. The filter coefficients of the graph filter are trained by fitting the estimated results to the known labels and minimizing the labeling uncertainty. The output graph signal after filtering is the posterior confidence matrix whose ith row will be the desired confidence vector $\hat{q}^{(i)}$.

Let $\mathcal{G} = (\mathcal{F}, P)$ be a graph with $\mathcal{F} = \{f^{(i)}\}_{i=1}^{N}$ a set of feature vectors in the given subband for the entire dataset and let $P \in \mathbb{R}^{N\times N}$ be a graph shift defined as

$$P_{i,j} \;=\; \frac{\exp\left(\frac{-\rho(f^{(i)},f^{(j)})}{\sigma}\right)}{\sum_{i=1}^{N}\exp\left(\frac{-\rho(f^{(i)},f^{(j)})}{\sigma}\right)}, \tag{9}$$

where ρ is a local distance measurement, such as the ℓ^2 norm or the cosine distance, and σ is a scaling coefficient, which controls the bandwidth. The graph shift we defined here is the Hermitian transpose of the *transition matrix* of the graph. The graph shift thus has a probabilistic interpretation: $P_{i,j}$ gives the probability that the jth node jumps to the ith node in one step [48].

We now build a graph filter as in (3), except that, because of the dependencies on the data in (9), this is an *adaptive graph filter*,

$$H \;=\; h(P) \;=\; \sum_{k=1}^{K} h_k P^k. \tag{10}$$

Note that we omit the 0th term since, as we will see, it does not contribute to classification. The graph filter thus represents the relational dependencies among signals represented via their feature vectors.

Let the graph signal be the confidence matrix of *all* the signals on the graph, called *prior confidence matrix*, that is, the following map:

$$s \;:\; \mathcal{F} \to \mathbb{R}^{N\times C},$$

defined as

$$(\hat{Q}_{\rm pr})_{j,c} \;=\; \begin{cases} 1, & \text{when } y^{(j)} = c; \\ 0, & \text{otherwise,} \end{cases}$$

or,

$$\hat{Q}_{\rm pr} \;=\; \begin{bmatrix} Q \\ \mathbf{0}_{U\times C} \end{bmatrix}.$$

In other words, the first L rows of $\hat{Q}$ are the confidence matrix Q representing the labeled dataset, while the other U rows are all zeros representing the unlabeled dataset. The prior confidence matrix thus starts with the prior knowledge on the labeled dataset (the ground-truth matrix) and without any knowledge on the unlabeled dataset.

By applying adaptive graph filtering now, the ground truth propagates from the labeled dataset to unlabeled dataset; the output graph signal, or, the *posterior confidence matrix*, is obtained as

$$\hat{Q}_{\rm ps} \;=\; H\hat{Q}_{\rm pr}, \tag{11}$$

where the ith row of the posterior confidence matrix $\hat{Q}_{\rm ps}$ is the desired confidence vector $\hat{q}^{(i)}$. Note that graph filtering propagates the labeling information not only from the labeled signals to the unlabeled signals, but among the labeled signals as well, giving the mislabeled signals a chance to be corrected, and consequently providing robustness in classification.

One issue left to address is how to choose the filter coefficients $h_1, h_2, \ldots h_K$. We rewrite (11) as

$$\hat{Q}_{\rm ps} \;\overset{(a)}{=}\; \sum_{k=1}^{K} h_k P^k \hat{Q}_{\rm pr} \;=\; \sum_{k=1}^{K} h_k \hat{Q}_k, \tag{12}$$

where (a) follows from (10) and $\hat{Q}_k$ is the $N \times C$ confidence matrix for the kth graph shift defined as $\hat{Q}_k = P^k\hat{Q}_{\rm pr}$. For each i, the desired confidence vector $\hat{q}^{(i)}$ (ith row of $\hat{Q}_{\rm ps}$) is thus a weighted linear combination of corresponding rows from each graph shift $\hat{Q}_k$. We can now see why we omitted the 0th term in (10); since $\hat{Q}_0 = \hat{Q}_{pr}$, it does not contribute to propagating information from the labeled to the unlabeled data. Fitting these estimates to the known labels and minimizing the labeling uncertainty is identical to the semi-supervised weighting we performed in the last section. We can thus use the same minimization as in (7),

$$\begin{aligned} h^* \;&=\; \arg\min_{h} \{\alpha \sum_{i=1}^{L} \|q^{(i)} - \hat{Q'}^{(i)} h\| \\ &+\; (1-\alpha)\|h - \gamma\|\}, \end{aligned} \tag{13}$$

where $h = \begin{bmatrix} h_1 & \ldots & h_K \end{bmatrix}^T$ with the constraint $||h||_1 = 1$, α is the labeling ratio as in (7), $\hat{Q'}^{(i)} = \begin{bmatrix} \hat{q}_1^{(i)} & \ldots & \hat{q}_K^{(i)} \end{bmatrix}$ is a $C \times K$ confidence matrix of the ith signal that collects the graph shifts' individual confidence vectors, and $\gamma = \begin{bmatrix} \gamma_1 & \ldots & \gamma_K \end{bmatrix}^T$ collects the discriminative powers of each graph shift,

$$\gamma_k \;=\; \frac{e^{-(\beta/U)\sum_{i=L+1}^{N} \mathrm{M}(\hat{q}_k^{(i)})}}{\sum_{j=1}^{K} e^{-(\beta/U)\sum_{i=L+1}^{N} \mathrm{M}(\hat{q}_j^{(i)})}},$$

where M is the uncertainty measurement as in (5) and β is the decay coefficient as in (6).

In the first term of (13), we fit the estimated confidence vectors to the ground-truth vectors by changing the filter coefficients. Since the 0th term always stays within the ground truth, it trends to have a larger value and does not contribute to the unlabeled data. This is another reason we omit the 0th term in (10).

B. Regression: Handling Unseen Data

As defined, adaptive graph filtering can only handle signals in the given dataset; should an unseen signal appear, the graph would need to be rebuilt and the filter coefficients retrained, at a significant computational cost. To handle unseen signals, we introduce regression.

We assume that each signal is randomly sampled from some continuous distribution and that the signals with the same label originate from the same distribution (recall that signals here are subband feature vectors). If we use the given signals and their posterior confidence vectors $\hat{q}^{(i)}$ to estimate the distributions, we can label those unseen signals originating from these distributions. The task is thus to design a regression

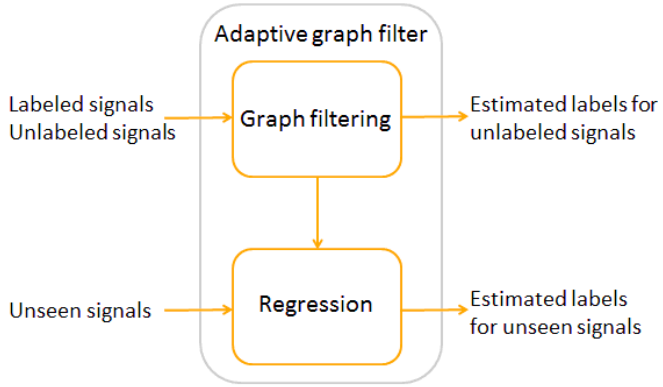

Fig. 4: Adaptive graph filter. Labeled and unlabeled signals are fed into the graph filtering block that outputs the estimated labels for unlabeled signals. Based on existing labeled and unlabeled signals, graph filtering generates the smoothing matrix for the regression block. Unseen signals are fed into the regression block that outputs the estimated labels for unseen signals.

function to map the unseen signals to their posterior confidence vectors.

Given the subband feature vectors $\mathcal{F} = \{f^{(i)}\}_{i=1}^{N}$ and their posterior confidence vectors $\hat{Q}_{\rm ps}$, for a batch of unseen signals $\mathcal{F}_{\rm us} = \{f^{(i)}\}_{i=N+1}^{N+M}$, the posterior confidence matrix for the unseen signals is

$$\hat{Q}_{\rm us} = L\hat{Q}_{\rm ps}, \tag{14}$$

where $L \in \mathbb{R}^{M\times N}$ is a generic form of the smoothing matrix determined by some regression technique, such as polynomial regression, spline regression or kernel regression [49]. For the regularized reproducing kernel regression, a valid kernel function $k(f, f')$ is first chosen to measure the inner product of f and f' in a higher-dimensional space, and then, the smoothing matrix is defined as

$$L = K'(K + \lambda I_N)^{-1},$$

where λ is a regularization parameter, I_N is an $N \times N$ identity matrix, $K'_{i,j} = k(f^{(N+i)}, f^{(j)})$, $i = 1, 2, \ldots, M$, $j = 1, 2, \ldots, N$, and $K_{i,j} = k(f^{(i)}, f^{(j)})$, $i, j = 1, 2, \ldots, N$. The smoothing matrix calculates the relational dependencies between the unseen and given signals. Closer dependencies lead to higher values in the smoothing matrix. Because of this, the generic regression model predicts the posterior confidence vector of the new signal by weighing the posterior confidence vectors of the given signals locally, which means that the signals close by have similar posterior confidence vectors.

Note that the forms of one step of graph filtering (11) and one step of regression (14) are similar. While both of these assume that similar signals have similar confidence vectors, their goals are different: the graph filtering step builds the relationship between unlabeled and labeled signals and produces labels from limited label information, while the regression step builds the relationship among the signals with the same label and connects unseen signals to their confidence

Algorithm 4 Adaptive graph filter

Input	$\mathcal{F}$	input dataset
	$\hat{Q}_{\rm pr}$	prior confidence matrix
Output	$\hat{Q}_{\rm ps}$	posterior confidence matrix
Parameters	$P_{i,j}$	graph shift
	ρ	local distance
	σ	scaling coefficient
	h_k	filter coefficients
Function	**AGF**$(\mathcal{F}, \hat{Q}_{\rm pr})$	
	$P_{i,j} =$ (9)	graph shift construction
	$h =$ (13)	filter coefficient optimization
	for $k = 1 : K$ **do**	diffusion
	$\hat{Q}_k = P^k \hat{Q}_{\rm pr}$	
	$\hat{Q}_{\rm ps} = \sum_{k=1}^{K} h_k \hat{Q}_k$	weighting
	return $\hat{Q}_{\rm ps}$	

vectors. Therefore, adding regression to graph filtering creates an adaptive graph filter that not only propagates the labeling information within a given dataset, but across unseen signals as well (see Figure 4).

Adaptive graph filter serves as a semi-supervised plugin method for classification. The traditional plugin methods estimate the unknown quantities in the Bayes' rule and plug them in for classification [50]; in the other words, they are equivalent to the regression block of adaptive graph filter. For instance, to label the ith signal, we first estimate the confidence vector $\hat{q}^{(i)}$ by using some regression techniques on the labeled dataset and then plug $\hat{q}^{(i)}$ in (2) to do classification. Without a large number of labeled signals, however, the traditional plugin methods fail to do robust semi-supervised classification [51]. Adaptive graph filter solves this problem by producing labels for unlabeled signals in the filtering step, such that both labeled and unlabeled signals contribute to the smooth matrix in the regression step.

C. Cost Analysis

The adaptive graph filter contains two steps: filtering and regression. In the filtering step, the computation involves the graph shift construction with the cost of $O((N-1)N/2 + N + N^2) = O((3N^2 + N)/2)$, the diffusion operation with the cost of $O(CKN^2)$, and the weighting operation with the cost of $O((K-1)N^2)$, for a total cost of $O(((C+1)K + 1/2)N^2 + N/2)$. In the regression step, the bulk of the cost comes from the construction of the smooth matrix and the inverse it involves with the cost of $O((N-1)N/2 + N + N^3 + MN + MN^2)$ and the matrix multiplication in (14) with the cost of $O(MCN)$, for a total cost of $O(N^3 + (M + 1/2)N^2 + (1/2 + M(C+1))N$.

D. Relation to Diffusion Maps

We now analyze the adaptive graph filter by connecting it to diffusion maps and show that it reconstructs a robust diffusion map with more flexibility.

1) Diffusion Maps: The diffusion maps are coordinates that provide efficient geometric descriptions of signals and are built based on the singular value decomposition of the transition matrix.

Let A be the adjacency matrix of a graph, D the diagonal matrix whose ith element is $D_{i,i} = \sum_j A_{i,j}$ and $T = D^{-1}A$ the transition matrix. Recall that each element of T measures the likelihood of getting from one data point to another in one step; each element of T^k measures the likelihood of getting from one data point to another in k steps. The diffusion distance between two signals x, y in step k is then defined as

$$\begin{aligned} D_k(x,y)^2 &= ||t_k^{(x)} - t_k^{(y)}||^2_{D^{-1}} \\ &= (t_k^{(x)} - t_k^{(y)})^T D^{-1} (t_k^{(x)} - t_k^{(y)}), \end{aligned} \tag{15}$$

where $t_k^{(x)}$ and $t_k^{(y)}$ are xth and yth rows of T^k, respectively. Since the diffusion distance takes into account all paths of length k from x to y, it is robust to noise perturbation and outliers. Note that the diffusion distance can also be calculated using the ℓ^2 norm.

Since the transition matrix is asymmetric, we introduce a normalized transition matrix $\hat{T} = D^{\frac{1}{2}} T D^{-\frac{1}{2}}$, which, since symmetric, can be factored as

$$\hat{T} \;=\; V \Lambda V^* \;=\; \sum_{i=1}^{N} v_i \lambda_i v_i^*,$$

where $V = \begin{bmatrix} v_1 & v_2 & \dots & v_N \end{bmatrix}$ is an orthogonal matrix and λ_i are the singular values; moreover, $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N \geq 0$. We then decompose the transition matrix as

$$T^k \;=\; \Psi \Lambda^k \Phi^* \;=\; \sum_{i=1}^{N} \psi_i \lambda_i^k \phi_i^*, \tag{16}$$

where $\Psi = \begin{bmatrix} \psi_1 & \psi_2 & \dots & \psi_N \end{bmatrix} = D^{-\frac{1}{2}} V$, and $\Phi = \begin{bmatrix} \phi_1 & \phi_2 & \dots & \phi_N \end{bmatrix} = D^{\frac{1}{2}} V$. Note that $\hat{T}$ and T share the same eigenvalues. The underlying diffusion map for step k and node x is now defined as

$$Y_k^{(x)} \;=\; \begin{bmatrix} \lambda_1^k \psi_{1,x} & \lambda_2^k \psi_{2,x} & \dots & \lambda_N^k \psi_{N,x} \end{bmatrix}^T. \tag{17}$$

If we define the diffusion space to be the space spanned by the columns of Φ, then, the diffusion map $Y_k^{(x)}$ gives the coordinates of x diffused k times in that space.

Using diffusion maps, the diffusion distance is simply

$$\begin{aligned} D_k(x,y)^2 &\overset{(a)}{=} (t_k^{(x)} - t_k^{(y)})^T D^{-1} (t_k^{(x)} - t_k^{(y)}), \\ &\overset{(b)}{=} \left[\sum_{i=1}^N \lambda_i^k (\psi_{i,x} - \psi_{i,y}) \phi_i^*\right] D^{-1} \\ &\quad \left[\sum_{i=j}^N \lambda_j^k (\psi_{j,x} - \psi_{j,y}) \phi_j\right] \\ &= \left[\sum_{i=1}^N \lambda_i^k (\psi_{i,x} - \psi_{i,y}) v_i^* D^{-\frac{1}{2}}\right] D^{-1} \\ &\quad \left[\sum_{i=j}^N \lambda_j^k (\psi_{j,x} - \psi_{j,y}) D^{\frac{1}{2}} v_j\right] \\ &\overset{(c)}{=} \sum_{i=1}^{N} \left[\lambda_i^k (\psi_{i,x} - \psi_{i,y})\right]^2 \\ &= ||Y_k^{(x)} - Y_k^{(y)}||_2^2, \end{aligned}$$

where (a) follows from (15); (b) from (16); and (c) from the orthogonality of V. Thus, by varying $k = 0, 1, \dots$, the diffusion maps allow us to find an alternate representation that might better separate the data, for example.

2) Relation to the Adaptive Graph Filter: We now show how, by using the adaptive graph filter, we construct diffusion maps that allow for more flexibility. Recall that our adaptive graph filter in (10) uses the graph shift P that is a Hermitian transpose of the transition matrix T. Thus, applying (16) to (10) , we get

$$\begin{aligned} h(P) &= \sum_{k=1}^{K} h_k P^k \;=\; \sum_{k=1}^{K} h_k (T^*)^k \;=\; \sum_{k=1}^{K} h_k (T^k)^* \\ &= \sum_{k=1}^{K} h_k (\Psi \Lambda^k \Phi^*)^* \;=\; \sum_{k=1}^{K} h_k \Phi \Lambda^k \Psi^* \\ &= \Phi \left(\sum_{k=1}^{K} h_k \Lambda^k \right) \Psi^* \;=\; \Phi\, h(\Lambda)\, \Psi^*, \end{aligned} \tag{18}$$

with $h(\Lambda) = \mathrm{diag}(\sum_{k=1}^K h_k \lambda_1^k, \dots, \sum_{k=1}^K h_k \lambda_N^k)$.

If we define the diffusion space to be the space spanned by the columns of Ψ, then, we define a diffusion map to be

$$Y_h^{(x)} \;=\; \begin{bmatrix} h(\Lambda)_{1,1} \phi_{1,x} & \dots & h(\Lambda)_{N,N} \phi_{N,x} \end{bmatrix}^T. \tag{19}$$

Thus, the construction of an adaptive graph filter allows for continuous change of the coordinates in the diffusion space, providing flexibility in finding the best representation in a data-adapted fashion; this adaptivity is reflected by subscript h in $Y_h^{(x)}$. This is in contrast to only discrete changes allowed by (17).

We illustrate the above discussion with an example. Let $\lambda_1 = 1$, $\lambda_2 = 0.5$, $Y_0^{(x)} = \begin{bmatrix} 1 & 1 \end{bmatrix}^T$, $Y_0^{(y)} = \begin{bmatrix} 2 & 1 \end{bmatrix}^T$, and $Y_0^{(z)} = \begin{bmatrix} 1 & 2.4 \end{bmatrix}^T$. The distance between x and y in this step is $D_0(x,y) = 1$, smaller than the distance between x and z, $D_0(x,z) = 1.4$. Diffusing once, we get $Y_1^{(x)} = \begin{bmatrix} 1 & 0.5 \end{bmatrix}^T$, $Y_1^{(y)} = \begin{bmatrix} 2 & 0.5 \end{bmatrix}^T$, and $Y_1^{(z)} = \begin{bmatrix} 1 & 1.2 \end{bmatrix}^T$. The distance between x and y in this step is $D_1(x,y) = 1$, larger than the distance between x and z, $D_1(x,z) = 0.7$. In other words, by changing the power k, distances change in the diffusion space. Since k can only be an integer, it is not possible, for example, to make the distances between x and y, and x and z be the same. By using the adaptive graph filter, however, we can find the optimal filter coefficients to match such a requirement using (13).

E. Relation to Diffusion Wavelets

We now analyze the adaptive graph filter by connecting it to diffusion wavelets and show that it performs multiresolution classification on graphs.

1) Diffusion Wavelets: Diffusion wavelets are a multiscale framework to analyze signals with complex structures [13]. They can be seen as an extension of the classical wavelet theory, where, the diffusion wavelet basis is learned from the geometry of the signal structure in a data-adapted fashion. The diffusion wavelet basis is constructed by dilation using the dyadic powers of the transition matrix, the idea being that they propagate local relationships to global relationships throughout the graph.

Given a graph, at the jth resolution level, we have T^{2^j}, as the transition matrix, $j = 1, 2, \dots$. Since the second singular

value of the transition matrix is less than 1 to keep the graph connected, only the first singular value is 1. Consequently, if the transition matrix is raised to a high power, all the singular values disappear except for the first one,

$$\lambda_i^{2^j} \;\to\; 0, \qquad i = 2, 3, \ldots, N \quad \text{as} \quad j \to \infty.$$

When $j = 1$, the transition matrix T measures local pairwise similarities; increasing the power j gradually decreases the rank of the transition matrix T^{2^j} and causes local information of the graph to be missed since the resolution on the graph changes from finest to the coarsest. Thus, by changing j, we can both perform a multiresolution analysis as well as do it in a computationally efficient manner.

2) Relation to the Adaptive Graph Filter: The adaptive graph filter in (10) is formed as a linear combination of graph shifts P raised to power k. When k is large, P^k become a low-rank matrix describing the global information of the graph, just as T^{2^j} does for diffusion wavelets. Each k corresponds to a different resolution on the graph, and thus, adaptive graph filter actually performs multiresolution classification on the graph. It weighs the classification results from each resolution to produce the global result. The filter coefficients represent the discriminative power of each resolution. This also explains why the objective functions to optimize the filter coefficients in (7) and the weights of each subband in (13) are the same.

V. Experimental Results

In this section, we validate the proposed framework on a real-world classification problem, indirect bridge structural health monitoring, and show that it performs remarkably better than previous approaches.

A. Dataset

To study the bridge behavior under various conditions comprehensively, a lab-scale bridge-vehicle dynamic system was built. Accelerometers were installed on a vehicle that moves across the bridge; acceleration signals were then collected from those accelerometers. We collected 30 acceleration signals for each of 13 different bridge conditions, 8 different speeds and 2 vehicles with different weights, for a total of 6240 acceleration signals [19]. The 13 bridge conditions include one pristine condition and 4 different damage severities for each of 3 different damage proxy scenarios.

B. Experimental Setup

Given a specific vehicle driven at a specific speed, we want to classify 13 bridge conditions, in particular, with a low labeling ratio. We consider 16 vehicle-speed combinations for each of which there are 30 acceleration signals for each of the 13 scenarios; the final accuracy is the average over the 13 scenarios; the baseline accuracy is thus $100/13 = 7.7\%$.

We choose a Coiflet filter bank [52] with 4 levels [21] in the multiresolution block, principal component analysis [20] in the feature extraction block, radius kernel support vector machine [53] in the supervised classification block, and adaptive graph filtering in the semi-supervised classification block. We construct the graph by choosing ρ as the cosine distance [54] and $\sigma = (1/N^2)\sum_{i,j}\rho(f^{(i)},(f^{(j)}))$ in (9). The length of the adaptive graph filter is $K = 30$. In the semi-supervised weighting function, we choose the penalty threshold $T = 0.02$ and the penalty function $\lambda(d) = 1 + 5(d/T - 1)^2$ in (5). The decay coefficient $\beta = 1$ is chosen to minimize (7) and (13). We performed a 30-fold cross-validation. Table II summarizes all the parameters at a glance.

Experimental setup		
dataset		
V	vehicles	2
S	speeds	8
	damage scenarios	13
	signals/scenario	30
multiresolution decomposition		
D	Coiflet filter bank	4 levels
S	number of subbands	31
feature extractions		
F	principal component analysis	
classification		
SC	kernel support vector machine	
SSC	adaptive graph filter	
ρ	local distance	cosine distance
σ	scaling coefficient	$(1/N^2)\sum_{i,j}\rho(f^{(i)},f^{(j)})$
β	decay coefficient	1
K	length of graph filter	30
weighting		
T	threshold	0.02
$\lambda(d)$	penalty function	$1 + 5((d/T) - 1)^2$

TABLE II: Parameters used in the experiments.

We compare our method, semi-supervised multiresolution classification with adaptive graph filtering, against:

- *Generic classification* with kernel support vector machine, diffusion functions, harmonic functions, and adaptive graph filtering.
- *Supervised multiresolution classification* with kernel support vector machine.
- *Semi-supervised multiresolution classification* with diffusion functions, harmonic functions, and adaptive graph filtering.

We use the following shorthands in figures and tables: S for supervised (if with another acronym) or speed (if without another acronym), SS for semi-supervised, MRC for multiresolution classification, SVM for kernel support vector machine, DF for diffusion functions, HF for harmonic functions, AGF for adaptive graph filtering, and V for vehicle.

C. Classification Results

We validate our method from three standpoints: (1) the performance of the semi-supervised multiresolution classification framework; (2) the ability of adaptive graph filtering to handle mislabeled signals; and (3) the ability of adaptive graph filtering to handle unseen signals.

1) Semi-Supervised Multiresolution Classification: We validate the proposed framework, semi-supervised multiresolution classification, by comparing it to generic classification and supervised multiresolution classification with a low labeling ratio of 10% (see Table III).

V	S	Generic				SMRC		SSMRC	
		SVM	DF	HF	AGF	SVM	DF	HF	AGF
1	1	57.97	83.70	88.92	87.98	81.19	99.63	99.80	**99.83**
	2	70.42	85.78	90.38	89.55	84.71	99.83	99.98	**99.99**
	3	74.29	86.25	90.67	89.91	84.59	99.11	**99.40**	99.38
	4	74.82	88.08	94.79	93.87	80.29	99.89	99.93	**99.98**
	5	70.68	74.87	78.92	77.64	72.35	93.76	**96.13**	94.52
	6	67.13	82.43	86.00	85.24	69.72	91.54	**93.47**	93.43
	7	59.48	65.31	66.05	66.23	59.75	77.44	**78.79**	78.74
	8	56.53	66.52	67.58	67.11	56.63	75.37	**77.93**	77.23
2	1	49.75	78.17	82.28	80.99	71.69	**85.36**	84.63	85.12
	2	53.64	67.30	71.38	70.32	60.69	80.19	80.51	**80.59**
	3	67.96	82.50	87.84	86.09	74.94	**95.12**	94.62	94.81
	4	61.52	79.57	82.85	82.86	65.27	86.65	86.44	**87.34**
	5	62.75	77.92	82.15	81.22	66.42	**88.86**	88.34	88.72
	6	66.89	80.17	81.65	81.73	69.75	84.08	83.61	**84.18**
	7	65.09	82.87	85.19	85.48	70.59	89.15	89.51	**89.66**
	8	48.57	80.32	83.01	82.57	53.67	92.03	**93.62**	93.31

TABLE III: Accuracy comparison of Vehicles (V) 1 and 2, with Speeds (S) $1, 2, \ldots, 8$, and labeling ratio of 10%.

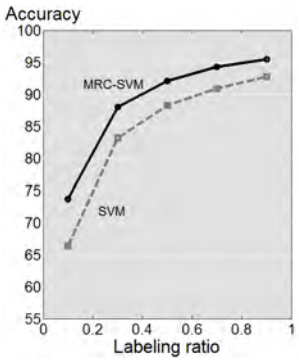

(a) Vehicle 1.

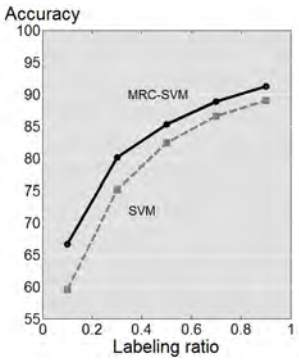

(b) Vehicle 2.

Fig. 5: Trend 1: Multiresolution framework improves classification accuracy. Comparison of supervised classifiers.

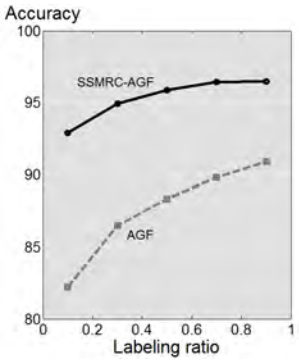

(a) Vehicle 1.

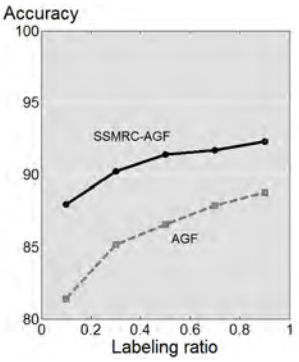

(b) Vehicle 2.

Fig. 6: Trend 1: Multiresolution framework improves classification accuracy. Comparison of semi-supervised classifiers.

We detect three trends, the first two of which validate each component of the framework (multiresolution classification and semi-supervised learning), while the third validates the entire framework: (1) Multiresolution framework improves classification accuracy: supervised multiresolution classifier performs better than the corresponding supervised generic classifier; and each semi-supervised multiresolution classifier (DF, HF, AGF) performs better than the corresponding semi-supervised generic classifier. (2) Semi-supervised learning improves classification accuracy: each semi-supervised generic classifier (DF, HF, AGF) performs better than the supervised generic classifier; and each semi-supervised multiresolution classifier performs better than the supervised multiresolution classifier. (3) Multiresolution framework with semi-supervised learning improves classification accuracy over the supervised generic classifier by $\sim 30\%$.

We further validate these trends under different labeling ratios. Figures 5–8 show the dependence of classification accuracy on the labeling ratio for 2 vehicles averaged across 8 speeds. Figures 5 and 6 validate Trend 1 for supervised (SVM) and semi-supervised (AGF) classifiers. In each case and for both vehicles, multiresolution framework improves classification accuracy across all labeling ratios.

Figure 7 validates Trend 2 for supervised (SVM) and semi-supervised (AGF) multiresolution classifiers. For both vehicles, semi-supervised learning improves classification accuracy across all labeling ratios. Moreover, as the labeling ratio decreases, accuracy drops sharply for multiresolution-based SVM; performance of the multiresolution-based adaptive graph filter stays relatively flat, however, even at very low labeling ratios.

Figure 8 validates Trend 3 for semi-supervised (AGF) multiresolution classifier and generic supervised classifier (SVM). For both vehicles, semi-supervised multiresolution classifier with adaptive graph filtering dramatically improves classification accuracy across all labeling ratios.

2) Ability of Adaptive Graph Filtering to Handle Mislabeled Signals: In real-world problems, some of the labeled signals could be unreliable for different reasons, for example, negligence or uncertainty. As mentioned in Section IV-A, one of the advantages of using adaptive graph filtering is to provide robustness to mislabeling. To validate that, we randomly mislabel a fraction of labeled signals, feed them into the classifiers together with correctly labeled signals,

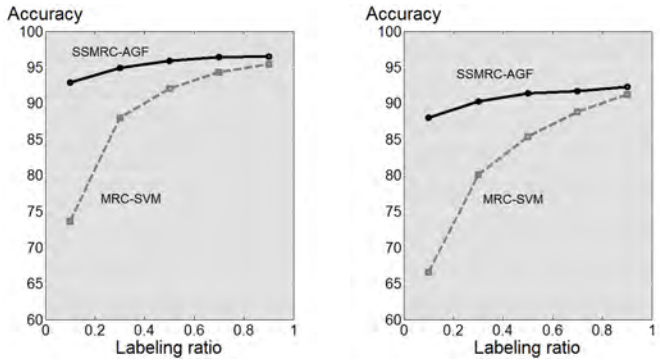

(a) Vehicle 1. (b) Vehicle 2.

Fig. 7: Trend 2: Semi-supervised learning improves classification accuracy. Comparison of multiresolution classifiers.

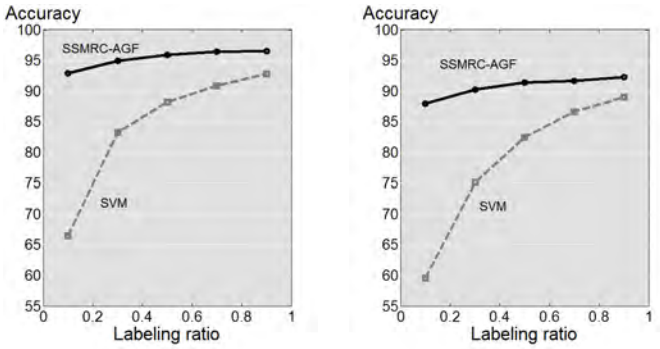

(a) Vehicle 1. (b) Vehicle 2.

Fig. 8: Trend 3: Multiresolution framework with semi-supervised learning improves classification accuracy.

V	S	Generic			SSMRC		
		DF	HF	AGF	DF	HF	AGF
1	1	83.88	85.15	88.53	95.92	96.60	**99.39**
	2	87.15	87.76	90.84	97.72	98.09	**99.79**
	3	87.28	88.77	91.03	97.47	97.71	**99.40**
	4	86.97	88.99	93.45	96.93	96.99	**99.52**
	5	74.52	73.99	77.72	94.41	95.28	**95.74**
	6	83.05	84.08	87.54	93.21	93.94	**95.75**
	7	63.99	63.45	66.08	78.37	79.21	**80.06**
	8	66.92	67.15	68.14	78.38	79.59	**80.44**
2	1	77.38	79.28	81.88	83.53	83.39	**85.68**
	2	66.78	67.65	69.98	77.97	78.32	**79.86**
	3	82.09	83.80	86.72	94.13	93.61	**95.07**
	4	79.99	81.00	84.56	86.67	86.61	**88.31**
	5	79.42	80.40	83.20	87.39	87.61	**89.01**
	6	80.40	80.97	83.34	84.25	83.43	**84.69**
	7	81.96	82.34	86.03	88.43	89.01	**89.85**
	8	79.59	81.47	83.31	92.31	92.94	**94.40**

TABLE IV: Robustness to mislabeled signals: accuracy comparison of Vehicles (V) 1 and 2, with Speeds (S) $1, 2, \ldots, 8$, with labeling ratio of 20% and misslabeling ratio of 15.38%.

V	S	Generic			SSMRC		
		DF	HF	AGF	DF	HF	AGF
1	1	72.55	73.41	81.99	87.16	87.43	**97.20**
	2	75.84	76.18	85.16	89.37	89.96	**98.14**
	3	79.36	79.40	87.17	90.33	90.73	**97.44**
	4	77.47	77.03	89.15	88.50	88.83	**97.90**
	5	66.32	62.75	72.87	86.78	85.67	**90.08**
	6	74.44	72.98	81.92	86.45	85.35	**90.98**
	7	57.70	55.89	62.02	72.43	72.69	**76.06**
	8	60.14	59.12	63.88	71.29	71.64	**75.26**
2	1	66.76	68.00	76.29	74.46	74.50	**81.90**
	2	57.62	57.50	63.41	71.08	70.85	**75.94**
	3	73.54	73.69	82.47	87.25	87.18	**93.19**
	4	71.27	70.05	79.78	78.49	78.70	**84.78**
	5	70.76	69.87	78.66	79.95	79.67	**85.80**
	6	72.39	71.84	80.41	77.18	77.13	**82.55**
	7	74.39	73.82	82.53	81.45	82.01	**87.04**
	8	72.88	71.12	80.04	85.82	86.29	**92.19**

TABLE V: Robustness to mislabeled signals: accuracy comparison of Vehicles (V) 1 and 2, with Speeds (S) $1, 2, \ldots, 8$, with labeling ratio of 20% and misslabeling ratio of 33.33%.

and compare the fault tolerances of the three semi-supervised classifiers. Tables IV and V show results where 20% of signals are labeled, with 15.38% and 33.33% of these labeled signals mislabeled, respectively. The three trends from before still hold: multiresolution framework on its own, semi-supervised learning on its own, and the two together, all improve classification accuracy. Among the semi-supervised multiresolution classifiers, adaptive graph filtering performs the best in each case. Moreover, as the ratio of mislabeled signals increases from 15.38% to 33.33%, the performance of adaptive graph filtering is relatively unaffected, while the performance of the other two semi-supervised classifiers, diffusion functions and harmonic functions, decreases dramatically. We thus conclude that the semi-supervised multiresolution classification using adaptive graph filtering is robust to mislabeled signals.

3) Ability of Adaptive Graph Filtering to Handle Unseen Signals: Finally, to validate the claim from Section IV-B that using regression allows us to handle unseen signals, we keep a portion of signals as unseen signals. For each vehicle and speed, we have 13 damage scenarios with 30 signals for each for a total of 390 available signals. We assign $M = 65$ of these to be unseen, leaving us with $N = 325$ total signals, out of which we label $L = 13$. We thus have 4% $(13/325)$ labeled and 20% $(65/325)$ unseen signals. Table VI shows results for both the generic classifier and semi-supervised multiresolution classifiers with adaptive graph filtering. In each case, we compare accuracies of $U = 312$ unlabeled signals as well as of $M = 65$ unseen signals. Although the unseen signals never appear in the filtering stage, classification accuracy is close to that of unlabeled signals. If the signal set is sufficiently large, the adaptive graph filter learns the distribution of signals from both labeled and unlabeled signals, which it then uses for the unseen signals.

VI. Conclusions and Future Work

We presented a novel classification framework combining multiresolution classification with semi-supervised learning; adaptive graph filtering for semi-supervised classification that allows for classifying unlabeled as well as unseen signals and for correcting mislabeled signals; and solution to indirect bridge structural health monitoring.

V	S	Generic-AGF		SSMRC-AGF	
		Unlabeled	Unseen	Unlabeled	Unseen
1	1	76.79	74.21	99.19	96.92
	2	80.34	74.72	99.26	96.77
	3	83.84	79.18	97.38	94.41
	4	85.67	84.51	99.70	98.77
	5	67.44	73.85	82.81	80.92
	6	75.20	74.97	85.35	80.97
	7	62.02	62.72	72.97	71.13
	8	57.97	58.05	67.39	66.97
2	1	75.63	75.38	81.99	82.41
	2	56.77	57.64	76.05	76.41
	3	77.13	77.08	93.25	92.00
	4	71.57	72.62	80.49	81.13
	5	69.31	68.77	83.73	82.62
	6	73.27	74.05	78.13	78.46
	7	77.98	78.10	83.60	83.59
	8	74.69	77.64	86.82	87.64

TABLE VI: Robustness to unseen signals: accuracy comparison of Vehicles (V) 1 and 2, with Speeds (S) $1, 2, \ldots, 8$, labeling ratio of 4% and ratio of unseen signals of 20%.

The proposed framework builds upon supervised multiresolution classification, which extracts hidden features in localized time-frequency subbands, and semi-supervised learning, which uses both labeled and unlabeled signals. We link the two via a novel weighting algorithm that combines information from all the subbands of all the signals to make a global decision in a semi-supervised fashion. We propose a novel semi-supervised classifier, adaptive graph filter, also, the first real application of signal processing on graphs. We further connect it to diffusion maps and diffusion wavelets and show that it performs multiresolution classification on graphs.

We validate the proposed framework on the task of indirect bridge structural health monitoring and show that: (1) multiresolution framework on its own, (2) semi-supervised learning on its own, and (3) the two together, all improve classification accuracy. Furthermore, we show that adaptive graph filtering has the ability to handle unlabeled, mislabeled as well as unseen signals.

Some near-future tasks are to use more features in each time-frequency subband, prune wavelet packet tree to achieve faster implementation and test the framework on real-world bridge-vehicle dynamic system.

VII. Reproducible research

We follow the principles of reproducible research. To that end, we created a reproducible research page available to reviewers [55]. Should the paper get accepted, that link will become live.

References

[1] S. Chen, F. Cerda, J. Guo, J. B. Harley, Q. Shi, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Multiresolution classification with semi-supervised learning for indirect bridge structure health monitoring," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Vancouver, Canada, May 2013.

[2] N. Saito, "Classification of geophysical acoustic waveforms using time-frequency atoms," *Proc. Am. Stat. Assoc. Stat. Comput.*, pp. 322–327, 1996.

[3] T. McConaghy, H. Leung, É. Bossé, and V. Varadan, "Classification of audio radar signals using radial basis function neural networks," *IEEE Trans. Instrum. Meas.*, vol. 52, no. 4, pp. 1771–1779, Dec. 2003.

[4] H. Sohn, C. R. Farrar, N. F. Hunter, and K. Worden, "Structural health monitoring using statistical pattern recognition techniques," *J. Dyn. Sys., Meas., Control*, vol. 123, no. 4, pp. 1519–1524, Dec. 2001.

[5] A. Chebira, Y. Barbotin, C. Jackson, T. E. Merryman, G. Srinivasa, R. F. Murphy, and J. Kovačević, "A multiresolution approach to automated classification of protein subcellular location images," *BMC Bioinform.*, vol. 8, no. 210, 2007.

[6] R. A. Kellogg, A. Chebira, A. Goyal, P. A. Cuadra, S. F. Zappe, J. S. Minden, and J. Kovačević, "Towards an image analysis toolbox for high-throughput Drosophila embryo RNAi screens," in *Proc. IEEE Int. Symp. Biomed. Imag.*, Arlington, VA, Apr. 2007, pp. 288–291.

[7] A. Chebira, J. A. Ozolek, C. A. Castro, W. G. Jenkinson, M. Gore, R. Bhagavatula, I. Khaimovich, S. E. Ormon, C. S. Navara, M. Sukhwani, K. E. Orwig, A. Ben-Yehudah, G. Schatten, G. K. Rohde, and J. Kovačević, "Multiresolution identification of germ layer components in teratomas derived from human and nonhuman primate embryonic stem cells," in *Proc. IEEE Int. Symp. Biomed. Imag.*, Paris, France, May 2008, pp. 979–982.

[8] T. Chen, W. Yin, X. Sean, Z. Dorin, C. Thomas, and S. Huang, "Total variation models for variable lighting face recognition and uneven background correction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, pp. 1519–1524, May 2006.

[9] X. Zhu, "Semi-supervised learning literature survey," Tech. Rep. 1530, Univ. Wisconsin-Madison, 2005.

[10] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, 2013.

[11] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, pp. 83–98, 2013.

[12] R. R. Coifman and S. Lafon, "Diffusion maps," *Appl. Comput. Harmon. Anal.*, pp. 5–30, 2006.

[13] R. R. Coifman and M. Maggioni, "Diffusion wavelets," *Appl. Comput. Harmon. Anal.*, pp. 53–94, July 2006.

[14] C.W. Lin and Y.B.Yang, "Use of a passing vehicle to scan the fundamental bridge frequencies: An experimental verification," *Engineering Structures*, vol. 27, pp. 1865–1878, 2005.

[15] F. Cerda, J. Garrett, J. Bielak, R. Bhagavatula, and J. Kovačević, "Exploring indirect vehicle-bridge interaction for bridge SHM," in *Proc. Int. Conf. Bridge Maint., Safety Manag.*, Philadelphia, PA, July 2010, pp. 696–702.

[16] F. Cerda, J. Garrett, J. Bielak, P. Rizzo, J. A. Barrera, Z. Zhang, S. Chen, M. T. McCann, and J. Kovačević, "Indirect structural health monitoring in bridges: scale experiments," in *Proc. Int. Conf. Bridge Maint., Safety Manag.*, Lago di Como, Italy, July 2012, pp. 346–353.

[17] A. Chebira, L. P. Coelho, A. Sandryhaila, S. Lin, W. G. Jenkinson, J. MacSleyne, C. Hoffman, P. Cuadra, C. Jackson, M. Püschel, and J. Kovačević, "An adaptive multiresolution approach to fingerprint recognition," in *Proc. IEEE Int. Conf. Image Process.*, San Antonio, TX, Sept. 2007, vol. 1, pp. 457 – 460.

[18] Y. B. Yang, C. W. Lina, and J. D. Yau, "Extracting bridge frequencies from the dynamic response of a passing vehicle," *J. Sound Vibrat.*, pp. 471–493, May 2004.

[19] F. Cerda, S. Chen, J. Bielak, J. H. Garrett, P. Rizzo, and J. Kovačević, "Indirect structural health monitoring of a simplified laboratory-scale bridge model," *Int. J. Smart Struct. Syst.*, 2012, Submitted.

[20] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, John Wiley & Sons, Englewood Cliffs, NJ, 2001.

[21] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*, Signal Processing. Prentice Hall, Englewood Cliffs, NJ, 1995, http://waveletsandsubbandcoding.org/.

[22] M. Vetterli, J. Kovačević, and V. K. Goyal, *Foundations of Signal Processing*, Cambridge Univ. Press, 2013.

[23] J. Kovačević, V. K. Goyal, and M. Vetterli, *Fourier and Wavelet Signal Processing*, Cambridge Univ. Press, 2014, http://www.fourierandwavelets.org/.

[24] J. Kovačević and A. Chebira, "Life beyond bases: The advent of frames (Part I)," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 86–104, July 2007.

[25] J. Kovačević and A. Chebira, "Life beyond bases: The advent of frames (Part II)," *IEEE Signal Process. Mag.*, vol. 24, no. 5, pp. 115–125, Sept. 2007.

[26] J. Kovačević and A. Chebira, *An Introduction to Frames*, Found. Trends Signal Process. Now Publishers, 2008.
[27] R. R. Coifman, Y. Meyer, S. Quake, and M. V. Wickerhauser, "Signal processing and compression with wavelet packets," Tech. Rep., Yale Univ., 1991.
[28] R. M. Haralick, "Statistical and structural approaches to texture," *Proc. IEEE*, vol. 67, pp. 786–804, 1979.
[29] D. Gabor, "Theory of communication," *J. IEE*, vol. 93, pp. 429–457, 1946.
[30] J.P. Jones and L.A. Palmer, "An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex," *J. Neuropsychol.*, vol. 58, pp. 1233–1258, 1987.
[31] J.G. Daugman, "Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, pp. 1169–1179, 1988.
[32] A.K. Jain and F. Farrokhnia, "Unsupervised texture segmentation using Gabor filters," *Pattern Recogn.*, vol. 24, pp. 1167–1186, 1991.
[33] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, pp. 971–987, 2002.
[34] T. Ahonen, A. Hadid, and M. Pietikäinen, "Face description with local binary patterns: Application to face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, pp. 2037–2041, 2006.
[35] G. Zhao and M. Pietikainen, "Dynamic texture recognition using local binary patterns with an application to facial expressions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, pp. 915–928, 2007.
[36] R. Bhagavatula, M. C. Fickus, J. W. Kelly, C. Guo, J. A. Ozolek, C. A. Castro, and J. Kovačević, "Automatic identification and delineation of germ layer components in H&E stained images of teratomas derived from human and nonhuman primate embryonic stem cells," in *Proc. IEEE Int. Symp. Biomed. Imag.*, Rotterdam, The Netherlands, Apr. 2010, pp. 1041–1044.
[37] A. Kuruvilla, J. Li, P. Hennings Yeomans, P. Quelhas, N. Shaikh, A. Hoberman, and J. Kovačević, "Otitis media vocabulary and grammar," in *Proc. IEEE Int. Conf. Image Process.*, Orlando, FL, Sept. 2012, pp. 2845–2848.
[38] M. T. McCann, R. Bhagavatula, M. C. Fickus, J. A. Ozolek, and J. Kovačević, "Automated colitis detection from endoscopic biopsies as a tissue screening tool in diagnostic pathology," in *Proc. IEEE Int. Conf. Image Process.*, Orlando, FL, Sept. 2012, pp. 2809–2812.
[39] A. Kuruvilla, N. Shaikh, A. Hoberman, and J. Kovačević, "Automated diagnosis of otitis media: A vocabulary and grammar," *Int. J. Biomed. Imag.*, 2013, Submitted.
[40] S. K. Narang and A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *IEEE Trans. Signal Process.*, vol. 60, pp. 2786–2799, 2012.
[41] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Appl. Comput. Harmon. Anal.*, vol. 30, pp. 129–150, 2011.
[42] K. Nigam, A. K. McCallum, S. Thrun, and T. Mitchell, "Text classification from labeled and unlabeled documents using EM," *Mach. Learn.*, vol. 39, no. (2/3), pp. 1–32, 2000.
[43] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proc. Conf. Computational Learning Theory*, Madison, WI, 1998.
[44] X. Zhu, *Semi-supervised learning with graphs*, Ph.D. thesis, Carnegie Mellon Univ., 2005.
[45] A. D. Szlam, R. R. Coifman, and M. Maggioni, "A general framework for adaptive regularization based on diffusion processes," *J. Mach. Learn. Res.*, , no. 9, pp. 1711–1739, Aug. 2008.
[46] X. Zhu and Z. Ghahramani, "Learning from labeled and unlabeled data with label propagation," Tech. Rep., Carnegie Mellon Univ., 2002.
[47] X. Zhu, J. Lafferty, and Z. Ghahramani, "Combining active learning and semi-supervised learning using Gaussian fields and harmonic functions," in *Proc. Int. Conf. Mach. Learn. Workshop on Continuum from Labeled to Unlabeled Data in Mach. Learn. Data Mining*, Washington, DC, 2003, pp. 58–65.
[48] F. R. K. Chung, *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*, Am. Math. Soc., 1996.
[49] L. Wasserman, *All of Nonparametric Statistics*, Springer Texts in Statistics. Springer, 2005.
[50] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning; Data Mining, Inference, and Prediction*, Springer-Verlag, 2001.
[51] L. Györfi, M. Kohler, A. Krzyzak, and H. Walk, *A Distribution-Free Theory of Nonparametric Regression*, Springer Texts in Statistics. Springer, 2002.
[52] I. Daubechies, "Orthonormal bases of compactly supported wavelets II. Variations on a theme," *SIAM J. Math. Anal.*, vol. 24, no. 2, pp. 499–519, Mar. 1993.
[53] C. M. Bishop, *Pattern Recognition and Machine Learning*, Information Science and Statistics. Springer, 2006.
[54] A. Singhal, "Modern information retrieval: A brief overview," *Bulletin IEEE Computer Soc. Tech. Committee Data Eng.*, vol. 24, no. 4, pp. 35–42, 2001.
[55] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring," 2013, http://jelena.ece.cmu.edu/repository/rr/13_ChenCRBGK/13_ChenCRBGK.html.

SIGNAL INPAINTING ON GRAPHS VIA TOTAL VARIATION MINIMIZATION

Siheng Chen[1,2], Aliaksei Sandryhaila[1], George Lederman[4], Zihao Wang[4], José M. F. Moura[1], Piervincenzo Rizzo[5], Jacobo Bielak[4], James H. Garrett[4] and Jelena Kovačević[3,1,2]

[1]Dept. of ECE, [2]Center for Bioimage Informatics, [3]Dept. of BME, [4]Dept. of CEE,
Carnegie Mellon University,
[5]Dept. of CEE, University of Pittsburgh
Pittsburgh, PA, USA

ABSTRACT

We propose a novel recovery algorithm for signals with complex, irregular structure that is commonly represented by graphs. Our approach is a generalization of the signal inpainting technique from classical signal processing. We formulate corresponding minimization problems and demonstrate that in many cases they have closed-form solutions. We discuss a relation of the proposed approach to regression, provide an upper bound on the error for our algorithm and compare the proposed technique with other existing algorithms on real-world datasets.

Index Terms— Signal processing on graphs, signal inpainting, total variation, semi-supervised learning.

1. INTRODUCTION

The problem of collecting, processing and analyzing data obtained from or represented by networks has been receiving a constantly increasing interest due to the abundance of such data in various research fields. An integral part of solving this problem is the development of new models and techniques that can be applied to datasets with complex irregular structures.

Recently, a theoretical framework called *signal processing on graphs* has emerged as a new approach to analyze signals with irregular structure [1, 2, 3, 4]. Its key idea is to represent the structure of a signal with a graph by associating signal coefficients with graph nodes and analyzing graph signals by using appropriately defined signal processing techniques, such as Fourier transform, filtering, and wavelets.

In this paper, we study the problem of signal recovery, that is, reconstruction or estimation of signal coefficients that are missing, unmeasurable, or corrupted by noise, a task often referred to as *signal inpainting* [5, 6, 7].

Previous work on this topic has primarily come when missing signal coefficients are unknown labels that need to be learned [8, 9]. Existing approaches are often based on the graph Laplacian operator and take roots in spectral graph theory [10]. From the perspective of signal processing, they seek to minimize a *total variation* function based on the graph Laplacian, which measures how much signal coefficients differ from other relevant coefficients. A major limitation of the graph Laplacian based method is its restriction to undirected graphs with real, non-negative edge weights.

We propose a novel approach to graph signal inpainting. It seeks to minimizes a total variation function that is based on a *graph shift*, which is a fundamental signal processing concept that can be defined for signals represented by any graph [1, 2]. As a result, our proposed technique is applicable to a much broader class of graphs. In this paper, we offer two formulations of the signal inpainting on graphs as minimization problems and demonstrate that in many cases these problems have exact closed-form solutions. We also discuss a connection of our approach to regression on graphs, derive an upper bound on the signal inpainting error produced by our technique, and demonstrate on real-world datasets that our approach can lead to better signal recovery and signal classification than other existing approaches.

2. DISCRETE SIGNAL PROCESSING ON GRAPHS

In this section, we briefly review relevant concepts of discrete signal processing on graphs; a thorough introduction can be found in [1, 2]. Discrete signal processing on graphs is a theoretical framework that generalizes classical discrete signal processing from regular domains, such as lines and rectangular lattices, to irregular structures that are commonly described by graphs.

Consider a graph $G = (\mathcal{V}, \mathrm{A})$, where $\mathcal{V} = \{v_n\}_{n=1}^N$ is the set of *nodes* and $\mathrm{A} \in \mathbb{C}^{N \times N}$ is a *graph shift*, that is, an elementary filtering operation that replaces a signal coefficient at a node with a weighted linear combination of coefficients at its neighboring nodes. The weights $\mathrm{A}_{i,j}$ characterize the relation between the ith node and the jth node. For example, they can quantify similarities and dependencies between nodes, or indicate communication patterns in networks. The graph shift operation is written as

$$\mathbf{s} \to \mathrm{A}\,\mathbf{s}.$$

This material is based upon work supported by the National Science Foundation under Grant no. 1130616 and 1017278, and by a National Science Foundation Graduate Research Fellowship under Grant No. 0946825. This research is also supported by a University Transportation Center grant (DTRT12-G-UTC11) from the US Department of Transportation and CMU Carnegie Institute of Technology Infrastructure Award.

where a *graph signal* $\mathbf{s}$ is defined as a mapping

$$\mathbf{s} \,:\, \mathcal{V} \rightarrow \mathbb{C}^n,$$

that assigns a signal coefficient $s_n \in \mathbb{C}$ to the graph node v_n.

A concept often used in signal processing is that of smoothness. Smoothness of graph signals is expressed by a *graph total variation* function

$$\mathrm{TV}_{\mathrm{A}}(\mathbf{s}) = \left|\left|\mathbf{s} - \frac{1}{|\lambda_{\max}|}\,\mathrm{A}\,\mathbf{s}\right|\right|_2^2, \tag{1}$$

where $\lambda_{max}(\mathrm{A})$ denote the eigenvalue of A with the largest magnitude.[1]

3. SIGNAL INPAINTING ON GRAPHS

Signal inpainting is a process of recovering missing or corrupted signal coefficients from a known part of the signal. In discrete-time signal processing, a signal, such as a time series or a digital image, is typically assumed to be smooth, and the missing part is recovered through regularization or lowpass filtering [5, 6, 7].

3.1. Problem formulation

The graph total variation (1) allows us to generalize the inpainting process from lines and rectangular lattices to arbitrary graphs. We work with graph signals of the form

$$\mathbf{s} = \begin{bmatrix}\mathbf{s}_{\mathcal{M}} \\ \mathbf{s}_{\mathcal{U}}\end{bmatrix}, \tag{2}$$

where $\mathbf{s}_{\mathcal{M}} \in \mathbb{C}^M$ is the known part of the signal and $\mathbf{s}_{\mathcal{U}} \in \mathbb{C}^{N-M}$ is the unknown part (without loss of generality, we assume that the known coefficients $s_1, \ldots, s_M$ correspond to the first M graph nodes $v_1, \ldots, v_M$; this arrangement can always be achieved by reordering nodes). Assuming that (2) is a smooth signal, that is, its variation is small, we recover the missing part $\mathbf{s}_{\mathcal{U}}$ by solving the following minimization problem:

$$\begin{aligned} \mathbf{s}^* &= \mathbf{argmin}\ \mathrm{TV}_{\mathrm{A}}(\widehat{\mathbf{s}}), && \text{(3a)} \\ \text{subject to} &\quad ||\widehat{\mathbf{s}}_{\mathcal{M}} - \mathbf{s}_{\mathcal{M}}||_2^2 \,\leq\, \epsilon^2. && \text{(3b)} \end{aligned}$$

The condition (3b) controls how well the known part of the signal is preserved.

3.2. Graph total variation regularization

Alternatively, the graph signal inpainting (3) can be formulated as an unconstrained problem,

$$\mathbf{s}^* = \mathbf{argmin}\, ||\widehat{\mathbf{s}}_{\mathcal{M}} - \mathbf{s}_{\mathcal{M}}||_2^2 + \lambda\,\mathrm{TV}_{\mathrm{A}}(\widehat{\mathbf{s}}), \tag{4}$$

[1] We omit the normalization factor $1/||\mathbf{s}||_2^2$ in (1) that was used in the original definition in [2], since signal inpainting algorithms we consider in this paper are not affected by this factor.

where the tuning parameter λ controls the trade-off between two parts of the objective function. Equation (4) is called as *graph total variation regularization* (GTVR). Small values of λ emphasize fitting the estimates to the known measurements more than the smoothness on the graph; large values of λ lead to smoother solutions of (4).

Note that (4) is a convex quadratic problem and has a closed-form solution. For many values of ϵ in (3) we can find a corresponding value of λ for which (4) yields an equivalent solution.

We now derive a closed-form solution to (4). For notational simplicity, assume that matrix A has been normalized to satisfy $\lambda_{max}(\mathrm{A}) = 1$, so we can write the variation function (1) as

$$\begin{aligned} \mathrm{TV}_{\mathrm{A}}(\mathbf{s}) &= ||\mathbf{s} - \mathrm{A}\,\mathbf{s}||_2^2. \\ &= \mathbf{s}^H(\mathrm{I}-\mathrm{A})^H(\mathrm{I}-\mathrm{A})\mathbf{s}, && \text{(5)} \end{aligned}$$

where I is the identity matrix. The objective function in (4) is thus a linear combination of two quadratic functions of $\widehat{\mathbf{s}}$. The derivative of the objective function in (4) is:

$$\begin{aligned} & \frac{\partial}{\partial \widehat{\mathbf{s}}}\left(||\widehat{\mathbf{s}}_{\mathcal{M}} - \mathbf{s}_{\mathcal{M}}||_2^2 + \lambda\,\mathrm{TV}_{\mathrm{A}}(\widehat{\mathbf{s}})\right) \\ = \;& \frac{\partial}{\partial \widehat{\mathbf{s}}}\left((\widehat{\mathbf{s}}-\mathbf{s})^H \begin{bmatrix}\mathrm{I}_M & \mathbf{0} \\ \mathbf{0} & \mathbf{0}\end{bmatrix}(\widehat{\mathbf{s}}-\mathbf{s}) + \lambda\widehat{\mathbf{s}}^H(\mathrm{I}-\mathrm{A})^H(\mathrm{I}-\mathrm{A})\widehat{\mathbf{s}}\right) \\ = \;& 2\begin{bmatrix}\mathrm{I}_M & \mathbf{0} \\ \mathbf{0} & \mathbf{0}\end{bmatrix}(\widehat{\mathbf{s}}-\mathbf{s}) + 2\lambda(\mathrm{I}-\mathrm{A})^H(\mathrm{I}-\mathrm{A})\widehat{\mathbf{s}}. && \text{(6)} \end{aligned}$$

By setting (6) to zero, we then obtain the closed-form solution as follows:

$$\mathbf{s}^* = \left(\begin{bmatrix}\mathrm{I}_M & \mathbf{0} \\ \mathbf{0} & \mathbf{0}\end{bmatrix} + \lambda(\mathrm{I}-\mathrm{A})^H(\mathrm{I}-\mathrm{A})\right)^{-1}\begin{bmatrix}\mathbf{s}_{\mathcal{M}} \\ \mathbf{0}\end{bmatrix}. \tag{7}$$

3.3. Graph total variation minimization

When the known part of the signal needs to be preserved intact, we must solve (3) directly for $\epsilon = 0$,

$$\begin{aligned} \mathbf{s}^* &= \mathbf{argmin}\ \mathrm{TV}_{\mathrm{A}}(\widehat{\mathbf{s}}), && \text{(8a)} \\ \text{subject to} &\quad \widehat{\mathbf{s}}_{\mathcal{M}} = \mathbf{s}_{\mathcal{M}}. && \text{(8b)} \end{aligned}$$

Equation (8) is called as *graph total variation minimization* (GTVM). Denote by $\widetilde{\mathrm{A}} = (\mathrm{I}-\mathrm{A})^H(\mathrm{I}-\mathrm{A})$. By writing $\widetilde{\mathrm{A}}$ in a block form as

$$\widetilde{\mathrm{A}} = \begin{bmatrix}\widetilde{\mathrm{A}}_{\mathcal{M}\mathcal{M}} & \widetilde{\mathrm{A}}_{\mathcal{M}\mathcal{U}} \\ \widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{M}} & \widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{U}}\end{bmatrix}, \tag{9}$$

we can rewrite the objective function in (8) as

$$\begin{aligned} \mathrm{TV}_{\mathrm{A}}(\widehat{\mathbf{s}}) &= \widehat{\mathbf{s}}^H\widetilde{\mathrm{A}}\widehat{\mathbf{s}} \\ &= \begin{bmatrix}\widehat{\mathbf{s}}_{\mathcal{M}}^H & \widehat{\mathbf{s}}_{\mathcal{U}}^H\end{bmatrix}\begin{bmatrix}\widetilde{\mathrm{A}}_{\mathcal{M}\mathcal{M}} & \widetilde{\mathrm{A}}_{\mathcal{M}\mathcal{U}} \\ \widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{M}} & \widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{U}}\end{bmatrix}\begin{bmatrix}\widehat{\mathbf{s}}_{\mathcal{M}}^H & \widehat{\mathbf{s}}_{\mathcal{U}}^H\end{bmatrix} \\ &= \widehat{\mathbf{s}}_{\mathcal{M}}^H\widetilde{\mathrm{A}}_{\mathcal{M}\mathcal{M}}\widehat{\mathbf{s}}_{\mathcal{M}} + \widehat{\mathbf{s}}_{\mathcal{U}}^H\widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{M}}\widehat{\mathbf{s}}_{\mathcal{M}} \\ &\quad + \widehat{\mathbf{s}}_{\mathcal{M}}^H\widetilde{\mathrm{A}}_{\mathcal{M}\mathcal{U}}\widehat{\mathbf{s}}_{\mathcal{U}} + \widehat{\mathbf{s}}_{\mathcal{U}}^H\widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{U}}\widehat{\mathbf{s}}_{\mathcal{U}}. && \text{(10)} \end{aligned}$$

Since $\widehat{\mathbf{s}}_{\mathcal{M}} = \mathbf{s}_{\mathcal{M}}$, as specified by (8b), the minimum of the objective function (10) is found by setting its derivative to zero, which yields the closed-form solution

$$\widehat{\mathbf{s}}_{\mathcal{U}} = -\widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{U}}^{-1}\widetilde{\mathrm{A}}_{\mathcal{U}\mathcal{M}}\mathbf{s}_{\mathcal{M}}. \tag{11}$$

4. DISCUSSION

We now discuss several properties of our proposed signal inpainting algorithms.

4.1. Connection to Graph-based Regression

We can construct a graph to represent a dataset and a set of labels by associating each dataset element with a node and viewing labels as a signal on this graph. In this setting, graph signal inpainting becomes analogous to a *regression*, since it estimates a regression function that assigns a label to each node.

In general, the graph-based regression has the form [11, 9, 8]

$$\mathbf{s}^* \;=\; \mathbf{argmin}\; ||\widehat{\mathbf{s}}_{\mathcal{M}} - \mathbf{s}_{\mathcal{M}}||^2 + \lambda \widehat{\mathbf{s}}^H \,\mathrm{G}\,\widehat{\mathbf{s}}, \tag{12}$$

where G is a smoothing matrix that can be defined in different ways depending on the context. For instance, if G is the graph Laplacian matrix, then (12) solves Laplacian regularization [9]. If, in addition, $\widehat{\mathbf{s}}_{\mathcal{M}} = \mathbf{s}_{\mathcal{M}}$, then the minimization (12) is based on the Laplacian harmonic functions [8]. Since our approach (4) uses the total variation defined by the graph shift, we use $\mathrm{G} = (\mathrm{I}-\mathrm{A})^H(\mathrm{I}-\mathrm{A})$ in (12).

4.2. Error Analysis

We now derive an upper bound on the estimation error of graph signal inpainting algorithm (3).

Let $\mathbf{s}^0$ be the *true* graph signal that we are trying to estimate by signal inpainting. Assume that $\mathrm{TV}_{\mathrm{A}}(\mathbf{s}^0) = \eta^2$ and $\mathbf{s}^0$ satisfies (3b), so that $||\mathbf{s}^0_{\mathcal{M}} - \mathbf{s}_{\mathcal{M}}||_2^2 \leq \epsilon^2$. Also, write A in a block form similarly to (9).

Lemma 1. *The estimation error* $\mathbf{e} = \mathbf{s}^0 - \mathbf{s}^*$ *of the signal inpainting algorithm* (3) *is bounded by the unmeasured part of the signal as*

$$||\mathbf{e}||_2 \;\leq\; \frac{q}{2}||\mathbf{e}_{\mathcal{U}}||_2 + p|\epsilon| + |\eta|, \tag{13}$$

where

$$p = \left|\left|\begin{bmatrix}\mathrm{I}_{\mathcal{M}\mathcal{M}}+\mathrm{A}_{\mathcal{M}\mathcal{M}}\\ \mathrm{A}_{\mathcal{U}\mathcal{M}}\end{bmatrix}\right|\right|_2,\;\; q = \left|\left|\begin{bmatrix}\mathrm{A}_{\mathcal{M}\mathcal{U}}\\ \mathrm{I}_{\mathcal{U}\mathcal{U}}+\mathrm{A}_{\mathcal{U}\mathcal{U}}\end{bmatrix}\right|\right|_2,$$

and $||\cdot||_2$ *for matrices denotes the spectral norm.*

Proof. Since $||\mathbf{e}_{\mathcal{M}}||_2 = ||\mathbf{s}^0_{\mathcal{M}} - \mathbf{s}^*_{\mathcal{M}}||_2 \leq ||\mathbf{s}^0_{\mathcal{M}} - \mathbf{s}_{\mathcal{M}}||_2 + ||\mathbf{s}_{\mathcal{M}} - \mathbf{s}^*_{\mathcal{M}}||_2 = 2|\epsilon|$, we obtain

$$\begin{aligned}
||(\mathrm{I}+\mathrm{A})\mathbf{e}||_2 &= \left|\left|\begin{bmatrix}\mathrm{I}_{\mathcal{M}\mathcal{M}}+\mathrm{A}_{\mathcal{M}\mathcal{M}} & \mathrm{A}_{\mathcal{M}\mathcal{U}}\\ \mathrm{A}_{\mathcal{U}\mathcal{M}} & \mathrm{I}_{\mathcal{U}\mathcal{U}}+\mathrm{A}_{\mathcal{U}\mathcal{U}}\end{bmatrix}\begin{bmatrix}\mathbf{e}_{\mathcal{M}}\\ \mathbf{e}_{\mathcal{U}}\end{bmatrix}\right|\right|_2 \\
&\leq \left|\left|\begin{bmatrix}\mathrm{I}_{\mathcal{M}\mathcal{M}}+\mathrm{A}_{\mathcal{M}\mathcal{M}}\\ \mathrm{A}_{\mathcal{U}\mathcal{M}}\end{bmatrix}\right|\right|_2 \cdot ||\mathbf{e}_{\mathcal{M}}||_2 \\
&\quad + \left|\left|\begin{bmatrix}\mathrm{A}_{\mathcal{M}\mathcal{U}}\\ \mathrm{I}_{\mathcal{U}\mathcal{U}}+\mathrm{A}_{\mathcal{U}\mathcal{U}}\end{bmatrix}\right|\right|_2 \cdot ||\mathbf{e}_{\mathcal{U}}||_2 \\
&= p||\mathbf{e}_{\mathcal{M}}||_2 + q||\mathbf{e}_{\mathcal{U}}||_2 \\
&= 2p|\epsilon| + q||\mathbf{e}_{\mathcal{U}}||_2.
\end{aligned}$$

Since both $\mathbf{s}^0$ and $\mathbf{s}^*$ satisfy (3b), then $\mathrm{TV}_{\mathrm{A}}(\mathbf{s}^*) \leq \mathrm{TV}_{\mathrm{A}}(\mathbf{s}^0)$, and we obtain

$$\begin{aligned}
||(\mathrm{I}-\mathrm{A})\mathbf{e}||_2 &= ||(\mathrm{I}-\mathrm{A})(\mathbf{s}^0 - \mathbf{s}^*)||_2 \\
&\leq ||(\mathrm{I}-\mathrm{A})\mathbf{s}^0||_2 + ||(\mathrm{I}-\mathrm{A})\mathbf{s}^*||_2 \\
&\leq 2||(\mathrm{I}-\mathrm{A})\mathbf{s}^0||_2 \\
&\leq 2|\eta|.
\end{aligned}$$

Combining these inequalities, we get

$$\begin{aligned}
||\mathbf{e}||_2 &= \left|\left|\frac{1}{2}(\mathrm{I}+\mathrm{A}+\mathrm{I}-\mathrm{A})\mathbf{e}\right|\right|_2 \\
&\leq \frac{1}{2}(||(\mathrm{I}+\mathrm{A})\mathbf{e}||_2 + ||(\mathrm{I}-\mathrm{A})\mathbf{e}||_2) \\
&\leq \frac{1}{2}(2p|\epsilon| + q||\mathbf{e}_{\mathcal{U}}||_2 + 2|\eta|),
\end{aligned}$$

which yields (13). □

Theorem 1. *If* $q < 2$*, then the estimation error on the unknown part is bounded by*

$$||\mathbf{e}_{\mathcal{U}}||_2 \;\leq\; \frac{2p|\epsilon| + 2|\eta|}{2-q}. \tag{14}$$

Proof. Using Lemma 1, we have

$$||\mathbf{e}_{\mathcal{U}}||_2 \;\leq\; ||\mathbf{e}||_2 \;\leq\; p|\epsilon| + \frac{q}{2}||\mathbf{e}_{\mathcal{U}}||_2 + |\eta|. \tag{15}$$

By rearranging the terms in (15), we obtain (14). □

The condition $q < 2$ in Theorem 1 may not hold for some matrices. However, if A is symmetric, we have $q \leq ||\,\mathrm{I}+\mathrm{A}\,||_2 \leq ||\,\mathrm{I}\,||_2 + ||\,\mathrm{A}\,||_2 = 2$, since $||\,\mathrm{A}\,||_2 = 1$. Also, note that the upper bound is related to the smoothness of the true graph signal and the noise level of the measured part. A central assumption of any inpainting technique is that the true signal $\mathbf{s}^0$ is smooth. If this assumption does not hold, then the upper bound is large and useless. When the noise level of the measured part is smaller, the measurements from the known part are closer to the true values, which leads to a smaller estimation error.

Method	λ	Ratio of known labels 0.5%	1%	2%	5%	10%
GTVM		80.25	**94.76**	**95.18**	**95.28**	95.20
GTVR	**0.01**	80.16	94.68	95.09	95.18	95.13
	0.1	79.99	94.47	95.11	95.19	95.15
	1	75.64	92.62	94.99	95.26	95.25
	10	61.68	68.13	89.02	95.15	**95.43**
	100	61.59	57.26	54.72	73.62	92.28
HF		51.58	53.15	60.75	86.05	94.68
LapR	**0.01**	51.47	53.48	60.79	85.90	94.62
	0.1	50.64	54.40	61.14	85.22	94.55
	1	50.25	56.49	60.50	79.29	93.94
	10	50.07	49.84	54.39	62.83	80.69
	100	51.29	49.74	50.16	52.04	55.95
AGF		**84.81**	88.36	94.08	95.00	95.10

Table 1: Accuracy of political blog classification.

Method	λ	Ratio of known masses 0.5%	1%	2%	5%	10%
GTVM		20.18	9.67	5.29	3.93	3.54
GTVR	**0.01**	14.71	9.28	4.92	3.88	3.52
	0.1	12.89	8.97	4.95	3.92	3.49
	1	10.12	9.11	5.14	3.58	3.24
	10	**8.84**	10.49	4.12	3.08	2.84
	100	15.70	**7.13**	**3.80**	**3.00**	**2.80**
HF		38.66	19.44	5.29	3.31	2.9
LapR	**0.01**	38.75	19.49	5.30	3.32	2.9
	0.1	39.52	19.99	5.42	3.34	2.9
	1	46.29	24.78	6.83	3.65	2.95
	10	74.38	52.64	22.90	9.44	5.13
	100	96.32	85.43	66.84	47.86	32.1

Table 2: The mean square error for the bridge condition identification.

5. EXPERIMENTS

We now apply the proposed algorithm to the classification of online blogs and to the bridge condition identification for indirect bridge structural health monitoring. We compare the proposed algorithm with Laplacian regularization (LapR) and harmonic functions (HF) discussed in Section 4.1. In classification of online blogs, we also compare the proposed algorithm with adaptive graph filtering (AGF) [12, 13], which is a semi-supervised classifier that combines the decisions from multiple graph filters using a semi-supervised weighting function.

5.1. Classification of online blogs

We consider the problem of classifying $N = 1224$ online political blogs as either conservative or liberal [14]. We represent conservative labels as $+1$ and liberal ones as -1.

The blogs are represented by a graph in which nodes represent blogs, and directed graph edges correspond to hyperlink references between blogs. For a node v_n its outgoing edges have weights $1/\deg(v_n)$, where $\deg(v_n)$ is the out-degree of v_n (the number of outgoing edges). We randomly labeled 0.5%, 1%, 2%, 5% and 10% of blogs and applied the inpainting algorithms to estimate the labels for remaining nodes. Estimated labels were thresholded around zero, so that positive values were set to $+1$ and negative to -1.

Classification accuracies of GTVM, GTVR, HF, LapR and AGF averaged over 30 tests for each labeling ratio are shown in Table 1. For proper evaluation, values of λ ranging between 0.01 and 100 were used for GTVR and LapR. In most cases, GTVM provides the most accurate classification. Note that GTVR is less sensitive to the value of λ than LapR, and our proposed methods achieve significantly higher accuracy than LapR and HF for low labeling ratios.

5.2. Bridge condition identification

We next consider the bridge condition identification problem [15, 16]. Monitoring infrastructure has become a major research subject within the civil engineering community. On operational bridges, signals can be gathered frequently but visual inspections typically occur every two years, which leads to a low labeling ratio. To validate the feasibility of indirect bridge structural health monitoring, a lab-scale bridge-vehicle dynamic system was built. Accelerometers were installed on a vehicle that travels across the bridge; acceleration signals were then collected from those accelerometers [16]. To simulate different bridge conditions in a lab-scale bridge, masses with various weights were put on the bridge. We collected 30 acceleration signals for each of 31 mass levels, with an interval of 5 grams from 0 to 150 grams, to simulate different severity of damages, for a total of 930 acceleration signals. For more details, see [17].

The recording are represented by an 8-nearest neighbor graph, in which nodes represent recordings, and each node is connected to eight other nodes that represent the most similar recordings. The graph shift A are constructed as $\mathrm{A}_{i,j} = \mathrm{P}_{i,j} \,/\, \sum_i \mathrm{P}_{i,j}$, where

$$\mathrm{P}_{i,j} \;=\; \exp \frac{-N^2 ||\mathbf{x}_i - \mathbf{x}_j||_2}{\sum_{i,j}(||\mathbf{x}_i - \mathbf{x}_j||_2)},$$

and $\mathbf{x}_i$ is a vector representation of the ith recording. We randomly assigned known masses to 0.5%, 1%, 2%, 5% and 10% of recordings and applied the inpainting algorithms to estimate the masses for remaining nodes.

The mean square errors for estimated masses averaged over 30 tests for each labeling ratio are shown in Table 2. The proposed GTVR approach yields the lowest errors and is less sensitive than LapR to the change of the tuning parameter λ. Overall, our proposed method achieve noticeably smaller errors than Laplacian-based method for low labeling ratios.

6. CONCLUSION

We presented a new algorithm for signal inpainting on graphs that, unlike previous approaches, is applicable to arbitrary graphs. We formulated corresponding minimization problems and derived closed-form solutions, as well as calculated an upper bound on the resulting error. We identified the relation between our approach and previous methods via regression on graphs. Experiments on real-world datasets of online political blogs and indirect bridge structural health monitoring showed that the proposed algorithm outperforms the graph Laplacian based method.

7. REFERENCES

[1] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, 2013.

[2] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Trans. Signal Process.*, 2013, Submitted.

[3] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, pp. 83–98, 2013.

[4] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Appl. Comput. Harmon. Anal.*, vol. 30, pp. 129–150, 2011.

[5] L. I. Rudin, Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1–4, pp. 259–268, 1992.

[6] T. F. Chan, S. Osher, and J. Shen, "The digital TV filter and nonlinear denoising," *IEEE Trans. Image Proc.*, vol. 10, no. 2, pp. 231–241, 2001.

[7] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, no. 1-2, pp. 89–97, Jan. 2004.

[8] X. Zhu, Z. Ghahramani, and J. Lafferty, "Semi-supervised learning using gaussian fields and harmonic functions," in *Proc. ICML*, 2003, pp. 912–919.

[9] M. Belkin, P. Niyogi, and P. Sindhwani, "Manifold regularization: A geometric framework for learning from labeled and unlabeled examples.," *J. Machine Learn. Research*, vol. 7, pp. 2399–2434, 2006.

[10] F. R. K. Chung, *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*, Am. Math. Soc., 1996.

[11] D. Zhou and B. Scholkopf, "A regularization framework for learning from graph data," in *ICML Workshop Stat. Rel. Learn.*, 2004, pp. 132–137.

[12] S. Chen, A. Sandryhaila, J. M. F. Moura, and J. Kovačević, "Adaptive graph filtering: Multiresolution classification on graphs," in *Proc. IEEE Glob. Conf. Signal Information Process.*, Austin, TX, Dec. 2013, To appear.

[13] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring," *IEEE Trans. Signal Process.*, June 2013, Submitted.

[14] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 u.s. election: Divided they blog," in *Proc. LinkKDD*, 2005, pp. 36–43.

[15] F. Cerda, J. Garrett, J. Bielak, P. Rizzo, J. A. Barrera, Z. Zhang, S. Chen, M. T. McCann, and J. Kovačević, "Indirect structural health monitoring in bridges: scale experiments," in *Proc. Int. Conf. Bridge Maint., Safety Manag.*, Lago di Como, Italy, July 2012, pp. 346–353.

[16] F. Cerda, S. Chen, J. Bielak, J. H. Garrett, P. Rizzo, and J. Kovačević, "Indirect structural health monitoring of a simplified laboratory-scale bridge model," *Int. J. Smart Struct. Syst., sp. iss. Challenge on bridge health monitoring utilizing vehicle-induced vibrations*, 2013, To appear.

[17] G. Lederman, Z. Wang, J. Bielak, H. Noh, J. H. Garrett, S. Chen, J. Kovačević, F. Cerda, and P. Rizzo, "Damage quantification and localization algorithms for indirect shm of bridges," in *Proc. Int. Conf. Bridge Maint., Safety Manag.*, Shanghai, China, 2014, Submitted.

Damage quantification and localization algorithms for indirect SHM of bridges

G. Lederman[1], Z. Wang[1], J. Bielak, H. Noh & J. H. Garrett
Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA

S. Chen & J. Kovačević
Department of Biomedical Engineering & Center for Bioimage Informatics

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

F. Cerda
Departamento de Ingeniería Civil, Universidad de Concepción, Concepción, 4070409, Chile

P. Rizzo
Department of Civil and Environmental Engineering, University of Pittsburgh, PA 15216, USA

ABSTRACT: This paper presents algorithms for diagnosing the severity and location of damage in a laboratory bridge model. We use signal processing and machine learning approaches to analyze the vibration responses collected both directly from the bridge model and indirectly from a vehicle passing over the model. Features are selected using principal component analysis (PCA), and a regression is performed using the kernel regression method. Various "damage" severities and positions are simulated on a laboratory bridge model by placing additional mass on the bridge. We perform two experiments; one to measure our ability to detect damage severity (i.e. size of the mass), and a second to measure our ability to detect damage location (i.e. position of the mass). In the first experiment, we vary the magnitude of the mass while keeping its location constant. In the second experiment, we vary the location of the mass while keeping its magnitude constant. In both cases, we use a portion of our data to train the algorithm, and another portion to test its validity. We report the accuracy of correctly quantifying the nature of the mass from the test data as a mean square error (MSE).

1 INTRODUCTION

With aging infrastructure both in the United States and abroad, structural health monitoring (SHM) for civil applications has become a focus of intense research as a means to objectively determine the condition of a structure.

The American Society of Civil Engineers reports that of the 600,000 bridges within the United States, one in nine bridges is rated as structurally deficient (ASCE 2013). In addition, the recent collapse of the I-5 Bridge in Washington State and the earlier collapse of the I-35 Bridge in Minneapolis demonstrated the need for advanced technologies to monitor bridges. Researchers in the SHM community have already made significant contributions towards developing sensing systems and damage detection algorithms (Doebling et al. 1996, Chang 2011, Frangopol et al. 2010, Casciati and Giordano 2010)

The ultimate goal of SHM is to determine the remaining useful life of the structure. The state of the structure can be determined through a five-step process: (1) existence, (2) localization, (3) type, (4) severity and (5) prognosis of the damage (Rytter 1993).

The vast majority of monitoring systems for bridges require either wired or wireless sensors placed directly on the structure of interest. These techniques have shown some promising results but they require significant capital investment. This paper focuses on an indirect monitoring paradigm where the sensors are placed on a passing vehicle. This indirect approach is more economical; a fleet of vehicles could potentially monitor a large bridge inventory (Lin et al. 2005, Cerda et al. 2010).

Previous work on the indirect monitoring paradigm has examined determining the state of the structure between two (binary classification) or several cases (multiclass classification) (Cerda et al. 2013). Cerda et al simulated damage by adding a 'proxy damage' to a laboratory scale model. In one experiment, they varied the size and location of the proxy damage to measure their ability to determine (2) localization and (4) severity of the damage. They quantified the accuracy using multiclass classification with several discrete classes.

In this paper, we expand on the work of Cerda et al. by performing a regression on a large dataset of proxy damage locations and proxy damage severities, using the same laboratory model as Cerda et al. Using this regression we can determine the

[1] These authors contributed equally to this work.

state of the structure for an infinite number of mass locations and sizes within our training set. To build the regression, we apply principal component analysis to the acceleration signals, and train the kernel regression model by the collected data. This model determines the size and the location of the damage proxy using the MSE as the evaluation score.

2 EXPERIMENTAL SETUP AND PROTOCOL

2.1 *Experimental setup (Cerda et al. 2013; Wang et al. 2013)*

A general view of the laboratory model used in this project is shown in Figure 1, and schematic of the setup is shown in Figure 2. The model consists of a vehicle that is pulled across the rails by a cable system. The vehicle starts on 'Ramp 1', accelerates up to a constant speed, crosses the middle section, the "bridge," then decelerates on 'Ramp 2.' The vehicle has wired accelerometers so there is a cable rail above to ensure these wires do not interfere with its motion.

The "bridge" is an aluminum plate, 2438 mm (8 feet) long, with two angle beams acting as girders and two rails to guide the vehicle. The vehicle model, as shown in Figure 3, has an independent suspension system. Both the vehicle and the bridge were instrumented with commercial accelerometers. On the vehicle, two sensors are on the sprung portion of the vehicle ('front chassis sensor' and 'rear chassis sensor'), and two sensors are on the unsprung portion of the vehicle, rigidly attached to the wheels ('front wheel sensor' and 'rear wheel sensor'), as shown in Figure 3. One sensor was placed underneath the bridge deck at midspan ('bridge sensor').

Figure 1. The general view of the laboratory setup.

The motors governing the movement of the vehicle and the data-acquisition systems are both controlled by National Instrument's® PXI system running LabView®. By using a single system, we can spatially align the time series data from different runs using the vehicle's position. More details about the experimental setup can be found in (Cerda et al. 2013).

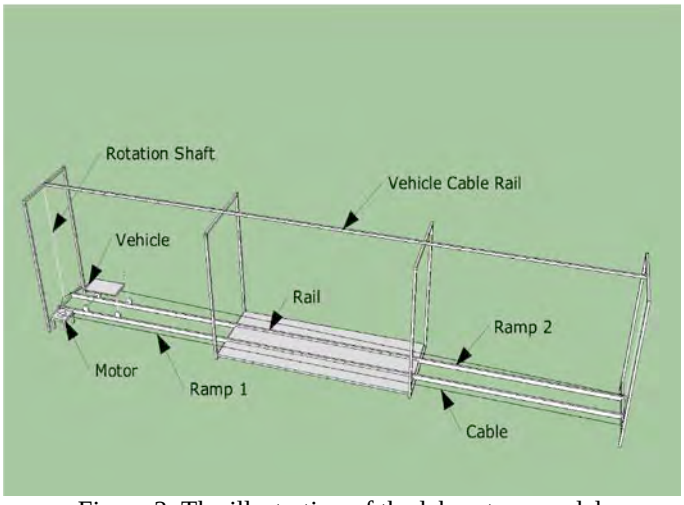

Figure 2. The illustration of the laboratory model.

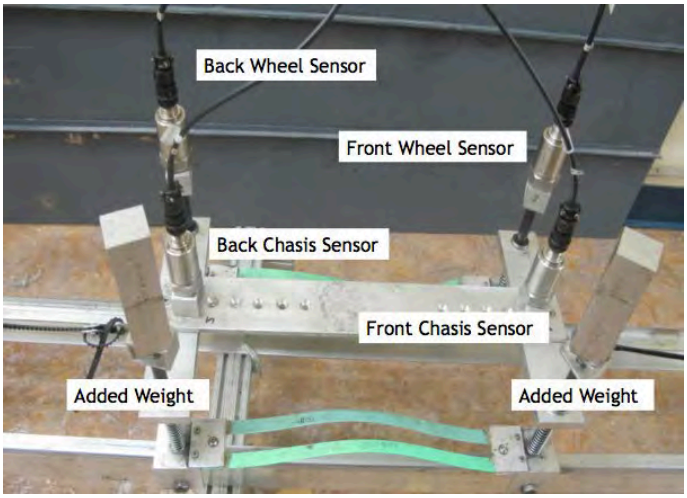

Figure 3. Details of the vehicle.

2.2 *Protocol*

In this experiment, the damage proxy is the presence of a mass on the deck. We assume that as the mass level changes gradually, the vibration characteristics will change accordingly. By mapping the relationship between changes in the acceleration signal to changes in the magnitude of the mass, we can determine the state of the bridge from an acceleration signal involving a change in mass size with our training range. The same assumption is applied to the change of the positions.

We assume a heavier mass means more severe damage as it is a more significant change from the baseline condition. In this paper, we use 31 mass levels, with an interval of 5 grams from 0 to 150 grams. We ran the experiments at 2 different speeds. The bridge itself weighs 15.5kg so the added mass

varies from 0%-1% of the mass the bridge. For these severity tests, we have 31 (mass) × 2 (speeds) × 30 (iterations) = 1860 (trials).

To investigate damage localization, a mass of 200 grams was placed at 30 locations, with an interval of 8 cm. We ran the experiment at 4 different speeds. The positions of the mass are shown in Figure 4. In total, for the localization tests, we ran 30 (locations) × 4 (speeds) × 30 (iterations) = 3600 (trials).

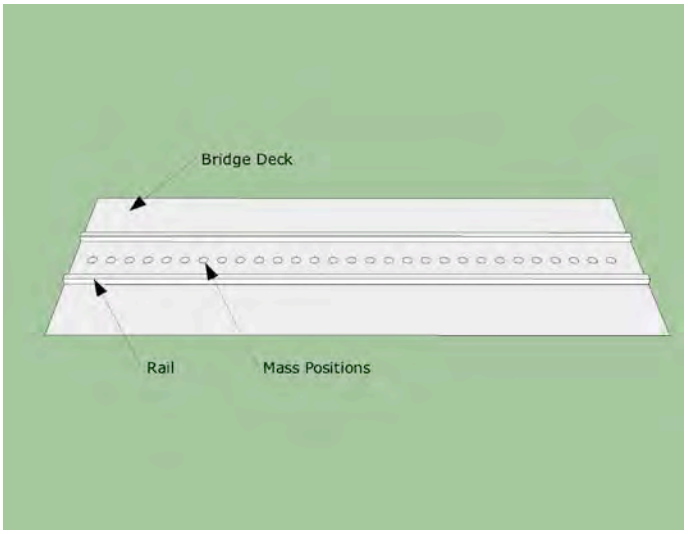

Figure 4. Illustration of mass positions on the deck.

3 REGRESSION FOR STRUCTURAL SCENARIOS

3.1 *The framework of the signal-processing system*

The goal of our signal processing approach for this experiment was to design a map to associate an acceleration signal with its corresponding bridge condition. In this experiment, we varied the bridge condition in small increments so that we could examine the evolution of the signal as the location or the severity of the proxy damage changed. By mapping this relationship, we were able to record a new acceleration signal of a previously unseen bridge condition, and predict that condition. There are two main challenges: the acceleration signals lie in a high-dimensional space, which is hard to visualize and further model; there is no closed-form formula to describe the relationship between the acceleration signals and the bridge conditions. We solve the first challenge by using PCA to reduce the dimensionality and solve the second one by using kernel regression to build a nonparametric regression model. The signal-processing system thus contains the dimensionality reduction block and the regression block shown in Figure 5.

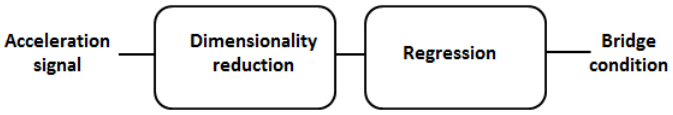

Figure 5. The system of regression for structural scenarios.

3.2 *Dimensionality reduction and visualization*

Each acceleration signal is sampled at 1667 Hz, and the vehicle takes roughly 2 seconds to cross the bridge (depending on its speed), it then contains over 3000 signal samples. The high dimensionality leads to the difficulty in visualizing and understanding the distribution of acceleration signals; it is also hard to perform further analysis because of the so-called *curse of dimensionality* (Duda et al. 2000). To solve this, we use PCA (Duda et al. 2000) to reduce the dimensionality. It finds an orthogonal linear transformation from the given dataset and transforms the signals into a new coordinate system such that the first coordinate captures the greatest variance, the second coordinate captures the second greatest variance, and so on. The algorithm to compute PCA is as follows.

$$\bar{x} = 1/m \sum_{i=1} x_i$$

PCA (extract top k eigenvectors as features):

Given data $X=\{x_1, \ldots, x_n\}$,

1. Calculate the mean of each column: $\bar{x} = (1/m) \sum_{i=1}^{m} x_i$.

2. Subtract the mean $\bar{x}$ from each vector x_i and get a new matrix **X**.

3. Calculate the covariance matrix Σ of X.

4. Calculate eigenvectors and eigenvalues of Σ and sort the eigenvectors ascending order based on the corresponding eigenvalues.

5. Select the first k principal components as features.

The implementation details of dimensionality reduction block are as follows: we first take the discrete Fourier transform of each acceleration signal and compute the magnitudes of their frequency spectrums; we then use PCA to analyze the magnitudes of all the acceleration signals. As an example, for visualization, we only preserve the first

three coordinates in Figure 6 to Figure 10. Each subplot is the feature spectrum extracted from the data collected from one position on the vehicle. The colors from red to blue indicate the increasing mass from 0 gram to 150 grams as can be seen on the colorbars. These figures show how the features of the acceleration signals change as the severity of the mass increases. We see a gradual change in the features as the size of the proxy damage increases; we can use these graphs to justify our mapping approach. For a given acceleration signal, we can plot its features in this space, and can deduce the condition of the bridge from the position of the features relative to known cases.

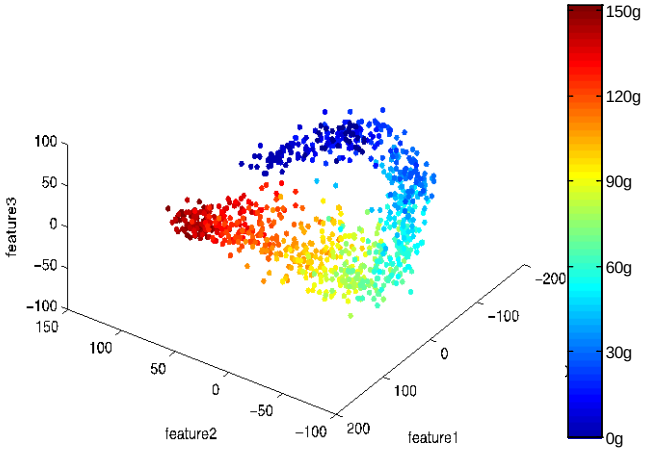

Figure 6. Visualization using first three principal components from bridge sensor accelerometer signals.

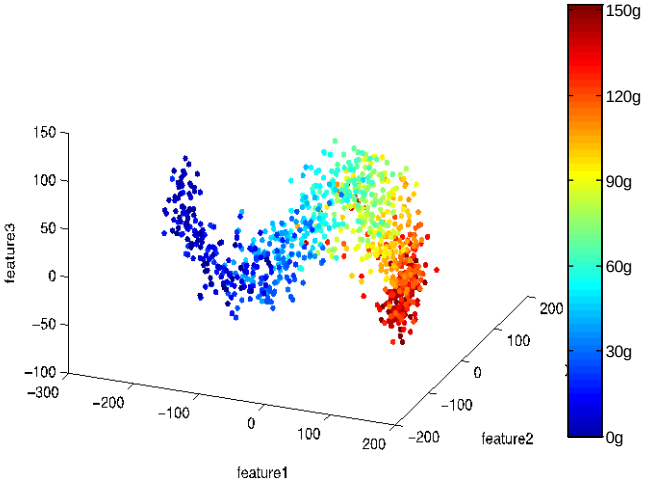

Figure 7. Visualization using three principal components from back wheel sensor accelerometer signals.

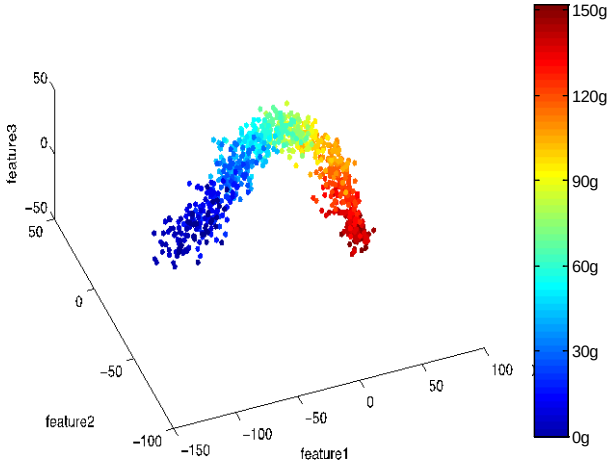

Figure 8. Visualization using first three principal components of the back chassis sensor accelerometer signals.

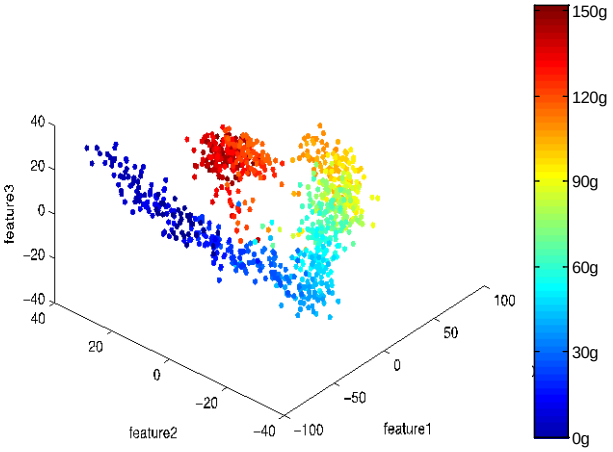

Figure 9. Visualization using first three principal components from front chassis sensor accelerometer signals.

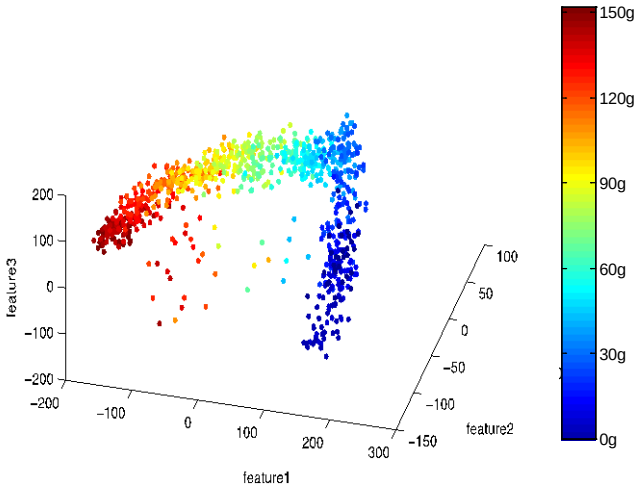

Figure 10. Visualization using three principal components from front wheel sensor accelerometer signals.

3.3 *Signal reconstruction for verification*

To verify that using only first three principal components will capture the characteristics of the original data, we reconstruct the signals based only on these principal components. In Figure 11, we demonstrate that the vast majority of the information in the signal is in the first several principal components—the singular values is a measure of information. Figures 12-16 show the reconstruction the signal for each of the sensor. They demonstrate that the characteristics of the vibration, such as peak occurrences, are very close to the original data, which means that to a certain extent, we can represent the original data by the features extracted from the PCA.

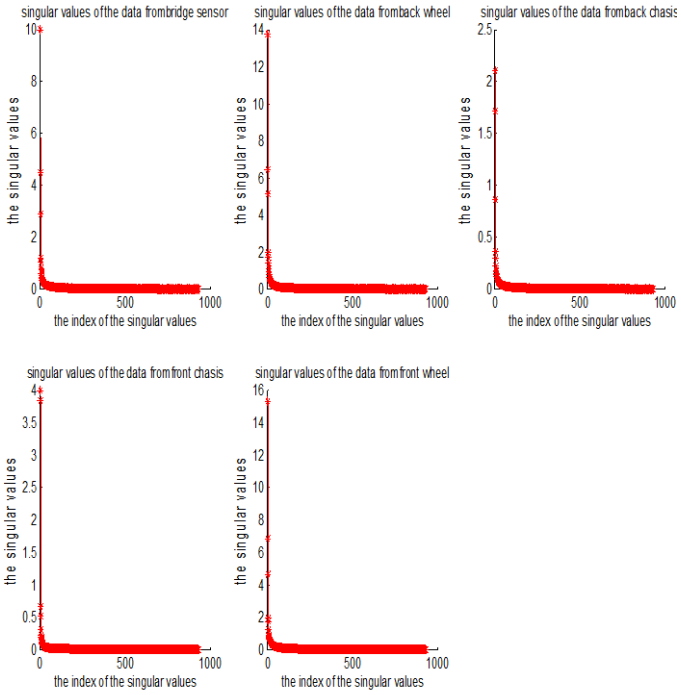

Figure 11. Singular value analysis.

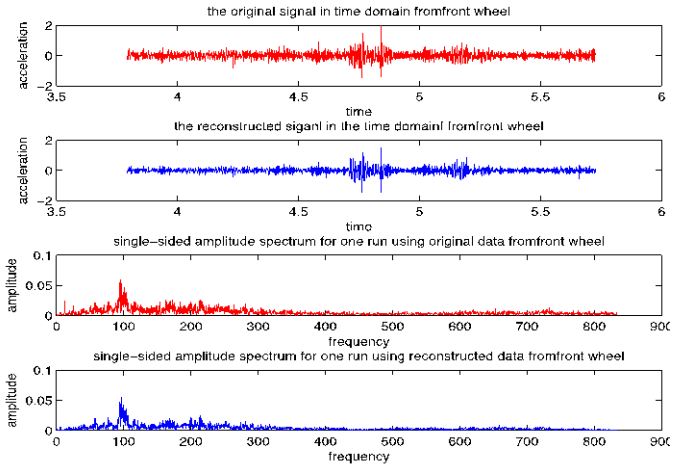

Figure 12. Original and reconstructed signals of the front wheel sensor.

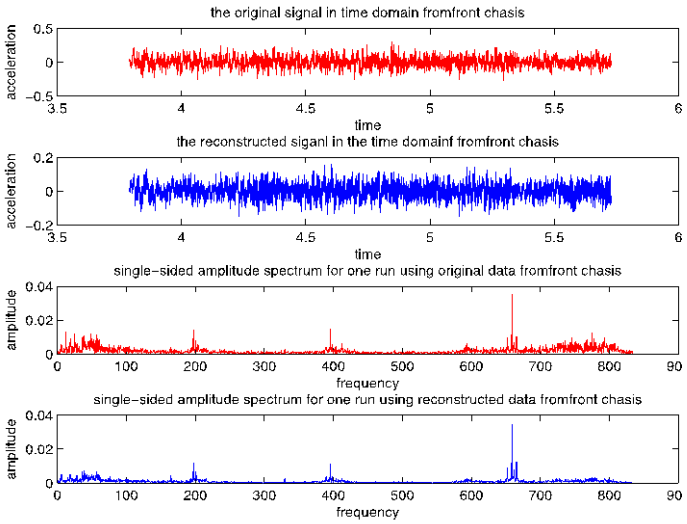

Figure 13. Original and reconstructed signals of the front chassis sensor.

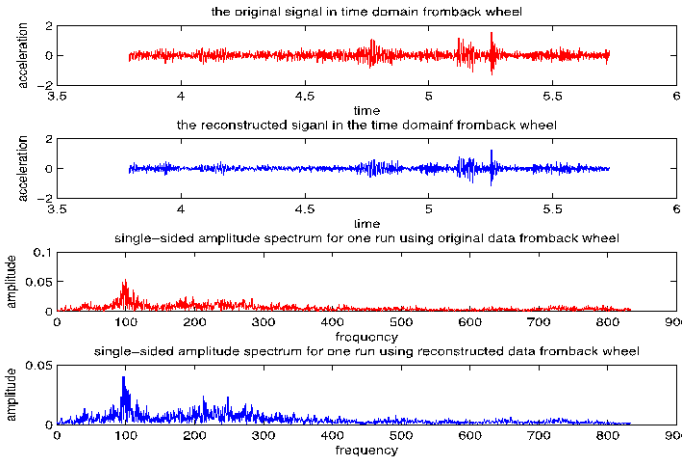

Figure 14. Original and reconstructed signals of the back wheel sensor.

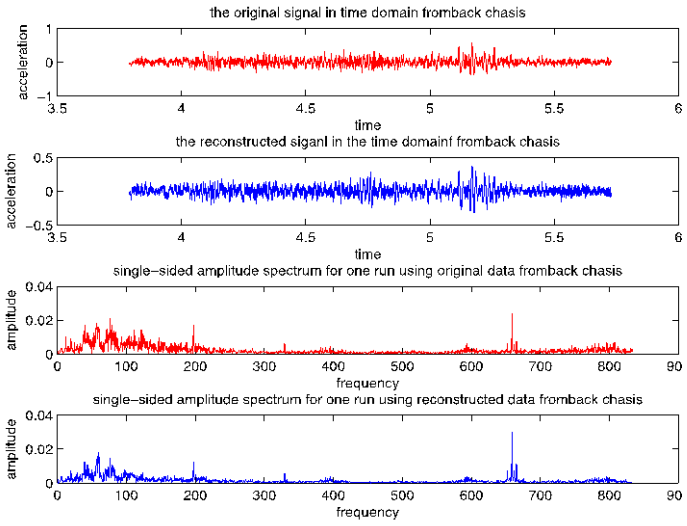

Figure 15. Original and reconstructed signals of the back chassis sensor.

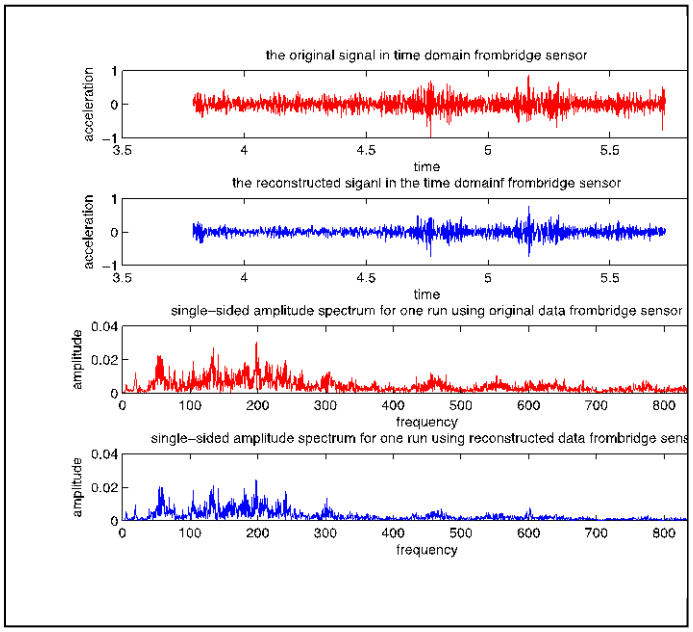

Figure 16. Original and reconstructed signals of the bridge sensor.

3.4 *Regression Model*

Regression analysis is an approach that finds possible connections among all the variables and is then used as the means of predicting when a new dataset is presented to the system. Many techniques for regression analysis have been developed. These techniques are categorized as parametric and nonparametric regression (Wasserman 2005). Parametric regression asks for parameters in the model and nonparametric regression relies on the functions that the user chooses, as the basis for data processing.

We use a nonparametric regression analysis method, called kernel regression. After applying PCA, the acceleration signals are represented in a more compact way by using the top three principal components. Based on this, we then look for a relationship between the acceleration signals and the bridge conditions. We train the regression model from the given dataset by using kernel regression (Wasserman 2005). The advantage of kernel regression is that it is a nonparametric method and it finds a nonlinear relationship between a pair of variables by averaging locally. Noh et al. (2012) used kernel regression to define fragility functions for damage classification purposes. Kernel regression works in two phases, the training phase and the testing phase. In the training phase, the inputs of the kernel regression are the first three coordinates of the acceleration signals after applying PCA and their corresponding bridge conditions; the output is the regression model. In the testing phase, the inputs are the unlabeled acceleration signals and the regression model trained previously, and the output is the predicted bridge condition.

Algorithm *(Regression for Structural Scenarios)*

Input: Labeled training dataset and unlabeled testing dataset.

Output: Predicted class labels for the testing dataset.

Training phase:

1. Compute the discrete Fourier transform of each signal in the training dataset.
2. Conduct the principal component analysis on the training dataset.
3. Preserve the first three components of each signal and the corresponding eigenvectors.
4. Train a kernel regression model by using the three components of each signal in the training dataset and the corresponding class labels.

Testing phase:

1. Compute the discrete Fourier transform of each signal in the testing dataset.
2. Represent signals in the testing dataset by projecting them to the eigenvectors trained in the training phase.
3. Feed the signals into the kernel regression model trained previously and get the predicted class label.

4 RESULTS

4.1 *Regression Testing Protocol*

To evaluate our regression system, we perform a series of cross-validation experiments. For each speed and each sensor, we randomly selected 90% of the acceleration signals from all bridge conditions and use them as the training set. The other 10% of the signals form the test set. This random selection is repeated in a 30-fold validation. For each case, we report MSE as the evaluation score.

4.2 *Severity Results*

The goal is to detect the weight changes with different masses put on the bridge through analyzing the acceleration signals. After the model is built, we randomly choose the signals with the mass on the bridge in the range of training data and calculate the

MSE. These results are shown in Figure 17. MSE for each sensor is found at two different speeds. An MSE of zero would denote a perfect regression, and an MSE of 75 (=30^2/12) would denote regression that was no better at identifying the size of the mass than a random guess. From the result, we see that signals from the wheel have high MSE error so they are less useful, while the signals from the chassis and the bridge have smaller prediction error.

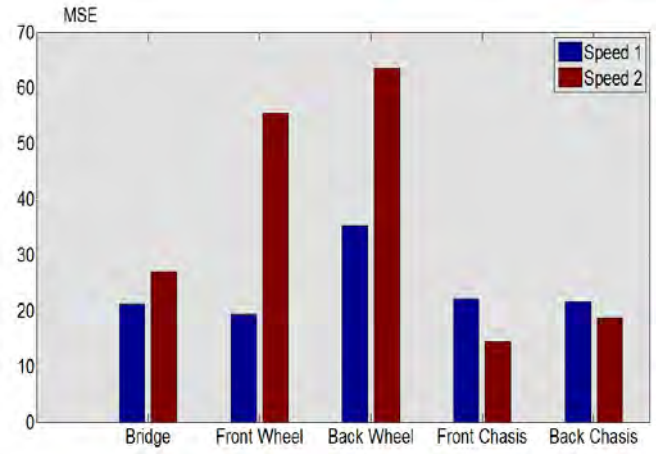

Figure 17. MSE result for the severity regression.

4.3 *Localization Results*

The goal here is to find the location of the mass through analyzing the acceleration signals. We randomly choose signals from locations within the training range and calculate the MSE as shown in Figure 18. Again, a perfect regression would have MSE of zero, and a regression that gave a completely random answer would have a MSE error of 80.1 (=31^2/12). From the result, it is indicated in the similar way with the location case that signals from the chassis have smaller prediction error than that of the signals from wheels. The MSE for the chassis signals is around 15 on average at different speeds.

4.4 *Analysis*

From the above results, we conclude that for the signal processing approach used in this paper, signals from the chassis perform better than signals from the wheels in terms of the prediction error. One possible explanation is the low-pass filtering function of the spring supporting the chassis. As the vehicle is traversing, the spring has the function to filtrate the signal to keep signals with relative low frequencies and filter out signals with relative high frequencies, which associate with noise.

It is also worth noting that the signals from the chassis outperform the signals from the accelerometer located at the midspan of the bridge. This supports our overarching hypothesis that an indirect monitoring approach where sensors are placed on the vehicle may be at least as effective as a direct monitoring approach.

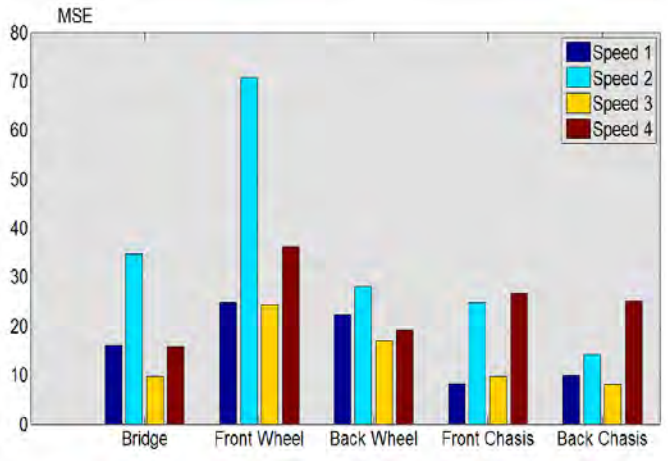

Figure 18. MSE result for the location regression.

5 CONCLUSIONS

We present the latest results of our research into indirect structural health monitoring. We expand on our previous work on locating and quantifying damage in a bridge by using the acceleration signal from a passing vehicle. While previous research has looked at several discrete locations or severity levels using multiclass classification, here we build a regression model that can handle an infinite number of possible damage locations and severities within a particular range. We use PCA to reduce the dimensionality of the signal, and kernel regression, a non-parametric approach, to map the signals to the bridge condition. We obtain low errors (quantified by MSE) for all sensors, in particular for chassis sensors. The error for the chassis sensors is lower than the error for the bridge sensor, which indicates that an indirect monitoring approach may be feasible. This work brings us one step closer to providing a bridge diagnosis in an indirect fashion.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of U.S. National Science Foundation awards 1130616 and 1017278 and of a University Transportation Center grant (DTRT12-G-UTC11) from the US Department of Transportation.

REFERENCES

ASCE. 2013. “Report Card for America’s Infrastructure.” http://www.infrastructurereportcard.org/ (Accessed 9/14/2013)

Casiciati, F. & Giordano, M. 2010. *Structural Health Monitoring 2010: Proceedings of the Fifth European*

Workshop on Strucutral Health Monitoring, Sorrento, Italy. DEStech Publications, Inc.

Cerda, F., Garrett, J., Bielak, J., Bhagavatula, R. & Kovačević, J. 2010. "Exploring Indirect Vehicle-Bridge interaction for SHM". In *Proc. of the Fifth International Conference on Bridge Maintenance, Safety and Management*, IABMAS2010, Philadelphia, USA, 696-702.

Cerda, F., Chen, S., Bielak, J., Garrett, J., Rizzo, P., Kovačević, J. 2013. "Indirect Structural Health Monitoring Of A Simplified Laboratory-scale Bridge Model". Submitted to *Smart Structures and Systems (Special Issue: Challenge on bridge health monitoring utilizing vehicle-induced vibrations.)*, in press.

Chang, F. 2011. *Structural Health monitoring 2011: Condition-Based Maintenance and intelligent Structures. Proceedings of the Eighth International Workshop on Structural Health Monitoring.* DEStech Pub.

Doebling, S. W., Farrar, C. R., Prime, M. B. & Shevitz, D. W. 1996. "Damage Identification and Health Monitoring of Structural and Mechanical Systems From Changes in Their Vibration Characteristics: A Literature Review". *Los Alamos National Laboratory Report.* Report Number: LA-13070-MS.

Duda, R., Hart, P., and Stork, D. (2000). Pattern Classification (2nd Edition). Wiley-Interscience.

Frangopol, D., Sause, R. & Kusko, C. S. 2010. *Bridge Maintenance, Safety and Management – IABMAS'10: in Proc. of the Fifth International IABMAS Conference (1st Edition).* CRC Press. Philadelphia, USA.

Lin, C. W. & Yang, Y. B. 2005. "Use of a passing vehicle to scan the fundamental bridge frequencies: An experimental verification". *Engineering Structures* 27, no. 13 (November): 1865-1878.

Noh, H. Y., Lignos, D. G., Nair, K. K., & Kiremidjian, A. S. 2012. "Development of fragility functions as a damage classification/prediction method for steel moment-resisting frames using a wavelet-based damage sensitive feature". *Earthquake Engineering & Structural Dynamics*, 41, no. 4: 681-696.

Rytter, A. 1993. Vibration Based Inspection of Civil Engineering Structures. Building Technology and Structural Engineering. Aalburg University, Aalborg, Denmark.

Wang, Z., Lederman, G., Chen, S., Cerda, F., Bielak, J., Garrett, J., Rizzo, P. & Kovačević, J. 2013. "Comparison of sparse representation and Fourier discriminant method: Damage location classification in Indirect Lab-scale bridge structural health monitoring". *The 2013 Structures Congress*, pp. 436-446. Doi: 10.1061/9780784412848.039.

Wasserman, L. All of Nonparametric Statistics, Springer Texts in Statistics. Springer, 2005.

Adaptive Graph Filtering: Multiresolution Classification on Graphs

Siheng Chen
Department of ECE
Center for Bioimage Informatics
Carnegie Mellon University
Pittsburgh, PA, USA
sihengc@andrew.cmu.edu

Aliaksei Sandryhaila
Department of ECE
Carnegie Mellon University
Pittsburgh, PA, USA
asandryh@andrew.cmu.edu

José M. F. Moura
Department of ECE
Carnegie Mellon University
Pittsburgh, PA, USA
moura@ece.cmu.edu

Jelena Kovačević
Department of BME & ECE
Center for Bioimage Informatics
Carnegie Mellon University
Pittsburgh, PA, USA
jelenak@cmu.edu

***Abstract*—We present an adaptive graph filtering approach to semi-supervised classification. Adaptive graph filters combine decisions from multiple graph filters using a weighting function that is optimized in a semi-supervised manner. We also demonstrate the multiresolution property of adaptive graph filters by connecting them to the diffusion wavelets. In our experiments, we apply the adaptive graph filters to the classification of online blogs and damage identification in indirect bridge structural health monitoring.**

I. Introduction

With the growing demand in analysis of signal generated by various networks, signals with complex structures arise that require novel processing techniques. Recently, signal processing on graphs has emerged as a new approach to analyze signals with irregular structure that reside on graphs [1]–[4]. Among many applications of this theory, we consider signal classification, which models each signal as nodes of a representation graph, and their labels as a graph signal [4]. Given a subset of initial labels, unknown labels are estimated by propagating known ones along the edges of the graph using graph filters.

In this paper, we propose an adaptive graph filter that extends the application of graph filters to signal classification. Our technique uses an adaptive weighting algorithm that combines results from different graph filters into a global decision. We show that the proposed technique possesses multiresolution properties similar to diffusion wavelets [5]. We also demonstrate applications of adaptive graph filters to the classification of online blogs and damage indentification in indirect bridge structural health monitoring.

II. Background and Problem Formulation

In this section, we introduce the background material necessary for the rest of the paper. We formulate the problem of signal classification and introduce the discrete signal processing on graphs theory that provides the foundation for the proposed adaptive graph filtering technique.

Classification. The purpose of classification is to assign each signal to one of the given classes [6]. Let $\mathcal{X} = \{x^{(i)} \in \mathbb{R}^d\}_{i=1}^N$ be the given dataset with $N = L+U$ elements. The first L elements form the *labeled dataset* $\mathcal{L} = \{x^{(i)} \in \mathcal{X}\}_{i=1}^L$, for which the *ground-truth labels* $\mathcal{Y} = \{y^{(i)} \in \{1, 2, \ldots, C\}\}_{i=1}^L$ are given. The other U elements form the *unlabeled dataset* $\mathcal{U} = \{x^{(i)} \in \mathcal{X}\}_{i=L+1}^N$. Hence, $\mathcal{X} = \mathcal{L} \cup \mathcal{U}$.

Notations		
$\mathcal{X} = \{x^{(i)}\}$	input dataset	$i = 1, \ldots, N$
$\mathcal{L} = \{x^{(i)}\}$	labeled dataset	$i = 1, \ldots, L$
$\mathcal{U} = \{x^{(i)}\}$	unlabeled dataset	$i = L+1, \ldots, N$
$\mathcal{Y} = \{y^{(i)}\}$	ground-truth labels for $\mathcal{L}$	$i = 1, \ldots, L$
$\hat{\mathcal{Y}} = \{\hat{y}^{(i)}\}$	estimated labels for $\mathcal{U}$	$i = L+1, \ldots, N$
$q^{(i)}$	ground-truth vector	$i = 1, \ldots, L$
Q	ground-truth matrix	$L \times C$
$\hat{q}^{(i)}$	confidence vector	$i = L+1, \ldots, N$

TABLE I: Notation used in this paper

A classifier system accepts the entire dataset $\mathcal{X}$ and the *ground-truth labels* $\mathcal{Y}$ as inputs, and produces the *estimated labels* $\hat{\mathcal{Y}} = \{\hat{y}^{(i)} \in \{1, 2, \ldots, C\}\}_{i=L+1}^N$ for the unlabeled dataset $\mathcal{U}$ (see Table I).

We formulate the classification problem as a mapping of the input signal to each class with a certain probability. We regard the label as a posterior probability vector $\hat{q} \in \mathbb{R}^C$, where the cth component, $\hat{q}_c$, is the probability that a signal belongs to the cth class. Since $\hat{q}$ expresses a confidence of an assigned label, we name it a *confidence vector*. The confidence vector for a labeled signal $i = 1, 2, \ldots, L$, is the *ground-truth vector*, $q \in \mathbb{R}^C$, with 0s everywhere except 1 in position c indicating membership in class c. The *ground-truth matrix* Q of size $L \times C$ collects all L ground-truth vectors as its rows, that is,

$$Q = \begin{bmatrix} q^{(1)} & q^{(2)} & \ldots & q^{(L)} \end{bmatrix}^T.$$

Discrete Signal Processing on Graphs. Discrete signal processing on graphs is an extension of the traditional signal processing theory to signals with complex structure residing on irregular domains [4]. The dataset is represented by a *graph* $\mathcal{G} = (\mathcal{V}, P)$, where $\mathcal{V} = \{v_i\}_{i=1}^N$ is the set of *nodes* that represents signals and $P \in \mathbb{C}^{N \times N}$ is an adjacency matrix of the graph called a *graph shift*. We extend the definition of graph signals from [4] from one-dimensional to D-dimensional signal as follows: a *graph signal* s is a map on the set of nodes

$$s : \mathcal{V} \to \mathbb{C}^{N \times D},$$

where D is the dimension of the graph signal at each node.

A linear shift-invariant system, or, a *graph filter*, is defined as

$$H \;=\; h(P) \;=\; \sum_{k=0}^{K} h_k P^k, \tag{1}$$

with $h_k \in \mathbb{C}$, $k = 0, 1, \ldots, K$. A graph filter $H \in \mathbb{C}^{N\times N}$ applied to a graph signal $s \in \mathbb{C}^{N\times D}$ produces an output, which is again a graph signal,

$$Hs \;=\; h(P)s.$$

Discrete signal processing on graphs then defines other fundamental signal processing concepts including the graph Fourier transform, frequency, spectrum, spectral decomposition, and impulse and frequency responses.

III. Adaptive Graph Filtering

In this section, we introduce adaptive graph filters. We also show their multiresolution properties by establishing a connection to diffusion wavelets.

Graph Filtering as Semi-Supervised Classification. Let $\mathcal{G} = (\mathcal{X}, P)$ be a graph with $\mathcal{X} = \{x^{(i)}\}_{i=1}^{N}$ a given datasest and let $P \in \mathbb{R}^{N\times N}$ be a graph shift defined as

$$P_{i,j} \;=\; \frac{\exp\left(\frac{-\rho(x^{(i)},x^{(j)})}{\sigma}\right)}{\sum_{i=1}^{N}\exp\left(\frac{-\rho(x^{(i)},x^{(j)})}{\sigma}\right)}, \tag{2}$$

where ρ is a local distance measurement, such as the ℓ^2 norm or the cosine distance, and σ is a scaling coefficient which scales the local distance. The graph shift defined here is the Hermitian transpose of the *transition matrix* of the graph. P thus has a probabilistic interpretation: constructed as in (2), $P_{i,j}$ gives the probability of transition from the jth node to the ith one in one step [7].

Now we define an *adaptive graph filter*

$$H \;=\; h(P) \;=\; \sum_{k=1}^{K} h_k P^k. \tag{3}$$

This adaptive graph filter is defined similarly to (1), except that its coefficients, h_k, depend on the data and are selected adaptively. We assign $h_0 = 0$, since P^0 does not contribute to the propagation of the labels in the classification problem as will be shown later.

Let the graph signal be the confidence matrix of *all* the signals on the graph, called *prior confidence matrix*, that is, the following map: $s \;:\; \mathcal{X} \to \mathbb{R}^{N\times C}$, defined as

$$\hat{Q}_{\rm pr} \;=\; \begin{bmatrix} Q \\ \mathbf{0}_{u\times C} \end{bmatrix}.$$

In other words, the first L rows of $\hat{Q}_{\rm pr}$ are the confidence matrix Q representing the labeled dataset, while the other U rows are all zeros representing the unlabeled dataset. The prior confidence matrix thus starts with the knowledge on the labeled dataset (the ground truth) and without any knowledge on the unlabeled dataset.

By applying an adaptive graph filter (3) to the prior confidence matrix $\hat{Q}_{\rm pr}$, the *posterior confidence matrix* is

$$\hat{Q}_{\rm ps} \;=\; H\hat{Q}_{\rm pr}, \tag{4}$$

where the ith row is the confidence vector $\hat{q}^{(i)}$ for the ith node.

The coefficients h_k of the adaptive filter (3) are then found as follows. Rewrite (4) as

$$\hat{Q}_{\rm ps} \;\overset{(a)}{=}\; \sum_{k=1}^{K} h_k P^k \hat{Q}_{\rm pr} \;=\; \sum_{k=1}^{K} h_k \hat{Q}_k, \tag{5}$$

where (a) follows from (3) and $\hat{Q}_k$ is the $N \times C$ confidence matrix for the kth graph shift defined as $\hat{Q}_k = P^k\hat{Q}_{\rm pr}$. For each i, the desired confidence vector $\hat{q}^{(i)}$ (ith row of $\hat{Q}_{\rm ps}$) is thus a weighted linear combination of corresponding rows from each graph shift $\hat{Q}_k$. The optimal filter coefficients are designed to fit the confidence vectors of the labeled signals from all the graph shifts to the ground truth and to minimize the labeling uncertainty of the unlabeled signals. We introduce the labeling uncertainty measure [8], [9] as

$$\mathrm{M}(\hat{q}) \;=\; \mathrm{H}(\hat{q})\,(\chi_{d>T} + \lambda(d)\,\chi_{d\le T}), \tag{6}$$

where χ_I is the indicator function of an interval I, $d \;=\; |\hat{q}_{(1)} - \hat{q}_{(2)}|$ with $\hat{q}_{(1)}, \hat{q}_{(2)}$ the first and second largest element in $\hat{q}$, respectively, T is the threshold, and $\lambda(d) = 1 + 5(d/T - 1)^2$ is a penalty function that is large when the first and second largest elements are close. The optimal filter coefficients are then given by the solution to the minimization problem

$$\begin{aligned} \hat{h} \quad &= \quad \arg\min_{h} \{\alpha \sum_{i=1}^{L} \|q^{(i)} - \hat{Q}^{(i)}h\| \\ &+ \quad (1-\alpha)\|h - \gamma\|\}, \end{aligned} \tag{7}$$

where $h = \begin{bmatrix} h_1 & \ldots & h_K \end{bmatrix}^T$ with the constraint $||h||_1 = 1$; $\alpha \;=\; L/(L+U)$ is the labeling ratio; matrix $\hat{Q}^{(i)} = \begin{bmatrix} \hat{q}_1^{(i)} & \hat{q}_2^{(i)} & \ldots & \hat{q}_K^{(i)} \end{bmatrix}$ is an $C \times K$ confidence matrix of the ith signal that collects the graph shifts' individual confidence vectors; and $\gamma = \begin{bmatrix} \gamma_1 & \gamma_2 & \ldots & \gamma_K \end{bmatrix}$ collects the discriminative powers of each graph shift,

$$\gamma_k \;=\; \frac{e^{-(\beta/U)\sum_{i=l+1}^{N}\mathrm{M}(\hat{q}_k^{(i)})}}{\sum_{j=1}^{K} e^{-(\beta/U)\sum_{i=l+1}^{N}\mathrm{M}(\hat{q}_j^{(i)})}}.$$

Here, β is the decay coefficient that controls the distribution of the discriminant power from all the graph shifts. The first term in (7) represents the contribution from all labeled signals. The second term in (7) represents the contribution from all unlabeled signals; to obtain it, we fit weights to the confidences from all the graph shifts . We use the labeling ratio to balance these two terms. Since this is a convex optimization problem, it is numerically efficient to solve. After weighting, we compute the global decision as $\hat{y}^{(i)} \;=\; \arg\max_c \hat{q}_c^{(i)}$, where $\hat{q}^{(i)} = \hat{Q}^{(i)}\hat{h}$ (see Algorithm 1). For more details, see [9].

Relation to Diffusion Functions. Diffusion functions are the graph-based semi-supervised classifiers [10]. The classifier operates by propagating known labels using the transition matrix a finite number of times. Given a transition matrix W and a prior confidence matrix $\hat{Q}_{\rm pr}$, the posterior confidence matrix $\hat{Q}_{\rm ps}$ is obtained by $\hat{Q}^*_{\rm ps} \;= \hat{Q}^*_{\rm pr}W^t$, where $*$ denotes the Hermitian transpose and t is the transition time. If the filter coefficient h is a Kronecker delta impulse at position t, then

Algorithm 1 Adaptive graph filter

Input	$\mathcal{X}$	input dataset
	$\hat{Q}_{\rm pr}$	prior confidence matrix
Output	$\hat{Q}_{\rm ps}$	posterior confidence matrix
Parameters	$P_{i,j}$	graph shift
	h_k	filter coefficients
Function	**AGF**$(\mathcal{X}, \hat{Q}_{\rm pr})$	
	$P_{i,j} = (2)$	graph shift construction
	$h = (7)$	filter coefficient optimization
	for $k = 1 : K$ **do**	filtering
	$\quad \hat{Q}_k = P^k \hat{Q}_{\rm pr}$	
	$\hat{Q}_{\rm ps} = \sum_{k=1}^{K} h_k \hat{Q}_k$	weighting
	return $\hat{Q}_{\rm ps}$	

$h = \delta_t$, the adaptive graph filter coincides with the diffusion function. Hence, diffusion functions form a subset of the adaptive graph filters that have only one non-zero coefficient. In practice, it can be hard to determine the required transition time; as a consequence, the performances of diffusion functions can be susceptible to initial conditions. In contrast, the adaptive graph filter learns the contribution from each transition time by using (7).

Relation to Diffusion Wavelets. Diffusion wavelets are a multiscale framework for the analysis of signals with complex structure [5]. They can be seen as an extension of the classical wavelet theory, where, the diffusion wavelet basis is learned from the geometry of the signal structure in a data-adaptive way. A diffusion wavelet basis is constructed by dilation using the dyadic powers of the transition matrix, the idea being that they propagate local relationships to global relationships throughout the graph.

Given a graph with the transition matrix T, the transition matrix at the jth resolution level is T^{2^j}, $j = 1, 2, \ldots$. Since the second singular value of the transition matrix is less than 1 to keep the graph connected, only the first singular value is 1. Consequently, if the transition matrix is raised to a high power, all the singular values disappear except for the first one,

$$\lambda_i^{2^j} \;\to\; 0, \qquad i = 2,\, 3,\, \ldots,\, N \quad \text{as} \quad j \to \infty.$$

When $j = 1$, the transition matrix T measures local pairwise similarities; increasing the power j gradually decreases the rank of the transition matrix T^{2^j} and causes local information of the graph to be missed since the resolution on the graph changes from finest to the coarsest. Thus, by changing j, we can both perform a multiresolution analysis as well as doing it in a computationally efficient manner.

The adaptive graph filter in (3) is formed as a linear combination of graph shifts P raised to power k. When k is large, P^k become a low-rank matrix describing the global information of the graph. Since the graph filter P is defined as the Hermitian transpose of the transition matrix T in (2), P^k works for the adaptive graph filter just as T^{2^j} does for diffusion wavelets. Each k corresponds to a different resolution on the graph, and thus, adaptive graph filter actually performs multiresolution classification on the graph. It weighs the classification results from each resolution to produce the global result. The filter coefficients represent the discriminative power of each resolution.

IV. Experimental Results

In this section, we apply adaptive graph filtering to classification of online blogs and damage identification in indirect bridge structural health monitoring.

A. Classification of Online Blogs

Dataset. We consider 1224 online political blogs that we wish to classify as conservative or liberal based on their context [11]. The graphs corresponding to each blog are obtained by tracing hyperlink references between blogs.

Experimental Setup. We perform a binary clasification, and vary the labeling ratio as 0.5%, 1%, 2%, 5% and 10%; in other words, we label 6, 12, 24, 60, 120 blogs, respectively. We considered two methods for which nodes to label initially: random selection, and selection of blogs with most hyperlinks. We ensure we have the same number of signals from both classes in the labeling dataset and compare the performances of the diffusion functions (DF) and adaptive graph filters (AGF). We choose the transition time $t = 2$, filter length $K = 10$ and the decay coefficient $\beta = 1$, and perform a 30-fold cross-validation.

Results. Fig. 1 shows the dependence of classification accuracy on the labeling ratio for the two selection mechanisms. Both show similar trends: adaptive graph filter works consistently better than diffusion functions. Another interesting observation is that when labeling ratio is low, initial labeling with most hyperlinks provides better accuracy, but when labeling ratio is sufficiently high, initial labeling with most hyperlinks has no advantage over initial labeling with random selection. The reason is that when labeling ratio is low, nodes with higher degree propagate labels to more nodes, but when labeling ratio is high, a large number of labeled nodes has the same effect.

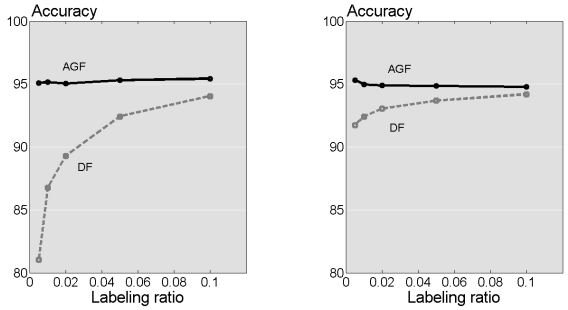

(a) Random selection. (b) Most hyperlinks.

Fig. 1: Accuracy as a function of the labeling ratio.

B. Damage Detection in Bridge Structure Monitoring

Dataset. We built a lab-scale bridge-vehicle dynamic system and put a sensor on a vehicle, and let it move across the bridge. We capture vibration characteristics of the bridge from the vibration of the traversing vehicle through the acceleration signal. We collected 30 acceleration signals for each of 13

V	S	AGF	DF	SVM
1	1	**87.98**	83.70	57.97
	2	**89.55**	85.78	70.42
	3	**89.91**	86.25	74.29
	4	**93.87**	88.08	74.82
	5	**77.64**	74.87	70.68
	6	**85.24**	82.43	67.13
	7	**66.23**	65.31	59.48
	8	**67.11**	66.52	56.53
2	1	**80.99**	78.17	49.75
	2	**70.32**	67.30	53.64
	3	**86.09**	82.50	67.96
	4	**82.86**	79.57	61.52
	5	**81.22**	77.92	62.75
	6	**81.73**	80.17	66.89
	7	**85.48**	82.87	65.09
	8	**82.57**	80.32	48.57

TABLE II: Accuracy comparison of Vehicles (V) 1 and 2, with Speeds (S) 1, 2, ..., 8, and labeling ratio of 10%.

different bridge conditions, 8 different speeds and 2 different vehicles, for a total of 6240 acceleration signals [8], [12].

Experimental Setup. Given a specific vehicle driven at a specific speed, we want to classify 13 bridge conditions, in particular with a low labeling ratio. We have 16 vehicle-speed cases for each of which there are 30 acceleration signals per 13 bridge conditions, and vary the labeling ratio as 10%, 30%, 50%, 70% and 90%; the final accuracy is the average over the 13 scenarios.

We ensure we have make the same number of signals from both classes in the labeled dataset. To reduce dimensionality of the raw acceleration signals, we conduct principal component analysis on all the signals and preverve 95% energy [6]. We then compare the performances of kernel support vector machine (SVM), which is a standard classifier that works through maximizing the margins of different classes [6], the diffusion functions and the adaptive graph filter. We construct a $k = 4$ regular graph (each vertex connects to 4 neighbors). We choose the local measurement ρ to be the cosine distance, scaling coefficient $\sigma = (1/N^2) \sum_{i,j} \rho(x^{(i)}, (x^{(j)}))$, transition time $t = 8$, filter length $K = 30$, the decay coefficient $\beta = 1$, and the penalty threshold $T = 0.02$, and perform a 30-fold cross-validation.

Results. Table II compares the performance of different classifiers with the low labeling ratio of 10%. We use V for vehicle and S for speed. In our experiments, kernel support vector machine performs poorly, diffusion function produce better accuracy, while the adaptive graph filters yields the highest accuracy.

Fig. 2 shows the dependence of classification accuracy on the labeling ratio for 2 vehicles averaged across 8 speeds. Both figures show similar trends; as the labeling ratio decreases, accuracy of SVM drops sharply; performances of diffusion function and adaptive graph filter stay relatively flat even at very low labeling ratios; and adaptive graph filter works consistently better than diffusion functions.

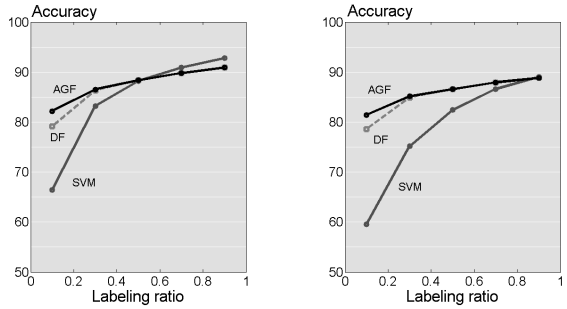

(a) Vehicle 1. (b) Vehicle 2.

Fig. 2: Accuracy as a function of the labeling ratio.

V. Conclusions

We presented an adaptive graph filtering technique for data classification based on a semi-supervised optimization algorithm that uses both labeled and unlabeled signals to learn the filter coefficients. Adaptive graph filters combine decisions from multiple graph filters to form a global decision. We also discussed the connection between adaptive graph filters and diffusion wavelets, and demonstrated their application to data classification and indirect bridge structural health monitoring.

Acknowledgements

The authors gratefully acknowledge support from the NSF through awards 1130616 and 1017278, as well as CMU Carnegie Institute of Technology Infrastructure Award.

References

[1] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, pp. 83–98, 2013.

[2] S. K. Narang and A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *IEEE Trans. Signal Process.*, vol. 60, pp. 2786–2799, 2012.

[3] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Appl. Comput. Harmon. Anal.*, vol. 30, pp. 129–150, 2011.

[4] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, 2013.

[5] R. R. Coifman and M. Maggioni, "Diffusion wavelets," *Appl. Comput. Harmon. Anal.*, pp. 53–94, July 2006.

[6] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, John Wiley & Sons, Englewood Cliffs, NJ, 2001.

[7] F. R. K. Chung, *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*, Am. Math. Soc., 1996.

[8] S. Chen, F. Cerda, J. Guo, J. B. Harley, Q. Shi, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Multiresolution classification with semi-supervised learning for indirect bridge structure health monitoring," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Vancouver, Canada, May 2013.

[9] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring," *IEEE Trans. Signal Process.*, June 2013, Submitted.

[10] A. D. Szlam, R. R. Coifman, and M. Maggioni, "A general framework for adaptive regularization based on diffusion processes," *J. Mach. Learn. Res.*, , no. 9, pp. 1711–1739, Aug. 2008.

[11] L. A. Adamic and N. Glance, "The political blogosphere and the 2004 u.s. election: Divided they blog," pp. 36–43, 2005.

[12] F. Cerda, S. Chen, J. Bielak, J. H. Garrett, P. Rizzo, and J. Kovačević, "Indirect structural health monitoring of a simplified laboratory-scale bridge model," *Int. J. Smart Struct. Syst.*, 2012, Submitted.

Comparison of Sparse Representation and Fourier Discriminant Methods: Damage Location Classification in Indirect Lab-scale Bridge Structural Health Monitoring

Z. Wang[1], S. Chen[2], G. Lederman[1], F. Cerda[3], J. Bielak[1], J. H.Garrett[1], P. Rizzo[4] and J. Kovačević[2]

[1]Department of Civil and Environmental Engineering, [2]Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.
[3]Departamento de Ingeniería Civil, Universidad de Concepción, Concepción, 4070409, Chile.
[4]Department of Civil and Environmental Engineering, University of Pittsburgh, Pittsburgh, PA 15261, USA.

ABSTRACT

This paper presents a novel method for interpreting data to improve the indirect structural health monitoring (SHM) of bridges. The research presented in the study is part of an ongoing study aimed at developing a novel SHM paradigm for the health assessment of bridges. In this paradigm, we envision the use of an instrumented vehicle that assesses a bridge's dynamic characteristics while traveling across the bridge. These characteristics are then correlated to the health of the structure by means of advanced signal processing and pattern recognition approaches. In this paper, we present and compare two classification algorithms that locate the presence of damages at well-defined locations on the structure: sparse representation and the Fourier discriminant methods, and find that the sparse representation method provides superior classification accuracy.

INTRODUCTION

The need for bridge structural health monitoring (SHM) has become clear in the aftermath of the I-35 Bridge collapse in Minneapolis. In the past two decades, the SHM community has done significant work to develop cost-effective technologies for the SHM of bridges and other large structures (Doebling et al. 1996; Chang 2011; Frangopol et al. 2010; Casciati and Giordano 2010). The majority of these works are based on direct methods, in which one or more arrays of sensors are mounted on the structure of interest to collect the dynamic and static characteristics of the bridge. These characteristics are then processed to infer the structure's health. The direct methods, however, may be too expensive to be implemented on an entire nation's bridge inventory.

The indirect approach, on the other hand, overcomes this limit by installing sensors on a moving vehicle that, by crossing the bridge of interest, can capture the dynamic vehicle-bridge interaction and thus, indirectly, the dynamic characteristics of the bridge itself (Lin et al. 2005; Cerda et al. 2010). Using sophisticated signal processing algorithms, information about the bridge condition can be extracted from the vehicle-bridge interaction signals, which may allow for the detection and location of structural damage.

In this paper, we present the latest advancements from an ongoing project at Carnegie Mellon University and the University of Pittsburgh aiming to develop a holistic indirect health monitoring paradigm for bridges. In these experiments, we have focused on a laboratory-scale model, where we simulate the vehicle-structure interaction; the structure consists of one simplified laboratory-scaled bridge, inspected by a schematized vehicle traveling along rigid rails. The work presented in this paper follows (Cerda et al. 2013) and builds on it by studying a new feature-selection technique to more accurately identify and locate the presence of simulated defects induced on a bridge model. We compare Fourier discriminant method, which has been used in previous work, to the sparse representation method to detect and locate damages on the bridge.

EXPERIMENTAL SETUP AND PROTOCOL

Experimental setup. The model used in this study is shown in Figure 1; it consists of a mechanical system, a motion-control system, and a data-acquisition system. The mechanical system includes the structure and the car. The structure has an acceleration ramp, a bridge deck with two rails, and a deceleration ramp. The ramps are built to guarantee that the speed of the car is constant when traveling over the deck. The deck consists of a simply supported aluminum plate, 2438 mm (8 feet) long. Two angle beams act as girders, and two rails, mounted on the top surface, guide the car. The car, shown in Figure 2, has four independent suspension systems and four pedestal bases. Four accelerometers are secured to the car. Two sensors are mounted on top of the front and rear suspensions and two sensors are mounted on the top of the front and rear traverse stands. Hereafter, the sensors on the front are labeled as A.V. and the ones on the rear as B.V.

The vehicle is pulled across the bridge by a cable-and-belt system, which is connected to the motor, as indicated in Figure 1. The motion-control and the data-acquisition systems consist of a National Instruments® PXI system running in LabView®, a PXI Chassis (NI PXI 1031) with a motion-control card (NI PXI 7342), a motion interface (UMI 7772), a stepper drive (P70360) and a dual shaft stepper motor (NEMA 34). The acceleration data are collected at a 1667 Hz sampling rate and stored for data processing. More details about the experimental setup are reported in (Cerda et al. 2013).

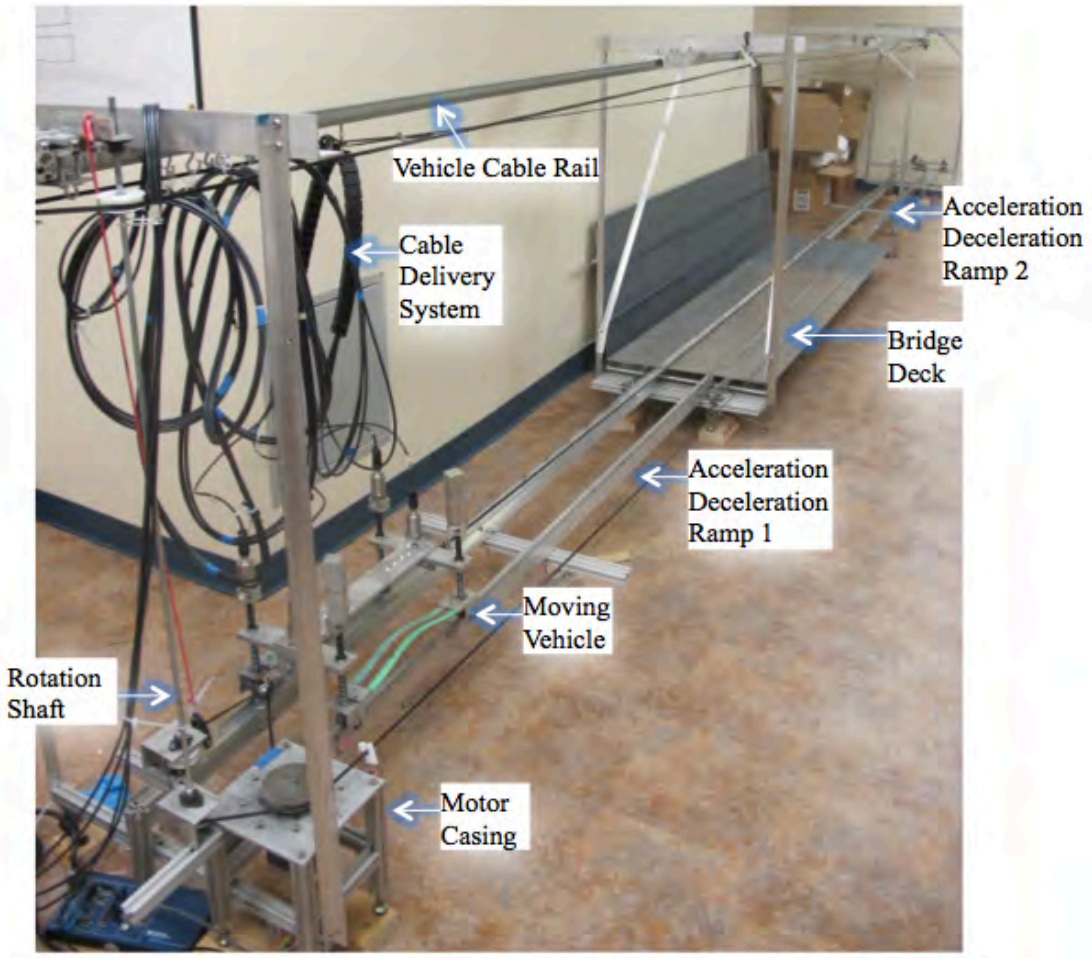

Figure 1. The lab-scale vehicle-bridge model.

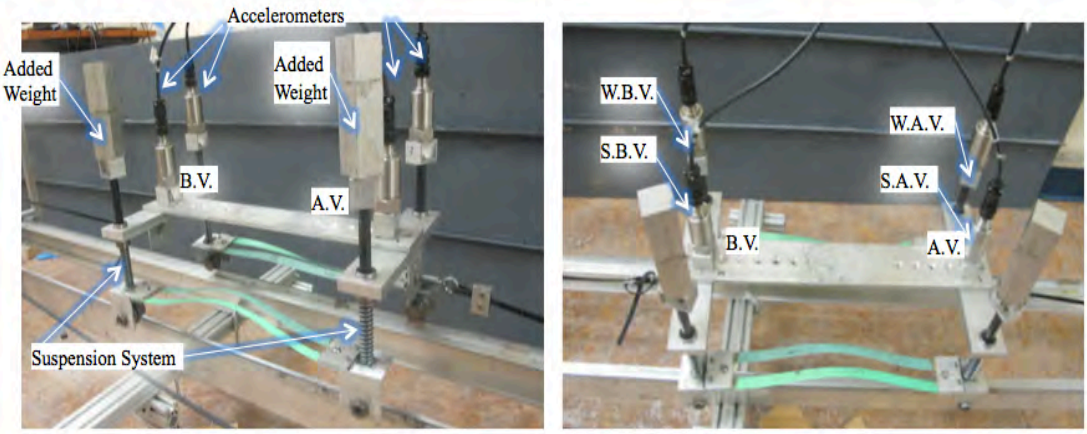

Figure 2. Close-up views of the vehicle.

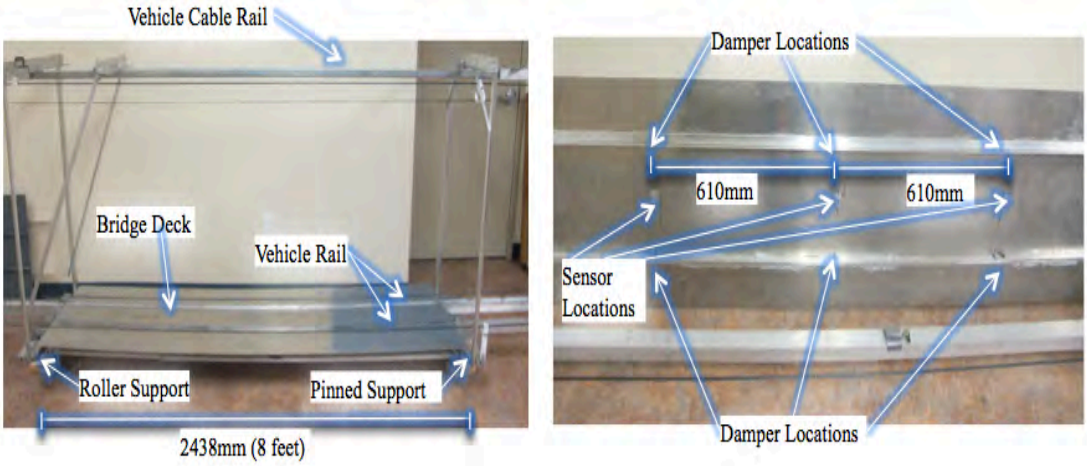

Figure 3. Close-up views of the bridge deck.

In addition to collecting data indirectly from the vehicle, we also collect data directly from the bridge itself to compare the two data collection techniques. Three sensors are mounted underneath the deck at the locations shown in Figure 3. They are located at 1/4L, 1/2L and 3/4L, where L is the deck's span. Figure 3 shows a close-up view of the deck.

Protocol. Damage is simulated by either placing mass, changing boundary restraints or installing dampers, each at four locations, on the bridge. Our aim is to find the accuracy with which we can classify the state of the bridge into these four categories (four locations). This four-class classification problem is repeated with two different masses, different rotational restraints and different damping conditions (see Table 1). The effects of these variations on the fundamental frequency and critical damping of the bridge are quantified by performing conventional free vibration tests. The results of these measurements are summarized in Table 1, which associates certain modal characteristics of the bridge to the simulated damage scenario. Each scenario is labeled to ease the identification of the damage. The first two numbers in the name string (scenario label) indicate a specific damage type and the last number the damage location, 0, 1, 2 or 3. Table 1 lists 24 scenarios. For each scenario, the car moved at 4 speeds, varying from 1.5 m/s to 2.25 m/s at 0.25 m/s. To assess repeatability, 32 measurements were taken. In total, we performed 24 (scenarios) x 4 (speeds) x 32 (iterations) = 3072 trials.

Table 1. Damage scenarios and bridge characteristics.

		Bridge characteristics			
Scenarios	**SC-**	**Frequency [Hz]**	**% f shift**	**% Critical damping**	**% d_{crit} shift**
	020	7.86	1.06	15.63	0.20
	021	7.87	1.30	19.64	0.51

	022	8.00	2.99	13.46	-0.03
	023	8.17	5.20	24.85	0.91
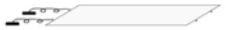	030	8.10	4.27	16.17	-0.24
	031	8.30	4.27	21.05	0.61
	032	8.40	8.14	19.53	-0.5
	033	8.10	6.85	9.98	-0.23
	070	7.77	0.00	18.93	0.45
	071	7.77	0.00	27.06	1.08
	072	7.77	0.00	18.54	-0.42
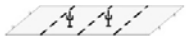	073	7.77	0.00	23.40	0.79
	080	7.67	-1.30	34.52	1.65
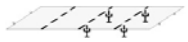	081	7.67	-1.30	33.80	1.59
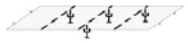	082	7.67	-1.30	35.51	1.72
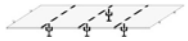	083	7.70	-0.88	37.03	1.84
	110	7.67	-1.30	11.82	-0.09
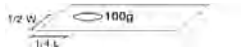	111	7.67	-1.30	12.28	-0.06
	112	7.67	-1.30	11.70	-0.10
	113	7.67	-1.30	13.81	0.06
	120	7.60	-2.16	11.25	-0.14
	121	7.70	-0.88	12.55	-0.04
	122	7.77	0.00	12.62	-0.03
	123	7.60	-2.16	11.28	-0.14

SIGNAL ANALYSIS AND CLASSIFICATION

The task of distinguishing various bridge conditions can be treated as a signal classification problem that requires associating an input signal with a pre-defined class label. In this paper, we tested two classification algorithms: Fourier discriminant method and sparse representation method. Fourier discriminant method searches for the most discriminative frequency features and then uses support vector machine (SVM) (Duda et al. 2000) to classify those features. Sparse representation method searches for a dictionary to represent each class and assigns the signal to the class whose corresponding dictionary gives the minimum representation error. We now discuss these in more detail.

(1) Classification based on Fourier discriminant method. The basic idea of this algorithm is to extract the most discriminative features of the vibrational signal in the frequency domain based on their discriminative power, and use those features to discriminate among different scenarios. Algorithm 1 as shown below summarizes the process. In the training phase, we first compute the discrete Fourier transform (DFT) of each signal in the training dataset. Then, we compute the discriminative power of each Fourier basis vector by using Fourier discriminant basis vector selection (Cerda et al. 2013) and recorde the order of the discriminative power. The k most discriminative Fourier basis vectors are then used as features. We then train an SVM classifier by using these features and their corresponding known labels. In the testing phase, we compute the DFT of each signal in the testing dataset, and then the Fourier discriminant features by retrieving the order of the discriminative power learned in the training phase. These features are then classified using the SVM classifier trained previously. More details can be found in (Cerda et al. 2013).

Algorithm 1 (Fourier discriminant method)
Input: labeled training dataset. **Output:** class labels.
Training phase

1. Compute the DFT of each signal in the training dataset.
2. Compute the discriminative power of each Fourier basis vector by using Fourier discriminant basis vector selection and record the order of the discriminative power.
3. Use k most discriminative Fourier basis vectors as features
4. Train SVM classifier by using features from Step 3 and their corresponding known labels.

Testing phase

1. Compute the DFT of each signal in the testing dataset.
2. Compute the Fourier discriminant features by retrieving the order of the discriminative power that learns in the training phase.
3. Classify the features by using SVM classifier trained previously.

(2) Classification based on sparse representation method. Sparse representations have been shown to be a strong tool for representing and denoising signals (Wright et al. 2009). Although a signal may have high dimensionality, it could have a compact representation in a low-dimensional space, which means the signal can be sparsely represented without losing information. The motivation to use sparse representations when analyzing vibrational signals is that we assume there are only a few key factors influencing a vehicle-bridge system. By squeezing most information into a small number of elementary signals, we find these key factors.

In sparse representations, signals can be represented by using a few expansion coefficients in a fixed over-complete dictionary, called sparse coding. The representation dictionary is updated by fitting the data, and usually referred to as dictionary learning. Each dictionary corresponds to a specific vehicle-bridge system. This classification can be done by labeling each signal and representing it by a dictionary that gives the minimum representation error. We now explain these concepts in more detail.

Sparse coding. Sparse coding computes the expansion coefficients based on the given signal and the dictionary. The exact solution is known as a nondeterministic polynomial hard problem, and thus, approximate algorithms are considered instead. Here, we use the orthogonal matching pursuit (OMP) algorithm (Chen et al. 1989). OMP is a greedy algorithm that selects the basis vector from a dictionary sequentially. In each iteration, the dictionary finds the best basis vector to fit the given signal and leaves the residue for the next iteration. Although it cannot provide a global optimum, it only involves the computation of inner products between the signal and basis vector, leading to a computationally efficient algorithm.

Algorithm 2 (Sparse representation)
Input: labeled training dataset. **Output:** class labels.
Training phase

1. Initialize a dictionary for each class and use OMP to represent the each signal in the training dataset.
2. For each class, update the dictionary to fit to training signals by using K-SVD.
3. Iterate Steps 1 & 2 in turn until the reconstruction error is below a given threshold and record the dictionaries for all classes.

Testing phase

1. Represent each signal in the testing dataset on the dictionaries for all classes by using OMP.
2. Label each signal to the class that the corresponding dictionary gives the minimum representation error.

Dictionary learning. Dictionary learning updates the dictionary iteratively to reach the goal of expanding the given signal sparsely. Unlike traditional transforms, such as Fourier and wavelets, the dictionary in sparse representation is flexible and can be learned by fitting itself to the given signals. Some state-of-the-art algorithms are methods of optimal directions (MOD) (Engan et al. 1999) and K-SVD (Aharon et al.

2006), which we use in this paper. It is designed to be a generalization of the K-means algorithm (MacQueen and J. B. 1967) and work towards minimizing an overall objective function.

We perform sparse coding and dictionary learning to get the dictionary and the representation coefficients for the given signals. Each signal is then assigned to the class for which the corresponding dictionary provides the minimum representation error. This algorithm is summarized in Algorithm 2.

CLASSIFICATION RESULTS AND DISCUSSION

We set up a similar damage location classification experiment in our previous work (Cerda 2012). To make the experiments more realistic, we lower the training ratio (the ratio of the number of the training samples and all samples) from 95% to 50%. Figures 4 and 5 shows the accuracies of classifying damage location for six different types of damage clusters, SC02-, SC03-, SC07-, SC08-, SC11- and SC12-, for the Fourier discriminant method and sparse representation method, respectively. Each type of a damage cluster has four different location scenarios indicated by the last number in the name string as in Table 1. The accuracies in the figures are the averaged accuracies over four speeds and all iterations. The red vertical bar represents the corresponding standard deviation. We show results associated with the transducers mounted on the bridge and the sensors mounted on the car's suspensions and wheels. On average, the sparse representation method provides about 10% higher accuracy than the Fourier discriminant method. The lowest accuracies for both methods are for SC08-; from Table 1, we see that the difference among the four scenarios in SC08- is only changing the locations of the two dampers, and only one out of four scenarios changes the frequency of the vibration by a small amount. Thus, these variations cannot be detected easily. Even then, sparse representation method is still able to improve accuracy by about 10% compared to Fourier discriminant method. Overall, we conclude that by using the sparse representation method, most of the accuracies are above 70% and all of them are above 60%.

We also evaluated the effect of speed on the classification accuracies of the two algorithms. Figure 6 shows the classification accuracy of both algorithms as a function of the vehicle speed. Two interesting trends are visible: (1) We see higher classification accuracy when using sparse representations as previously noted. (2) We observe a decrease in the classification accuracy as the speed increases. This accuracy decrease can be explained as follows: as the speed goes up, the length of the analyzed signal is shorter and therefore less information from the bridge can be extracted. Also, as the excitation frequency increases, the vehicle may excite higher modes that may be less sensitive to the presence of the simulated damage.

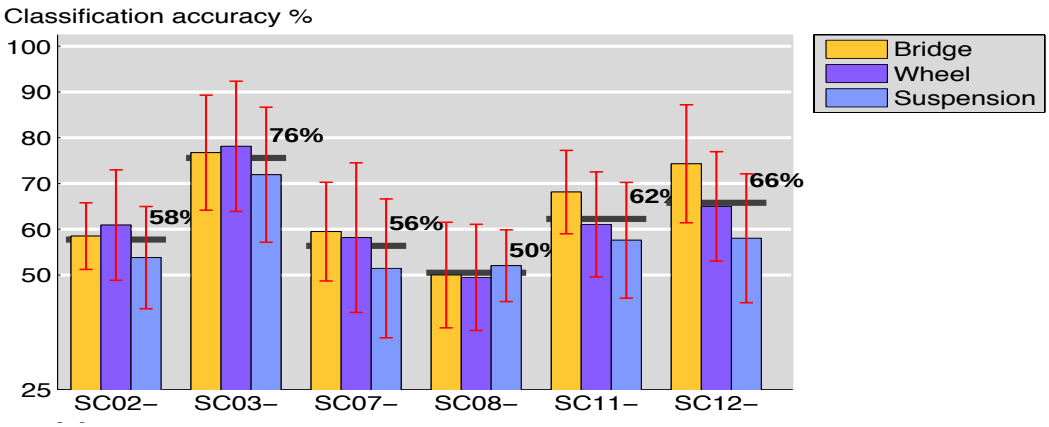

Figure 4. Classification accuracy using Fourier discriminant method.

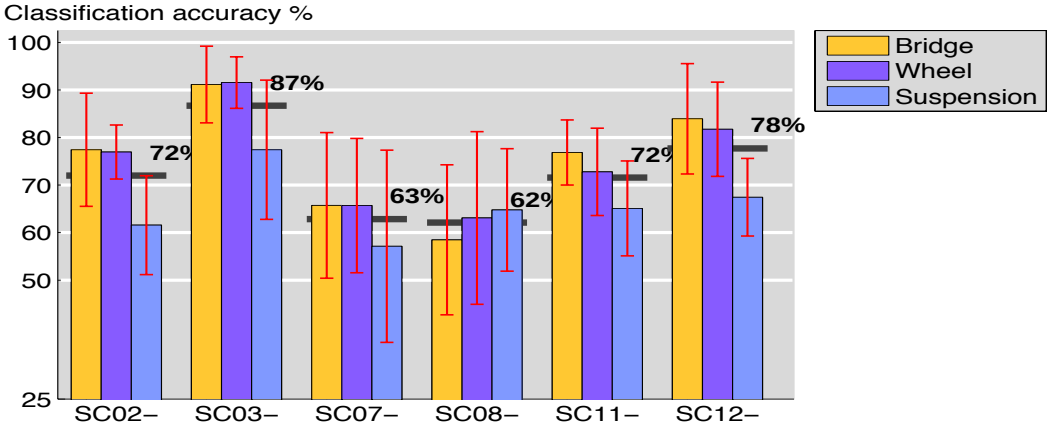

Figure 5. Classification accuracy using sparse representation method.

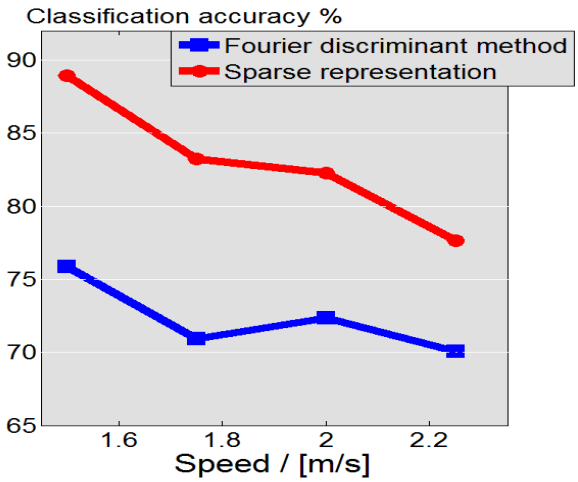

Figure 6. Classification accuracy as a function of speed.

CONCLUSIONS

We presented the latest advancement of an ongoing project at Carnegie Mellon University and at the University of Pittsburgh that aims at developing a holistic approach for the indirect health monitoring of bridges. We compared a proposed novel classification algorithm with a previous method (Cerda 2012). The proposed method, using sparse representations, was tested on vibrational signals collected from a laboratory-scale bridge model, which simulates the vehicle-bridge interaction. The results are compared to Fourier discriminant method that has been used previously. We found that the sparse representations allows for higher classification accuracy when compared with the Fourier discriminant method, and that the accuracy of each algorithm is inversely proportional to the speed of the car.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of U.S. National Science Foundation awards 1130616 and 1017278 and of a University Transportation Center grant (DTRT12-G-UTC11) from the US Department of Transportation.

REFERENCES

Aharon, M., Elad, M. and Bruckstein, A. (2006). "The K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation". *IEEE Transactions on Signal Processing*, vol. 54, no. 11: 4311 –4322. doi: 10. 1109/ TSP. 2006. 881199.

Casciati, F., and Giordano, M. (2010). *Structural Health Monitoring 2010: Proceedings of the Fifth European Workshop on Structural Health Monitoring*, Sorrento, Italy. DEStech Publications, Inc.

Cerda, F., Garrett, J., Bielak, J., Bhagavatula, R. and Kovačević, J. (2010). "Exploring Indirect Vehicle-Bridge interaction for SHM". In *Proc. of the Fifth International Conference on Bridge Maintenance, Safety and Management*, IABMAS2010, Philadelphia, USA, 696-702.

Cerda, F. (2012). *Indirect Structural health Monitoring of Bridges: Mathematical Models and Laboratory Experiments.* Ph.D Dissertation, Carnegie Mellon University. Pittsburgh, PA.

Cerda, F., Chen, S., Bielak, J., Garrett, J., Rizzo, P., Kovačević, J. (2013). "Indirect Structural Health Monitoring Of A Simplified Laboratory-scale Bridge Model". Submitted to *Smart Structures and Systems (Special Issue: Challenge on bridge health monitoring utilizing vehicle-induced vibrations.),* in review.

Chang, F. (2011). *Structural Health Monitoring 2011: Condition-Based Maintenance and Intelligent Structures. Proceedings of the Eighth International Workshop on Structural Health Monitoring.* DEStech Pub.

Chen, S., Billings, S. A., and Luo, W. (1989). "Orthogonal Least Squares Methods And Their Application to Non-linear System Identification". *Int. J. Contr.*, vol. 50, no. 5, pp. 1873–96.

Doebling, S. W., Farrar, C. R., Prime, M. B., Shevitz, D. W. (1996). "Damage Identification and Health Monitoring of Structural and Mechanical Systems From Changes in Their Vibration Characteristics: A Literature Review". *Los Alamos National Laboratory Report.* Report Number: LA-13070-MS.

Duda, R., Hart, P., and Stork, D. (2000). *Pattern Classification* (2nd Edition). Wiley-Interscience.

Engan, K., Aase, S. O., and Hakon-Husoy, J. H. (1999). "Method of Optimal Directions for Frame Design". In *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, vol. 5, pp. 2443–2446.

Frangopol, D., Sause, R., and Kusko, C. S. (2010). *Bridge Maintenance, Safety and Management - IABMAS'10: in Proc. of the Fifth International IABMAS Conference* (1st Edition). CRC Press. Philadelphia, USA.

Lin, C.W., and Yang, Y.B. (2005). "Use of a passing vehicle to scan the fundamental bridge frequencies: An experimental verification". *Engineering Structures* 27, no. 13 (November): 1865-1878.

Macqueen, J. B. (1967). "Some Methods for Classification and Analysis of Multivariate Observations". 1:281–297. *University of California Press.*

Wright, J., Yang, A. Y., Ganesh, A., Sastry, S. S. and Ma, Y. "Robust Face Recognition via Sparse Representation." *IEEE Transactions on Pattern Analysis and Machine Intelligence.* Vol. 31, no. 2: 210–227. doi: 10. 1109/ TPAMI. 2008. 79.

MULTIRESOLUTION CLASSIFICATION WITH SEMI-SUPERVISED LEARNING FOR INDIRECT BRIDGE STRUCTURAL HEALTH MONITORING

Siheng Chen[1,2], *Fernando Cerda*[3,5], *Jia Guo*[2], *Joel B. Harley*[1], *Qing Shi*[1], *Piervincenzo Rizzo*[4], *Jacobo Bielak*[3], *James H. Garrett*[3] *and Jelena Kovačević*[1,2]

[1]Dept. of ECE, [2]Dept. of BME and Center for Bioimage Informatics
[3]Dept. of CEE, Carnegie Mellon University, Pittsburgh, PA, USA
[4]Dept. of CEE, University of Pittsburgh, Pittsburgh, PA, USA
[5]Dept. Ingeniería Civil, Universidad de Concepción, Concepción, Chile

ABSTRACT

We present a multiresolution classification framework with semi-supervised learning for the indirect structural health monitoring of bridges. The monitoring approach envisions a sensing system embedded into a moving vehicle traveling across the bridge of interest to measure the modal characteristics of the bridge. To enhance the reliability of the sensing system, we use a semi-supervised learning algorithm and a semi-supervised weighting algorithm within a multiresolution classification framework. We show that the proposed algorithm performs significantly better than supervised multiresolution classification.

Index Terms— multiresolution classification, semi-supervised learning, bridge structural health monitoring

1. INTRODUCTION

Bridge structural health monitoring (SHM) has been an intense research area for some time. Traditional, direct approaches, are to collect acceleration signals by installing sensors on a bridge. The drawback of such direct approaches is that they require a sophisticated and expensive electronic infrastructure with installation, maintenance and power support. Recently, indirect approaches have been proposed [1, 2, 3], suggesting the use of moving vehicles to collect data from accelerometers inside the vehicles, a far less expensive and complex solution (see Figure 1).

In indirect approaches less data is collected and the data is noisier (as it is farther from the source); thus, data analysis plays a crucial role. Moreover, although it is easy to get a large number of data samples, it is expensive to label them (which involves physically inspecting the bridge and determining its health); thus, very few data samples are actually labeled. This real-world constraint turns the indirect bridge SHM into a semi-supervised classification problem.

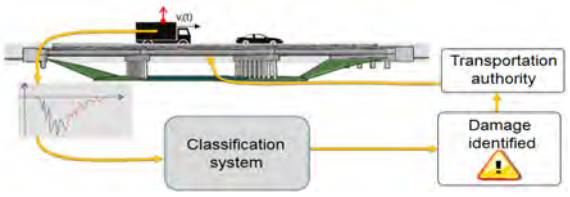

Fig. 1: Indirect bridge SHM system. Acceleration signals are collected from a moving vehicle and sent to a classification system, which identifies the bridge status and reports it to a transportation authority.

We propose a novel semi-supervised classification framework that takes advantage of supervised multiresolution classification (MRC) [4], which extracts hidden features in localized time-frequency regions (subbands), and a semi-supervised learning algorithm [5], which uses both labeled and unlabeled samples. This is followed by a semi-supervised weighting algorithm that combines information from all the subbands of all the signals to make a global decision.

The outline of the paper as follows: Section 2 states the problem and gives a brief overview of MRC and semi-supervised learning; Section 3 describes our proposed algorithm, which is validated in Section 4 on acceleration signals collected from a lab-scale bridge-vehicle dynamic system. Section 5 concludes with pointers to future directions.

2. BACKGROUND AND PROBLEM FORMULATION

Classification. Our task in bridge SHM is to label acceleration signals as belonging to different classes of structural change or damage, a task known as classification [6]. Let $\mathcal{X} = \{x^{(i)} \in \mathbb{R}^N\}_{i=1}^n$ be the given dataset with n signals, l labeled and u unlabeled; $\mathcal{Y} = \{y^{(i)} \in \{1, 2, \ldots, C\}\}_{i=1}^{\ell}$

The authors gratefully acknowledge support from the NSF through awards 1130616 and 1017278, as well as CMU Carnegie Institute of Technology Infrastructure Award.

i	sample index	
$\mathcal{X} = \{x^{(i)}\}$	input dataset	$i = 1, \ldots, n$
$\mathcal{Y} = \{y^{(i)}\}$	ground-truth labels for $\mathcal{L}$	$i = 1, \ldots, \ell$
$\mathcal{L} = \{(x^{(i)}, y^{(i)})\}$	labeled dataset	$i = 1, \ldots, \ell$
$\mathcal{U} = \{x^{(i)}\}$	unlabeled dataset	$i = \ell+1, \ldots, n$
$\mathcal{F}$	feature extraction function	
$f^{(i)}$	feature vector	$i = 1, \ldots, n$
$q^{(i)}$	ground-truth vector	$i = 1, \ldots, \ell$
$\hat{y}^{(i)}$	estimated label	$i = \ell+1, \ldots, n$
$\hat{q}^{(i)}$	confidence vector	$i = \ell+1, \ldots, n$
$\hat{\mathcal{Y}} = \{\hat{y}^{(i)}\}$	estimated labels for $\mathcal{U}$	$i = \ell+1, \ldots, n$

Table 1: Parameters used in a generic classification system.

the *ground-truth labels* for the *labeled dataset* $\mathcal{L} = \{(x^{(i)} \in \mathcal{X}, y^{(i)} \in \mathcal{Y})\}_{i=1}^{\ell}$; and $\mathcal{U} = \{x^{(i)} \in \mathcal{X}\}_{i=\ell+1}^{n}$ the *unlabeled dataset*. Then, the problem can be formulated as designing a map that associates an input signal to a class label with a certain probability. Typically, a generic classification system will have an intermediate block between the two, a *feature extractor* $\mathcal{F}$ *(FE)*, aimed at reducing the dimensionality of the problem; this is followed by a *classifier* $\mathcal{C}$. The outputs of the classifier are the *estimated labels* $\hat{\mathcal{Y}} = \{\hat{y}^{(i)} \in \{1, 2, \ldots, C\}\}_{i=\ell+1}^{n}$ for the unlabeled dataset $\mathcal{U}$. If the classifier is supervised, we denote the block by SC (see Figure 2).

Note that a label can also be viewed as a posterior probability vector $\hat{q}$ of size $C \times 1$, where the cth component of the vector, $\hat{q}(c)$, is the probability that a sample belongs to the cth class. Since $\hat{q}$ gives a confidence to an assigned label, we name it a *confidence vector*. The confidence vector for a labeled sample is called the *ground-truth vector*, q.

MRC. MRC is a supervised classification framework (see Figure 2), originally proposed for bioimaging applications [4, 7, 8]. It decomposes images into S localized space-frequency subbands using wavelet packets, a data-adaptive MR technique [9]. In each subband, MRC extracts features, classifies them, and produces a local classification decision. A supervised weighting algorithm combines all local decisions into a global decision (see Algorithm 1).

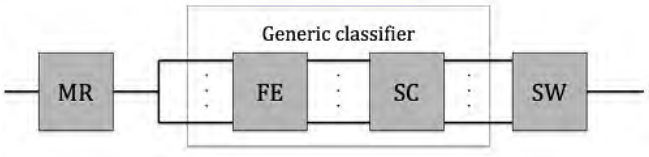

Fig. 2: Supervised MRC decomposes images into localized space-frequency subbands using wavelet packets (MR), followed by feature extraction (FE) and supervised classification (SC) in each subband, yielding a local classification decision. A supervised weighting algorithm (SW) combines all local decisions into a global decision.

Algorithm 1 MRC

Input	$\mathcal{X} = \{x^{(i)}\}$	input dataset	
Output	$\hat{\mathcal{Y}} = \{\hat{y}^{(i)}\}$	estimated labels for $\mathcal{X}$	
	s	subband index	$s = 1, 2, \ldots, S$
	$\mathcal{D}_s$	MR function	
	$a_s^{(i)}$	MR coefficients	
	$f_s^{(i)}$	feature vector	
	$\mathcal{C}_s$	supervised classification function	
	$\hat{q}_s^{(i)}$	confidence vector	
	$\mathcal{W}$	supervised weighting function	
	w	weighting vector	$\|w\|_1 = 1$
	$\hat{q}^{(i)}$	confidence vector after weighting	
MRC$(\mathcal{X})$			
	MR	$a_s^{(i)} = \mathcal{D}_s(x^{(i)})$	
	FE	$f_s^{(i)} = \mathcal{F}(a_s^{(i)})$	
	SC	$\hat{q}_s^{(i)} = \mathcal{C}_s(f_s^{(i)})$	
	SW	$\hat{q}^{(i)} = \mathcal{W}([\hat{q}_1^{(i)}, \hat{q}_2^{(i)}, \ldots, \hat{q}_s^{(i)}])$	
		$\hat{y}^{(i)} = \arg\max_c \{\hat{q}^{(i)}(c)\}$	
	return	$\hat{\mathcal{Y}}$	

(Only parameters different from Table 1 are listed.)

MRC provides various options: the choice of the filter bank $\mathcal{D}$ used in the MR block [10], the feature extraction method $\mathcal{F}$ used in the FE block, and the supervised classifier $\mathcal{C}$ used in the SC block. In the sth subband, the function producing the filter bank output is denoted by $\mathcal{D}_s$ and the supervised classifier by $\mathcal{C}_s$ (different classification boundaries in different subbands). To combine the subbands' classification decisions, we collect subbands' individual confidence vectors $\hat{q}_s^{(i)}$ into a $C \times S$ confidence matrix $\hat{Q}^{(i)}$, and define the weighting function $\mathcal{W}$ as

$$\hat{q}^{(i)} \;=\; \mathcal{W}(\hat{Q}^{(i)}) \;=\; \hat{Q}^{(i)} w,$$

where the weighting vector w, which assigns weight to each subband according to its discriminative power, is chosen by optimizing a supervised weighting objective function

$$w \;=\; \arg\min_{\omega} \{ \sum_{i=1}^{\ell} \|q^{(i)} - \hat{Q}^{(i)} \omega\| \}. \qquad (1)$$

The optimization is performed over all labeled data samples with the constraint that $\|w\|_1 = 1$.

Semi-Supervised Learning. Semi-supervised learning is a technique for training classifiers with both labeled and unlabeled data that assumes that unlabeled data can provide distribution information to build a stronger classifier. It includes generative mixture models with expectation maximization, co-training, transductive support vector machine and graph-based approaches [5]; we focus here on label propagation [11, 12], one of graph-based approaches. It assumes that, while the measured samples exist in a high-dimensional space, they are distributed in a low-dimensional manifold. Based on this, a graph is constructed to analyze the distribution of all samples; by understanding how labels propagate on this graph, classification can be achieved (see Algorithm 2).

Algorithm 2 Label propagation

Input	$\mathcal{X} = \{x^{(i)}\}$	input dataset
Output	$\hat{\mathcal{Y}} = \{\hat{y}^{(i)}\}$	estimated labels for $\mathcal{X}$
	A	adjacency matrix
	ρ	local measurement
	σ	scaling coefficient
	P	transition matrix
	t	transition time
Label propagation$(\mathcal{X})$		
	Construction	$A_{ij} = \exp\left(-\rho(x^{(i)}, x^{(j)})/\sigma\right)$
	Normalization	$P = D^{-1}A, \quad D_{ii} = \sum_j A_{ij}$
	Initialization	$\hat{Q}$, with $\hat{Q}_{y^{(i)},i} = 1, \quad i = 1, \ldots, \ell$
	Diffusion	$\hat{Q} \leftarrow \hat{Q}P^t$
	Labeling	$\hat{y}^{(i)} = \arg\max_c \{\hat{q}^{(i)}(c)\}$
	return	$\hat{\mathcal{Y}}$

(Only parameters different from Algorithm 1 are listed.)

Algorithm 3 Semi-supervised MRC

Input	$\mathcal{X} = \{x^{(i)}\}$	input dataset
Output	$\hat{\mathcal{Y}} = \{\hat{y}^{(i)}\}$	estimated labels for $\mathcal{X}$
	$\mathcal{C}$	semi-supervised classification function
	$\mathcal{W}$	semi-supervised weighting function
SSMRC$(\mathcal{X})$		
	MR	$a_s^{(i)} = \mathcal{D}_s(x^{(i)})$
	FE	$f_s^{(i)} = \mathcal{F}(a_s^{(i)})$
	SSC	$\hat{q}_s^{(i)} = \mathcal{C}_s(f_s^{(i)})$
	SSW	$\hat{q}^{(i)} = \mathcal{W}([\hat{q}_1^{(i)}, \hat{q}_2^{(i)}, \ldots, \hat{q}_s^{(i)}])$
		$\hat{y}^{(i)} = \arg\max_c \{\hat{q}^{(i)}(c)\}$
	return	$\hat{\mathcal{Y}}$

(Only parameters different from Algorithm 1 are listed.)

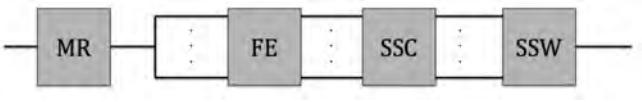

Fig. 3: Semi-supervised MRC. Classification and weighting algorithm in Figure 2 are replaced with their semi-supervised versions so that unlabeled data can contribute to classification.

3. PROPOSED ALGORITHM

MRC analyzes data to uncover hidden information; in its original form, it uses supervised classification, and can thus train on labeled samples only. When the labeled set is small or contains improperly labeled samples, the classification boundary and the weights assigned to subbands can be unreliable. Semi-supervised learning, on the other hand, uses the entire dataset to help classification, but works on one resolution level only. We thus propose to merge these two frameworks and gain the best of both worlds: a semi-supervised MRC.

Semi-Supervised Classification. The MR and FE blocks from Figure 2 work as before. The first change is that the supervised classifier block, SC, is replaced by a semi-supervised one, SSC (see Figure 3), using both labeled and unlabeled samples to make a labeling decision in each subband.

Semi-Supervised Weighting. We now explain how to weigh decisions from all the subbands to get a global decision in a semi-supervised manner. Labeled samples contribute to weighting directly by fitting their confidence vectors from all the subbands to the ground truth; unlabeled samples cannot do the same as they do not have the ground truth. We could use Shannon entropy to measure the confidence of labeling an unlabeled sample; if the entropy is small (less uncertainty, high confidence), it is easy to assign a label to the sample, and vice versa. In the label propagation algorithm, Algorithm 2, we label each sample by finding the largest element in its confidence vector; we could thus normalize each confidence vector to sum to 1 to measure its entropy. We encounter a problem, however; for example, let $\hat{q}^{(1)} = \begin{bmatrix}0.5 & 0.5 & 0\end{bmatrix}^T$ and $\hat{q}^{(2)} = \begin{bmatrix}0.5 & 0.25, & 0.25\end{bmatrix}^T$ be confidence vectors. While we can label $\hat{q}^{(2)}$ as Class 1 but cannot make a decision for $\hat{q}^{(1)}$, the entropy measure tells us that we can label $\hat{q}^{(1)}$ with higher confidence (less uncertainty) because its entropy is lower. To resolve this issue, we define a new uncertainty measure,

$$M(\hat{q}) \;=\; H(\hat{q})\,(\chi_{d>T} + \lambda(d)\,\chi_{d\leq T}),$$

where $H(\hat{q})$ is the entropy of confidence vector $\hat{q}$, χ_I is the indicator function of an interval I, $d = |\hat{q}_{(1)} - \hat{q}_{(2)}|$ with $\hat{q}_{(1)}, \hat{q}_{(2)}$ the first and second largest element in $\hat{q}$, respectively, T is the threshold, and $\lambda(d)$ is a penalty function that is large when the first and second largest elements are close. Let $M_s^{(i)}$ be the uncertainty of the sth subband to label the ith sample. Since entropy is additive, the total uncertainty of a subband when classifying unlabeled samples is the mean uncertainty over all the unlabeled samples in this subband. Thus, the total uncertainty of the sth subband is

$$m_s \;=\; \frac{1}{u}\sum_{i=l+1}^{l+u} M_s^{(i)}.$$

Define the normalized confidence of the sth subband as

$$g_s \;=\; \frac{e^{-\beta m_s}}{\sum_{k=1}^{S} e^{-\beta m_k}},$$

where β is the decaying coefficient. When the uncertainty of a subband is large, the confidence is small and the subband gets assigned a low weight, and vice versa. This confidence is the discriminative power of a given subband.

To combine confidences of all subbands into a global decision, we choose the weighting vector by optimizing a semi-supervised weighting objective function,

$$\begin{aligned} w \quad &= \quad \arg\min_{\omega} \{\frac{\alpha}{\ell}\sum_{i=1}^{\ell} \|q^{(i)} - \hat{Q}^{(i)}\omega\| \\ &+ \quad (1-\alpha)\|\omega - g\|\}, \end{aligned} \tag{2}$$

where w is as defined before and α is the labeling ratio defined as $\alpha = \ell/(\ell+u)$. The first term in (2) is exactly (1), representing the contribution from labeled samples. In the second term, we fit weights to subbands' confidences, representing the contribution from the unlabeled samples. We use the labeling ratio to balance these two terms; when we have a large number of labeled samples, the first term dominates, otherwise, the second one does. Since this is a convex optimization problem, it is numerically efficient to solve. After getting the weights, we can get the global decision as $\hat{y}^{(i)} = \arg\max_c \hat{q}^{(i)}(c)$, where $\hat{q}^{(i)} = \hat{Q}^{(i)} w$ (see Algorithm 3).

4. EXPERIMENTAL RESULTS

Dataset. We built a lab-scale bridge-vehicle dynamic system. put a sensor on a vehicle, and let it move across the bridge. We collected 30 samples for each of 13 different bridge damage scenarios, 8 different speeds and 2 different vehicles [13].

Experimental Setup. Given a specific vehicle driven at a specific speed, we want to classify 13 scenarios, in particular with a low labeling ratio. We consider 16 vehicle-speed cases for each of which there are 30 samples per 13 of the scenarios, and vary the labeling ratio as 10%, 30%, 50%, 70% and 90%; the final accuracy is the average over the 13 scenarios. We compare the performance of our proposed algorithm to the supervised MRC as well as the label propagation algorithm. We choose a Coiflet filter bank [14] with 4 levels in the MR block, principal component analysis in the FE block, naive Bayes, logistic regression and radius kernel SVM in the SC block [15], and label propagation in the SSC block. For label propagation, we choose the local measurement ρ to be the cosine distance, scaling coefficient $\sigma = 1$, transition time $t = 8$. For computational efficiency, we construct a $k = 4$ regular graph (each vertex connects to 4 neighbors). In semi-supervised weighting function, we choose penalty threshold $T = 0.02$ and the penalty term $\lambda(d) = 1 + 5(d/T - 1)^2$. We performed a 30-fold cross-validation and found that parameters do not influence the results too much.

Results. Table 2 compares the performance of different classifiers with the low labeling ratio of 10%. We use V for vehicle, S for speed, SMRC for supervised MRC, LR for logistic regression, NB for naive Bayes, KSVM for kernel SVM, LP for label propagation, SSMRC for semi-supervised MRC and LP-W for label propagation with a semi-supervised weighting algorithm in (2). We see that when the labeling ratio is low, supervised MRC performs poorly, label propagation works well, and semi-supervised MRC works the best.

Figure 4 shows the dependence of classification accuracy on the labeling ratio for 2 vehicles averaged across 8 speeds. Both figures show similar trends; as the labeling ratio decreases, accuracy drops sharply for all algorithms except for semi-supervised MRC, which consistently outperforms them all and whose performance stays relatively flat even at very low labeling ratios.

V	S	SMRC LR	SMRC NB	KSVM	LP	SSMRC LP-W
1	1	47.8	62.4	84.1	81.3	**99.8**
	2	57.0	62.9	84.5	86.2	**99.9**
	3	52.4	60.9	84.7	86.0	**99.4**
	4	66.2	63.2	89.1	90.8	**99.9**
	5	48.0	46.5	81.8	85.2	**94.5**
	6	34.2	45.7	74.8	86.0	**93.5**
	7	37.4	46.8	66.0	69.0	**72.2**
	8	38.8	43.4	59.0	75.9	**82.5**
2	1	40.9	58.5	76.2	75.6	**85.9**
	2	37.7	57.7	61.6	68.3	**80.5**
	3	58.1	65.6	81.3	81.7	**94.7**
	4	46.7	56.2	73.1	80.0	**87.3**
	5	47.9	59.0	72.5	76.5	**88.1**
	6	44.4	54.5	73.3	78.6	**83.8**
	7	48.3	63.7	76.9	83.3	**88.2**
	8	54.2	62.3	79.1	90.4	**93.8**

Table 2: Accuracy comparison of Vehicles (V) 1 and 2, with Speeds (S) 1, 2, ..., 8, and labeling ratio of 10%.

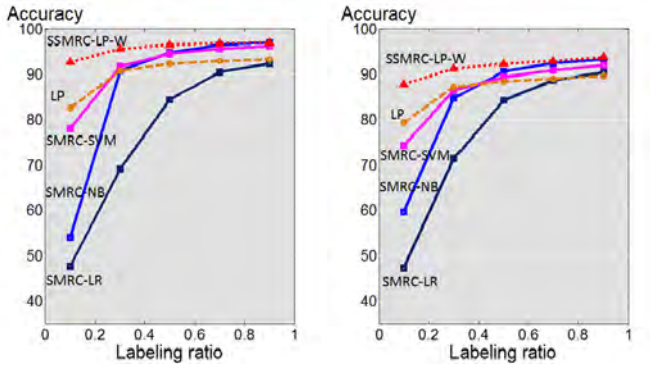

(a) Vehicle 1. (b) Vehicle 2.

Fig. 4: Accuracy as a function of the labeling ratio.

5. CONCLUSIONS AND FUTURE WORK

We presented a framework and algorithm for indirect bridge SHM. Compared to our previous work [2, 3], we added two new ingredients: (1) We analyze signals in MR spaces, instead of only in frequency domain. (2) We focus on semi-supervised learning setting, instead of supervised learning. Our system combines MR techniques and semi-supervised learning by using a semi-supervised weighting algorithm. The new system performs consistently better than supervised MRC, and significantly better when the labeling ratio is small.

Some near-future tasks are to use more features in each time-frequency subband, prune wavelet packet tree to achieve faster implementation, use stronger semi-supervised classifiers and test the framework on real-world bridge-vehicle dynamic system.

6. REFERENCES

[1] Y. B. Yanga, C. W. Lina, and J. D. Yaub, "Extracting bridge frequencies from the dynamic response of a passing vehicle," *Journ. of Sound and Vibration*, pp. 471–493, May 2004.

[2] F. Cerda, J. Garrett, J. Bielak, R. Bhagavatula, and J. Kovačević, "Exploring indirect vehicle-bridge interaction for bridge SHM," in *Proc. Int. Conf. Bridge Maintenance, Safety and Management*, Philadelphia, PA, July 2010, pp. 696–702.

[3] F. Cerda, J. Garrett, J. Bielak, P. Rizzo, J. A. Barrera, Z. Zhang, S. Chen, M. McCann, and J. Kovačević, "Indirect structural health monitoring in bridges: scale experiments," in *Proc. Int. Conf. Bridge Maintenance, Safety and Management*, Lago di Como, Italy, July 2012, pp. 346–353.

[4] A. Chebira, Y. Barbotin, C. Jackson, T. E. Merryman, G. Srinivasa, R. F. Murphy, and J. Kovačević, "A multiresolution approach to automated classification of protein subcellular location images," *BMC Bioinformatics*, vol. 8, no. 210, 2007.

[5] X. Zhu, "Semi-supervised learning literature survey," Tech. Rep. 1530, Univ. Wisconsin-Madison, 2005.

[6] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, John Wiley & Sons, Englewood Cliffs, NJ, 2001.

[7] R. A. Kellogg, A. Chebira, A. Goyal, P. A. Cuadra, S. F. Zappe, J. S. Minden, and J. Kovačević, "Towards an image analysis toolbox for high-throughput Drosophila embryo RNAi screens," in *Proc. IEEE Int. Symp. Biomed. Imaging*, Arlington, VA, Apr. 2007, pp. 288–291.

[8] A. Chebira, J. A. Ozolek, C. A. Castro, W. G. Jenkinson, M. Gore, R. Bhagavatula, I. Khaimovich, S. E. Ormon, C. S. Navara, M. Sukhwani, K. E. Orwig, A. Ben-Yehudah, G. Schatten, G. K. Rohde, and J. Kovačević, "Multiresolution identification of germ layer components in teratomas derived from human and nonhuman primate embryonic stem cells," in *Proc. IEEE Int. Symp. Biomed. Imaging*, Paris, France, May 2008, pp. 979–982.

[9] R. R. Coifman, Y. Meyer, S. Quake, and M. V. Wickerhauser, "Signal processing and compression with wavelet packets," Tech. Rep., Yale Univ., 1991.

[10] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*, Signal Processing. Prentice Hall, Englewood Cliffs, NJ, 1995, http://waveletsandsubbandcoding.org/.

[11] X. Zhu and Z. Ghahramani, "Learning from labeled and unlabeled data with label propagation," Tech. Rep., Carnegie Mellon Univ., 2002.

[12] A. D. Szlam, R. R. Coifman, and M. Maggioni, "A general framework for adaptive regularization based on diffusion processes," *Journ. Mach. Learn. Res.*, , no. 9, pp. 1711–1739, Aug. 2008.

[13] F. Cerda, S. Chen, J. Bielak, J. H. Garrett, P. Rizzo, and J. Kovačević, "Indirect structural health monitoring of a simplified laboratory-scale bridge model," *Int. Journ. Smart Struct. Syst.*, 2012, Submitted.

[14] I. Daubechies, "Orthonormal bases of compactly supported wavelets II. Variations on a theme," *SIAM Journ. Math. Anal.*, vol. 24, no. 2, pp. 499–519, Mar. 1993.

[15] C. M. Bishop, *Pattern Recognition and Machine Learning*, Information Science and Statistics. Springer, 2006.

Indirect structural health monitoring in bridges: scale experiments

F. Cerda[1], J.Garrett [1], J. Bielak [1], P. Rizzo [2], J. Barrera [1], Z. Zhuang [1], S. Chen [1], M. McCann [1] & J. Kovačević [1]

[1] *Carnegie Mellon University, Pittsburgh, Pennsylvania, USA*

[2] *University of Pittsburgh, Pittsburgh, Pennsylvania, USA*

ABSTRACT:

In this paper, we use a scale model to experimentally validate an indirect approach to bridge structural health monitoring (SHM). In contrast to a traditional direct monitoring approach with sensors placed on a bridge, the indirect approach uses instrumented vehicles to collect data about the bridge. Indirect monitoring could offer a mobile, sustainable, and economical complementary solution to the traditional direct bridge SHM approach. Acceleration signals were collected from a vehicle and bridge system in a laboratory-scale experiment for four different bridge scenarios and five speeds. These signals were classified using a simple short-time Fourier transform technique meant to detect shifts in the fundamental frequency of the bridge due to changes in the bridge condition. Results show near-perfect detection of changes when this technique is applied to signals collected from the bridge (direct monitoring), and promising levels of detection when one uses signals from sensors on the vehicle (indirect monitoring) instead of those recorded on the bridge itself.

1 INTRODUCTION

In this paper, we explore whether the acceleration signals from vehicles moving over a bridge can be used for diagnostic purposes. The acceleration signals from the bridge and the vehicle are affected mainly by three different factors. These factors are: the dynamic properties of the bridge structure; the motion characteristics of the passing vehicle; and the dynamic properties from the vehicle.

Traditionally, the Structural Health Monitoring (SHM) community uses the data measured *directly* from a structural system for diagnostic purposes. In such an approach, a number of sensors are deployed on the structure. We refer to this as a *direct* SHM approach. In contrast, the use of data not recorded directly from a structure is referred to as an *indirect* SHM approach (Lin et al. 2005, Cerda et al. 2010). The top block in Figure 1 shows this distinction in terms of the data acquisition approach.

There are several practical reasons that drive our research on the *indirect* SHM approach. There is a need in most countries, and especially in the US, to monitor a large bridge stock in a reliable, objective and economically feasible way. The traditional *direct* SHM approach requires installation, power and maintenance of an expensive electronic infrastructure on top of the physical bridge infrastructure.

The indirect approach can gain leverage by using the equipment already located on board newer models of vehicles, or on a fleet of vehicles that can be equipped with sensors to collect the desired information as they undergo their daily routines.

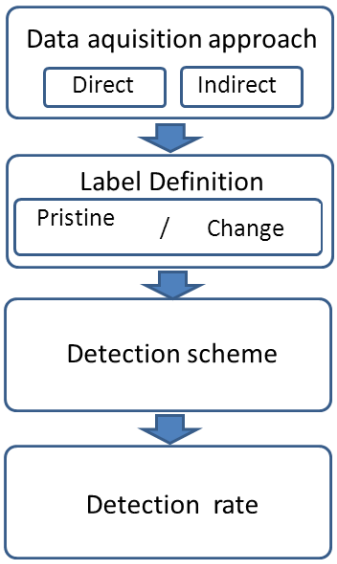

Figure 1. Block diagram of proposed system

The *indirect approach,* however, has the disadvantage that the data are influenced by the motion of the vehicle and its dynamic properties.

Through the use of an experimental setup which is described in the next section, we collected acceleration data from a particular scale bridge structure and a particular scale vehicle in order to compare the *indirect* and *direct* methods. These two data acquisition approaches are depicted in the first block of Figure 1 since both approaches for SHM were tested as part of the experiments described in this paper. As far we know, such comparisons have not been carried out previously.

To test the detection capability of the different approaches, we created four experimental scenarios. The first scenario considers the structure in a pristine while in each of the other three, we added mass to the bridge to simulate damage. Two different scenarios are compared in each test. A pristine one, and a second one with the induced change to modify its dynamic properties. The second block in Figure 1 shows the two scenarios considered.

The data from the two scenarios is classified by means of a short-time Fourier transform-based change detection scheme aimed at detecting changes in the fundamental frequency of a bridge. We report the results of this work in terms of the detection rate. This quantity reflects the fraction of cases in which an actual change is detected. These two steps are represented in the last two blocks in Figure 1.

Previous work includes a theoretical solution for the simplified case of a single degree of freedom oscillator travelling over a beam structure (Yang et al. 2004). In a subsequent paper, Lin et al. (2005) were able to experimentally determine the natural frequency of an actual bridge structure by analyzing the acceleration signals of a passing vehicle.

An experimental setup was used by Kim to simulate the vehicle-bridge interaction and identify damage scenarios (Kim et al. 2010). A particular methodology, referred as the "pseudo static approach", was used to identify damage using vibration data taken from the bridge structure at different locations along the bridge span (¼, ½ and ¾ of span length). This damage identification approach shows good accuracy at determining a change of stiffness of the bridge. Being inspired by this experimental work, we decided to further pursue the identification of changes in the bridge structure using the indirect approach.

In this paper, we concentrate on studying the influence of different vehicle velocities and sensor locations on the classification accuracy of different scenarios. The scenarios are produced in a laboratory using a scaled physical model of a moving vehicle over a simply supported bridge. The data is obtained through multiple runs of the vehicle over the structure. Hereafter, we refer to this particular experimental setup as the "scale bridge structure".

The following section contains a description of the experimental setup. This description includes the structural model, the vehicle model, the vehicle motion control system, and the data acquisition equipment. The third section contains a description of the different scenarios that were compared in the detection experiments. In the fourth section, we describe the Fourier transform-based change detection approach, and in the fifth section, we present and discuss preliminary results. Our initial conclusions are given in the last section.

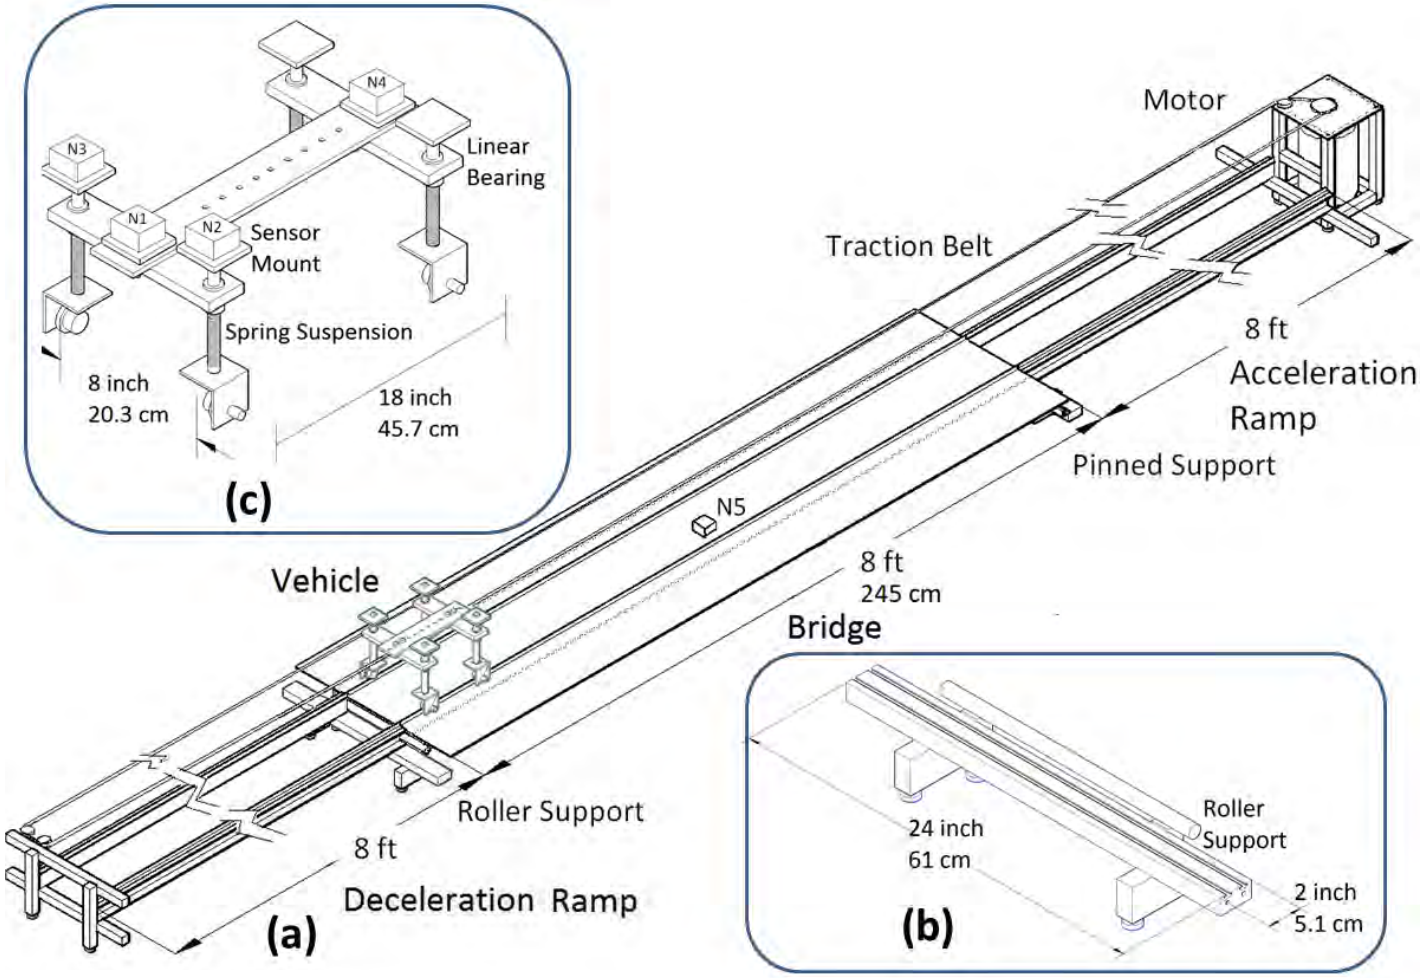

Figure 2. Structural components and vehicle from experimental setup

2 EXPERIMENTAL SETUP

The experimental setup simulates a passing vehicle over a simply supported bridge structure. The vehicle-bridge interaction is studied by recording accelerations at four different locations on the vehicle as well as at the midspan of the bridge. The whole system consists of several components: 1) the mechanical components consisting of the bridge, its approaches, and the vehicle, 2) the vehicle motion control system and 3) the data acquisition system.

2.1 *Mechanical components*

An overview of the mechanical components is given in Figure 2 (a). It consists of an acceleration ramp and a deceleration ramp that provide the running path for the vehicle, and a simply supported bridge. The acceleration and deceleration ramp are made from C Shape aluminum extrusions (2 x 1 x 1/8 in). They are supported on each end by aluminum slotted extrusions. The slotted extrusion shown in Figure 2 (b) allows one to fix the ramps at different locations along the slots. This flexibility will allow further research that will explore placing the ramps in the right or left lanes to study the effect of traffic-induced torsion.

2.1.1 *Vehicle*

The vehicle used on the experiment has two axles. A scheme of the vehicle is shown in Figure 2(c). The vehicle has four independent wheel suspensions. In this paper the vehicle properties are maintained at constant values shown in Table 1. The dynamic properties of the vehicle were obtained by capturing the dynamic response after an impulse force is applied.

Table 1. Vehicle properties

Properties with added mass	
Bouncing frequency	5 Hz
Front damping	5.9%
Rear damping	5.9%
Extra weight at midspan	5 lb)

The suspension system is designed so it can be easily modified to simulate different vehicle characteristics. For example, a heavily loaded 2 axle vehicle can be simulated by replacing the spring in the suspension shown in Figure 2(c) with a stiffer spring.

2.1.2 *Bridge*

The bridge structure is composed of an aluminum plate and two aluminum angles that act as beams. A cross-sectional view of the bridge is shown in Figure 3.

.

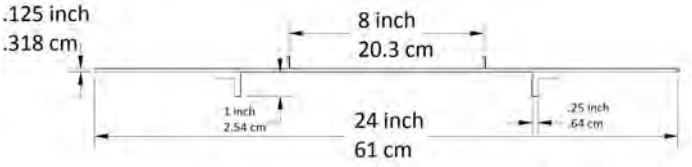

Figure 3. Bridge section

The bridge deck also has two angles on top that are used as rails for the vehicle. The properties of the bridge are shown in Table 2.

Table 2. Bridge properties

Deck dimensions	8 x 2 x 1/8 in
Beams dimensions	8 x 1 x 1/4 in
Fundamental frequency	7.18 Hz
Damping	1.35%

2.2 *Motion control*

A set of National Instrument® components was configured to reliably control the speed of the vehicle. The individual components are shown in Figure 4. The PXI 7342 motion controller commands the NI 70360 driver that provides the signals to a double shaft stepper motor, model NEMA 34. The encoder attached to the shaft of the motor provides position feedback and closes the loop for the motion control system.

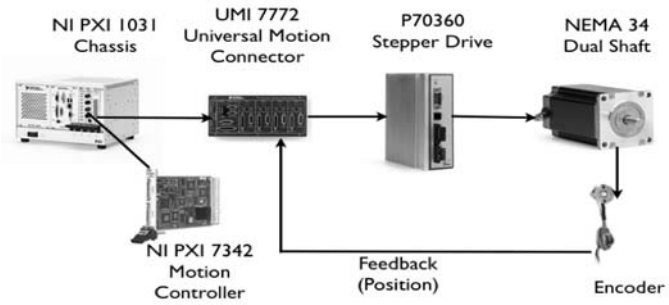

Figure 4.Motion control equipment scheme (Images from www.ni.com)

A list of the individual components is shown in the Table 3.

Table 3. Motion equipment components

Driver	NI 70360
Motor	NEMA 34
Motion Controller	NI PXI 7342
Interface	UMI 7772

2.3 *Data acquisition equipment*

The data acquisition system for the moving vehicle and bridge is wireless. Figure 5 shows the Microstrain® acceleration sensor nodes that communicate to a base directly connected to the PXI controller. The data is transmitted in packets to the base after digitalization. The resolution of the accel-

eration data is 1.5mg RMS. The resolution of the digitization is 12 bit.

The components used in the data acquisition process are described in Table 4.

Table 4. Data acquisition components

Acceleration nodes	MicroStrain G-Link mXrs 2G
Wireless base	MicroStrain WSDA mXrs

Figure 5. Data acquisition equipment

The sensor on the bridge is located at the center of the midspan, as depicted in Figure 2 (a) by the node labeled N5. The vehicle sensor locations are labeled as N1, N2, N3, and N4 in Figure 2 (c). Table 5 lists the sensor locations and the corresponding node names shown in Figure 2 (a) and (c).

Table 5. Node location

Node Name	Node location
N1	Front Suspension Front
N2	Left Wheel front
N3	Right Wheel front
N4	Rear Suspension
N5	Bridge

The nodes located at the top of the suspension shaft, N2 and N3, transmit the vertical motion at the wheel level through the suspension shaft.

3 EXPERIMENTAL SCENARIOS

In the work of Yang et al (2004), the authors derived an explicit analytical solution for the interaction of a simply supported beam with a traveling single degree of freedom oscillator. In this solution, the main interaction parameters are defined as S and μ: (1) $S = \pi v/L\omega_b$;(2) $\mu = \omega_b/\omega_v$, where S is a normalized vehicle velocity; v = vehicle velocity; ω_v = the vehicle (oscillator) vertical natural frequency; L = length of beam, and ω_b= fundamental natural frequency of the beam.

We explored the influence of these parameters by inducing changes to the bridge structure and running the moving vehicle at different travelling speeds over the bridge. The bridge was changed by adding mass at the midspan. Figure 6 shows the procedure by which different amounts of mass were clamped at the midspan of the bridge. For our structure, ω_b is the fundamental natural frequency of the complete bridge system.

The different conditions of the bridge and the corresponding changes that they produce in terms of the fundamental frequency of the bridge are summarized in Table 6.

Table 6. Bridge scenarios

Scenario	Total added mass (lbs, % of bridge mass)	fundamental freq (Hz)
1	0, 0%	7.18
2	6, 16%	6.28
3	10, 27%	5.93
4	14, 38%	5.57

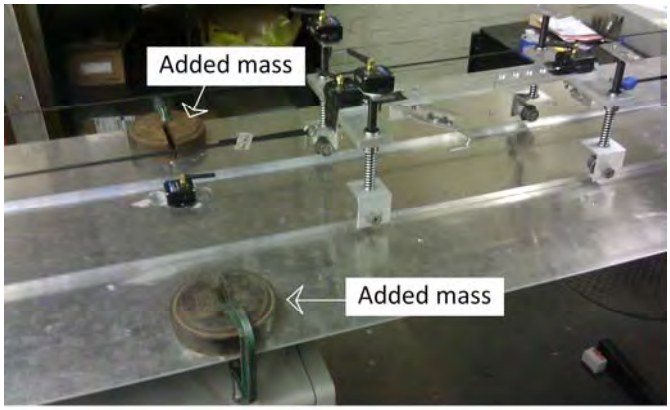

Figure 6. Added mass to the midspan of the experimental bridge structure.

The speed of the vehicle under the different bridge scenarios ranged from 1 m/s to 3 m/s at 0.5 m/s intervals. A total of 5 different speeds were studied.

The range of scenarios and speeds can be plotted in terms of the parameters S and μ as depicted in Figure 7.

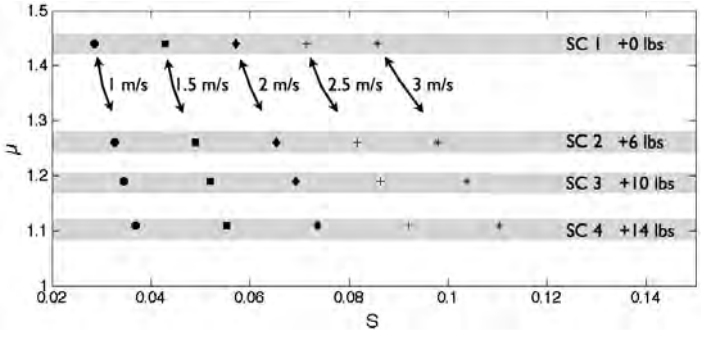

Figure 7. S and μ cases studied experimentally.

Each of the 20 dots in Figure 7 reflects a data set composed of 60 runs of the vehicle over the bridge. We compare the variation of the different data sets by performing detection experiments, as described in the next section.

4 DETECTION OF STRUCURAL SCENARIOS

The detection task is as follows. First, we use a training set of acceleration signals from a vehicle traveling over a known-to-be-healthy (pristine) scale bridge structure to set a baseline for the system. We then use a test set of different acceleration signals from the same bridge under different scenarios, to determine if the detection approach can detect whether the bridge has sustained a significant change since the collection of the training set. For the specific scenarios we are exploring, note that the induced change causes a decrease in the fundamental frequency of the bridge. We therefore hypothesize that a classification scheme based on detecting shifts in this frequency should work well.

Figure 8 shows a typical signal obtained through the experimental setting and the corresponding portion of the signal used for the classification, which is when the vehicle is completely between the supports of the bridge.

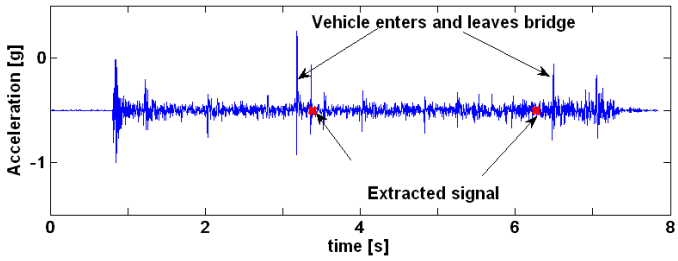

Figure 8 Original acceleration signal.

4.1 *STFT Calculation.*

To detect changes in the bridge, we first extract the frequency spectrum of the bridge/vehicle system from an acceleration signal. As seen in Figure 9, taking the Fourier transform of the acceleration signal results in a frequency spectrum with a large amount of noise and little consistency between runs; this is because the time-domain signal contains numerous spikes and other transient signals. We instead compute the spectrogram of the acceleration signal with a short-time Fourier transform (STFT) with overlapping windows that are 250 samples in length. The spectrogram is shown in Figure 10. We then average the spectrogram over time, creating a frequency spectrum as depicted in Figure 11. This technique exploits the fact that the frequencies of interest are not transient, while much of the noise is. The time averaging should therefore remove noise and preserve the signal.

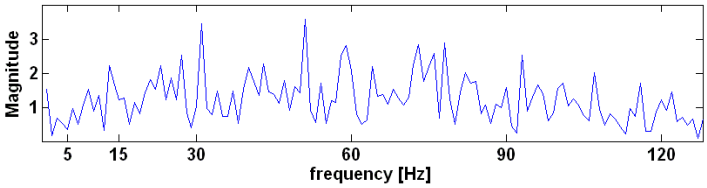

Figure 9. Fourier transform of an acceleration signal.

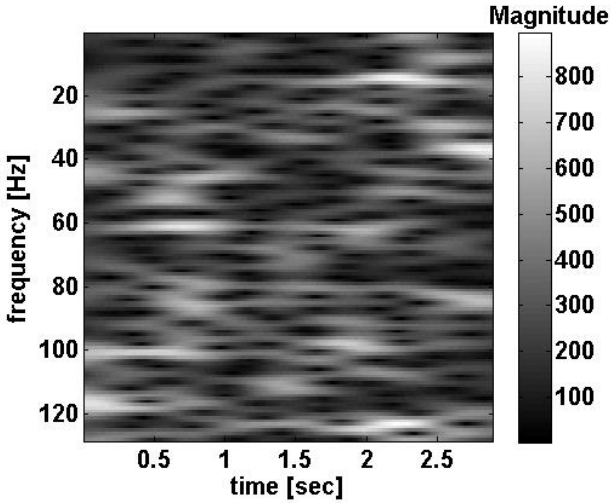

Figure 10. Acceleration spectrogram.

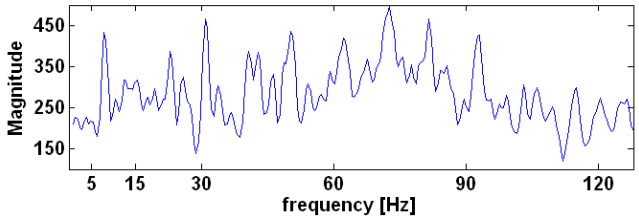

Figure 11. Spectrogram averaged across time.

4.2 *Classification System.*

For the training of our classification system, the data from each scenario at a particular speed is divided into training and testing sets. The spectrum of each signal in the training set is calculated first, then all such signals are averaged across frequency. The result is a single reference spectrum representing the undamaged bridge. To determine whether a testing set of signals indicates a change in the structure, the spectrum is computed for each signal in the testing set and again averaged across frequency to create a signal candidate spectrum for each scenario at a particular speed. The portion of each spectrum corresponding to the range 3.5-13 Hz is then extracted. The averaged candidate spectrum with the corresponding extracted section is shown in Figure 12.

The considered range is where our theoretical calculations suggest the fundamental frequency of the bridge should appear, and experimental explora-

tion confirms it. Finally, the trimmed reference and candidate spectra are shifted to be zero mean and their cross-correlation is computed. If the maximum value of the cross-correlation occurs when the spectra are not shifted, then the test set is labeled as unchanged. If the maximum value occurs at a shifted location, the test set is labeled as changed.

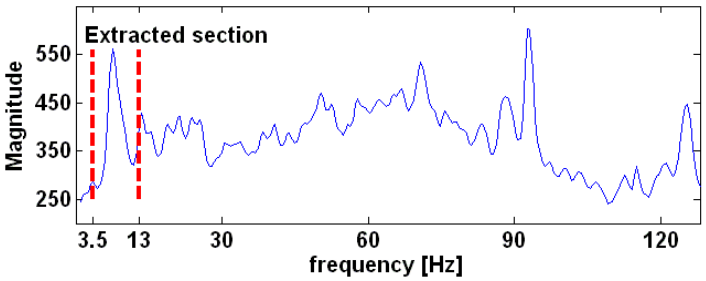

Figure 12. Extracted candidate spectrum of unchanged scenario.

The plots in Figure 13 show the correlation coefficients of two candidate spectrums compared against the unchanged reference spectrum. The correlation for the unchanged candidate has a strong peak at the zero shift location, indicating that it matches the reference spectrum. The peak in the correlation for the changed candidate spectrum is at a shifted location.

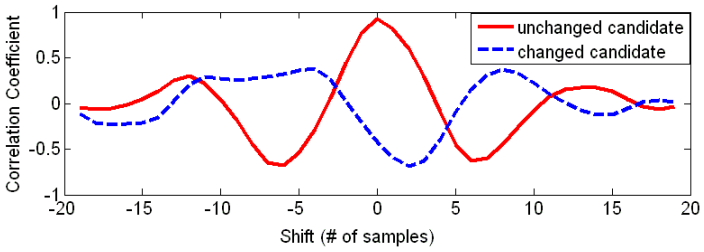

Figure 13. Cross correlation between candidates for the unchanged and changed scenario (SC3) with respect to the unchanged reference spectrum.

5 RESULTS AND DISCUSSION

The proposed detection approach was tested using a data set generated with the experimental setup previously described. We now show some of the preliminary findings of this approach and the corresponding results.

5.1 *Dataset.*

Our dataset consists of 60 acceleration signals from each of the five sensors (front suspension, back suspension, left wheel, right wheel, and bridge) collected under four different change scenarios (unchanged, +6 kg, +10 kg, +14 kg) and vehicle speeds (1 m/s, 1.5 m/s, 2 m/s 2.5 m/s 3 m/s). The signals were sampled at 256 Hz and vary in length from 1 to 3 seconds. The variation in signal length occurs because of the time it takes the vehicle to travel over the bridge at different speeds.

5.2 *Experimental Setup.*

To evaluate our classification system, we performed a series of cross-validation experiments. We first fixed a group size, N = 3, 4, …, 35. For each speed and each sensor, we randomly selected 20 signals from the undamaged bridge and used them as our training set. We then randomly selected N signals from each of the changed scenarios and used them to form test sets. Additionally, we selected N of the remaining 40 signals from the unchanged scenario and formed a test set with them. This random selection was repeated in a 1,000-fold validation.

For each scenario, we report the detection rate as the percentage of folds in which the test set was labeled as changed. For the three changed scenarios, the detection rate represents the true positive rate (TPR), while for the unchanged bridge, it represents the false positive rate (FPR). A perfect system would have a TPR of 1 for each damage condition and a FPR of 0.

5.3 *Results and Discussion.*

Figure 14 shows the damage detection rate for each scenario and sensor location plotted across speed, with N=35. The four lines represent the detection rate for the undamaged bridge and for each of the three damage scenarios. Ideally, we would see the FPR of 0 for the undamaged bridge, and the TPR of 1 for each of the damage scenario.

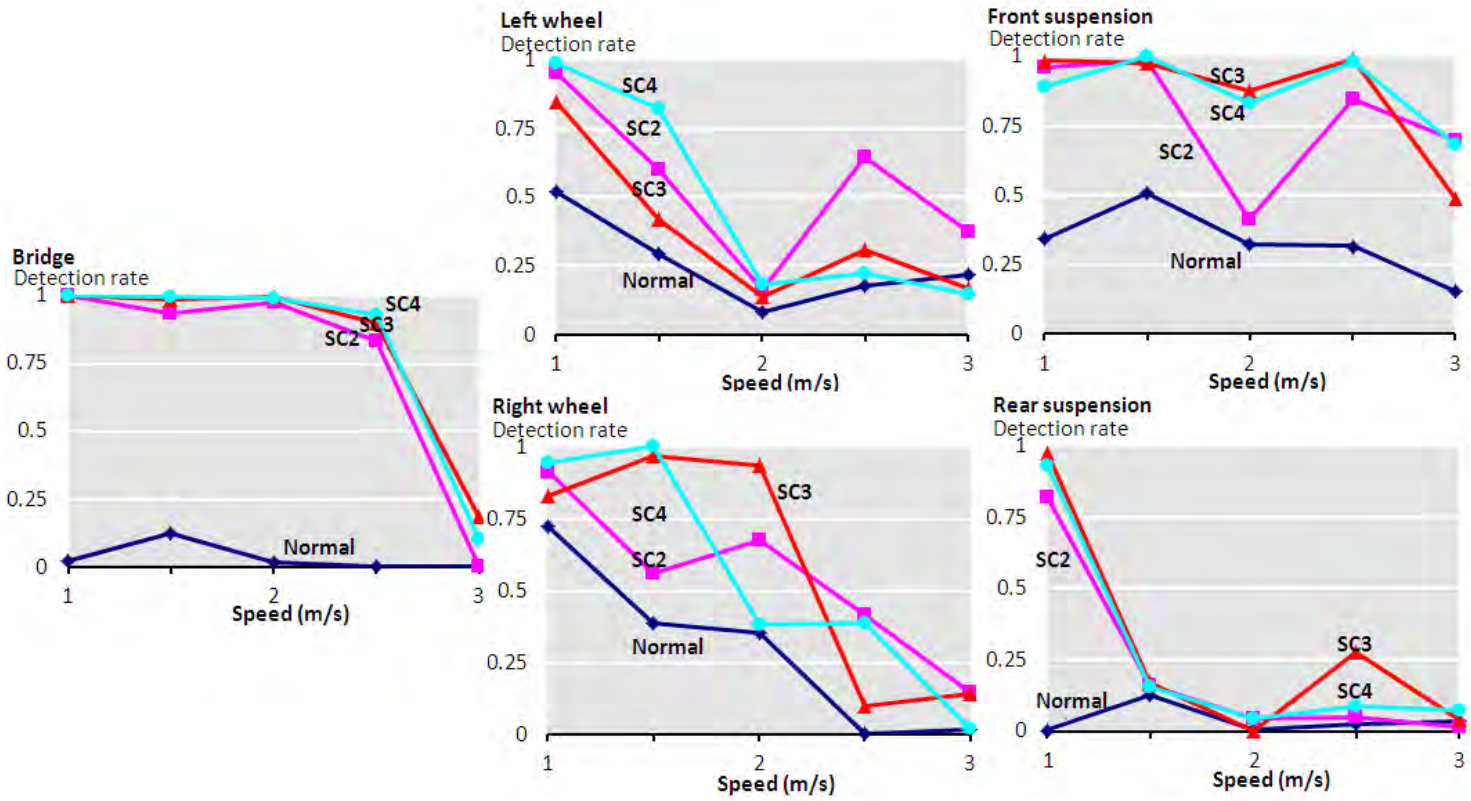

Figure 14. Detection rate for each scenario and sensor location plotted across speed, averaging 35 runs. For each damage scenario, the curve represents the true positive rate (TPR), while for the normal (undamaged) case, the curve represents the false positive rate (FPR). Accuracy is consistently good at the lowest speed, with the sensors on the suspension providing the best accuracy of the indirect sensors.

The lowest speed consistently produces the best detection for all sensors, while the accuracy for other speeds is inconsistent across sensors. This may be due to increased noise or shorter signal duration (which reduce the effectiveness of our spectrum averaging technique) for higher speeds. The sensor on the bridge detects damage nearly perfectly, validating our classification technique for the easier, direct-monitoring case. Of the indirect sensors, those on the suspension were better than those on the wheel. This is likely because the suspension acts as a low-pass filter, reducing the noise while preserving the low fundamental frequency of the bridge.

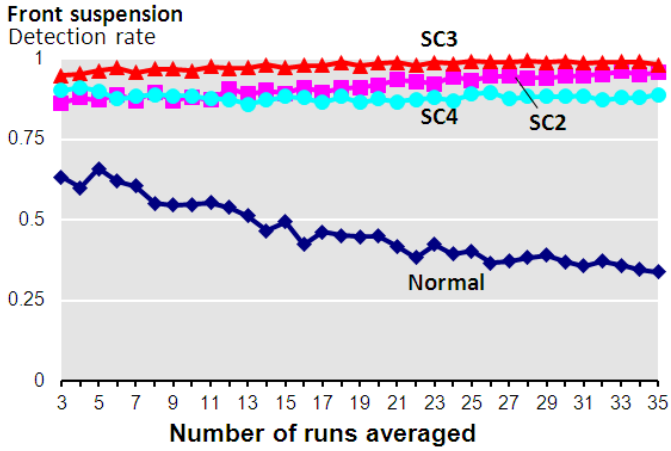

Figure 15.The effect of the number of averaged runs on the detection rate for the front suspension sensor at a speed of 1 m/s. Averaging a larger number of runs increases the TPR, while lowering the FPR.

Figure 15 shows the effect of N on the classification accuracy for the front suspension sensor at 1 m/s. In general, accuracy increases as N increases. For some sensors and speeds, there was a clear diminishing-returns effect, while for others there was not. Contrary to what we would expect, scenario 4 is not consistently the easiest to detect. Inspection of the spectra reveal that the peak for the fundamental frequency appears wider and shorter as more mass is added to the bridge, decreasing the accuracy of the correlation matching method.

6 CONCLUSIONS AND FUTURE WORK

This paper presents initial work for detecting changes in bridge structures based on acceleration data from passing vehicles. We refer to this approach as *indirect*. We compare the results of the *indirect* with the traditional *direct* approach in which sensors are located on the bridge structure.

An experimental setup that resembles a moving vehicle passing over a simply supported bridge was used to generate dynamic interaction data from several physical scenarios. The scenarios consisted of

changes in the mass of the bridge structure generated by adding a localized mass at midspan.

A detection procedure was developed to capture the shifts in the fundamental frequency of the bridge.

The detection capability of the proposed signal processing approach is more stable across different speeds for acceleration data gathered in a *direct* fashion rather than in the *indirect one.*

For the particular experimental setup used in this work and the different scenarios simulated, the sensor location on the vehicle has a strong influence in terms of the detection capability of the different scenarios. The sensor located at the front of the vehicle over the suspension system outperformed those of all other sensor locations.

Lower travelling speeds of the vehicle seem to be better for identifying changes in the natural frequency of the bridge than higher traveling speeds.

Regarding the number of runs averaged to calculate the true positive detection rate (TPR), there is a mild linear increasing trend over the quality of the detection. More significant is the reduction of false positive rate when increasing the number of runs.

Future work will include an exploration of the sensitivity of the approach to smaller changes in the bridge structural system.

We are also interested in examining the consistency of the TPR trends across different speeds. We will study this by populating our experimental data with smaller speed intervals.

We will also explore other types of changes that resemble damage scenarios in a real bridge structure. Such scenarios will consider frozen bearings and cracks. The first will be modeled by increasing the rotational restraint at the supports of the simply supported bridge. The latter will be simulated by a section reduction of the supporting beam elements of the bridge.

7 ACKNOWLEDGMENTS

This research was supported by a grant of the National Science Foundation (Award No. CMMI1130616) and by the Traffic 21 initiative at Carnegie Mellon University. We are grateful for this support. We also thank Ranny Zhao and Colin Rutenbar, who participated in the construction of the experimental setup.

8 REFERENCES

Cerda F., Garrett J., Bielak J., Bhagavatula R. & Kovačević J. 2010. Exploring Indirect Vehicle-Bridge interaction for SHM. *Proceedings of the Fifth International Conference on Bridge Maintenance, Safety and Management,* IABMAS2010, Philadelphia, USA, 696-702.

Kim, C.W. Kawatani, M. & Fujimoto T, (2010), Identifying bending stiffness change of a beam under a moving vehicle, *Proceedings of the Fifth International Conference on Bridge Maintenance, Safety and Management,* IABMAS2010, Philadelphia, USA, 761-767.

Lin, C.W., & Yang Y.B. 2005. Use of a passing vehicle to scan the fundamental bridge frequencies: An experimental verification. *Engineering Structures* 27, no. 13 (November): 1865-1878.

Yang, Y. B., Lin C. W., & Yau J. D. 2004. Extracting bridge frequencies from the dynamic response of a passing vehicle. *Journal of Sound and Vibration* 272, no.3-5(May 6):471-493.

Exploring Indirect Vehicle-Bridge Interaction for Bridge SHM.

F. Cerda[1], J.Garrett, J. Bielak, R. Bhagavatula & J. Kovačević
Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.

ABSTRACT: In this paper, we explore an indirect measurement approach for bridge structural health monitoring (SHM) that collects sensed information from the dynamic responses of many vehicles travelling over a bridge and then makes extensive use of advanced signal processing techniques to determine information about the state of the bridge. We refer to this approach as vehicle-data driven and indirect. We discuss some of the advantages of this indirect approach over direct monitoring of structures. We simplified the vehicle-bridge interaction and used a numerical oscillator-beam interaction model to generate some preliminary interaction response data with which to begin to assess the validity of this approach. A Multiresolution image classifier was used to analyze the preliminary data. We present the basic idea behind this approach and preliminary results that demonstrate its viability.

1 INTRODUCTION

1.1 *The Need for Structural Health Monitoring*

There are approximately 600,000 highway bridges within the U.S., and approximately 25 percent of them are currently rated as structurally deficient or functionally obsolete (Federal Highway Administration 2006)

Currently, bridges are inspected visually every two years. There is a strong interest to aid these inspection efforts with a more continuous, reliable, physics-based and less subjective procedure. This has led to a great deal of activity in structural health monitoring. Most of the current approaches consider data acquisition of bridges in a direct form, that is, by putting sensing devices at different specific locations on the structure. This poses a number of practical challenges, such as vandalism or involuntary damage of installed equipment, the need for a power source or complex energy harvesting, the initial and recurring costs associated with the monitoring system, and the need for extensive data processing and management at the bridge. Thus, there is an urgent need to explore alternative, more cost-effective means to monitor our complete stock of bridges on a regular basis.

1.2 *Overview of Proposed Approach*

In this paper, we describe a possible approach for performing structural health assessment that takes a markedly different tack. This approach is based on the collection of, and multiresolution pattern analysis of, data in the form of dynamic responses of vehicles passing over bridge structures. This approach can be considered as indirect, since it acquires information about the bridge from sensor-equipped vehicles moving over the bridge, as described by Lin & Yang 2005. Since some bridges over which such vehicles travel will be monitored by sensor systems installed on the bridge, direct data gathered from the bridge itself can be used for validation and calibration of the vehicle-based monitoring system. Thus, this indirect approach can also be considered as complementary to the direct approach. The vehicle-based approach will allow for much broader coverage of the entire bridge population, as only a fraction of the bridges will likely by sensed directly due to initial and long-term maintenance costs of the installed monitoring systems.

The data will be acquired from many passing vehicles (cars, buses and trucks) that are able to timestamp and locate themselves with respect to the bridge and make that data available for structural analysis. The data can then be processed and analyzed with advanced signal processing and

[1] Also affiliated with the Universidad de Concepción, Concepción, Chile.

pattern recognition techniques, including state-of-the-art multiresolution techniques such as wavelets, to identify the existence, location and severity of damage. The idea is to infer damage from changes in global and local properties of the bridge and its structural response characteristics that are present in the vehicle dynamic response. Such structural response characteristics include resonant frequencies, mode shapes, local deflections, etc.

2 PROPOSED APROACH

As described in the previous section, our approach is based on merging two main concepts: 1) the sensed data will be collected from many vehicles moving over the structure of interest; and 2) the data will be collected and processed using advance image multiresolution techniques. We now present a more detailed discussion of the added features of this approach with respect to a direct monitoring approach. We start by reviewing some of the literature on direct approaches, some preliminary research on indirect approaches, the advantages of using an indirect monitoring approach, damage identification based on moving loads, the practical advantages of having mobile monitoring and a description of the multiresolution classifier.

2.1 *Direct Approaches for SHM*

During the past two decades, structural health monitoring and damage assessment have been very active research areas, and have motivated several excellent review and overview papers, which highlight some of the most relevant approaches (e.g., Van der Auweraer & Peeters 2003; Farrar & Worden 2007). Brownjohn (2006) describes some general and fundamental objectives for monitoring civil infrastructure and points out some historical applications. More specific review topics include wireless, structural health monitoring, design of devices, and the trend for localized processing (Lynch, 2007); vibration-based condition monitoring (Doebling et al. 1998, Carden & Fanning 2004); damage identification using inverse methods (Friswell 2006); unsupervised learning (Fulgate et al. 2000, Worden & Dulieu-Barton 2004, Worden & Manson 2007); and vibration-based condition monitoring methods (Carden & Fanning 2004).

One of the widely used classifications for structural health damage identification is based on the level of detection attempted (Rytter 1993): Level 1: determine presence of damage; Level 2: determine location of the damage; Level 3: quantify the severity of the damage; and Level 4: predict the remaining service life of the structure. A modification to these four levels, as described by Farrar & Worden (2007) considers the determination of the "type of damage" as an intermediate level between Levels 3 and 4. This incremental identification definition is suitable for the proposed approach, as it identifies the difficulty of detecting local failures.

Most of the existing literature addresses direct measurement approaches, in which sensors are placed on the structural elements from which one wishes to collect information to be used for the damage identification. The next section discusses research that has been done on indirect approaches to SHM, where data about the structure is collected from other sources, such as vehicles moving over the structure.

2.2 *Indirect Approaches for SHM*

Yang et al. presented an indirect approach in 2004, with the sole objective of extracting bridge frequencies from the dynamic response of a moving vehicle. They considered the bridge structure as a simply supported beam and the vehicle as a sprung mass. They derived an approximate analytical closed-form solution based only on the beam's first mode, and decoupled the bridge and the vehicle by neglecting the terms that contain the ratio of the oscillator mass to the beam total mass. This solution allows for the identification of a few significant dimensionless parameters that dominate the vehicle response, such as: $S = \pi v/L\omega_b$, a normalized vehicle velocity, where v = vehicle velocity, L = length of beam, and ω_b = bridge's natural fundamental frequency; and $\mu = \omega_b/\omega_v$, where ω_v is the vehicle (oscillator) vertical natural frequency. By performing a general finite element study, the concept was shown to be extendable to more complex structures. Later, Lin & Yang (2005) presented the experimental verification of the approach by using a four-wheel commercial light truck, towing a small two-wheel cart. They used accelerometers and velocity meters near the center of gravity of the cart to sense its vertical motion. The experiment also considered the use of a heavy truck that played the role of ongoing traffic. The authors concluded that it is feasible to scan the natural frequencies using the cart-based approach as the numerical study anticipated.

Another paper by Yang et al. (2005) explores the potential applications of an indirect approach to SHM. In this paper, Yang and his colleagues focused on the participation of the different modes of the bridges vibration and the complexity of dealing with multiple oscillators traveling at different speeds. The position of each vehicle within the bridge is crucial for determining the contribution

of the different excitation sources (passing vehicles) to the dynamic response of the bridge when considering multiple oscillators. Yang et al. (2005) also concluded that the first mode of the bridge was dominant in the dynamic response.

Yang's promising idea was not used for damage detection, but only for extracting the natural frequencies of the structure. Moreover, work done by Farrar indicates that natural frequencies by themselves are not good damage predictors (Farrar & Jauregui1998). The study by Farrar consisted of experiments on the I-40 Bridge used to compare five different damage assessment methods against the same set of data in order to contrast their detection capability. Different levels of damage were inflicted to a girder to test the sensitivity of the five methods considered. The studies found that resonant frequencies and modal damping are insensitive to low levels of damage, but experimentally determined mode shapes are more sensitive indicators. They also found that changes caused by environmental conditions can be as significant as the ones caused by damage.

2.3 *Damage Identification Approaches Using Moving Oscillators as Excitation Sources*

We now briefly describe some of the research efforts regarding moving loads for damage detection and experimental validation. Law & Zhu (2004, 2005) explored the changes in different damage indicators and the possibility of capturing those changes when considering the excitation of a moving oscillator on a beam. The flexural stiffness has been used as a damage index measure that has a good correlation with a vehicle's response (Law & Zhu 2005). Other authors have presented a damage detection approach based on both the vehicle's and the bridge's response in the time domain (Majumber & Manohar 2003). Yet others report the dynamic response of damaged beams subjected to moving masses (Mahmoud & Abou Zaid 2002, Bilello & Bergman 2004), but these studies do not take into account the suspension system of a vehicle. This simplification can be well justified as dynamic response of the vehicle is far less important in terms of the overall load of the vehicle when considering a static and dynamic load separately. Experiments with moving masses over a sliding rail have been performed to validate mathematical models of damaged beams (Bilello & Bergman 2004).

2.4 *Practical Advantages of Indirect Measurements from Passing Vehicles*

In this section, we point out some of the issues that cause significant practical challenges for direct monitoring, which are absent or mitigated when using an indirect approach. Using an indirect approach to SHM will have a number of potential advantages.

The first issue is related to the powering of the sensors. Since direct monitoring requires that sensors be deployed on the bridge being monitored, there is a need to provide power for the sensors and their associated electronics and data transmission and storage devices. A sustainable approach to providing this will likely consider energy harvesting in various forms, such as optimized solar energy (Alippi & Galperti 2008) or vibration based power collection systems (Beedy et al. 2006). In the case of indirect monitoring, there is readily available energy from the vehicle's electric system that completely eliminates the concern for how to provide power to the sensors while the vehicle is in operation.

The useful life of structures is much greater than the current reliable lifespan of most sensors. An indirect monitoring approach mitigates this issue because it will use data collected from many passing vehicles, which will have a variety of ages and thus a variety of ages of their sensor systems. As vehicles are replaced, the sensors in them will be replaced as well. In addition, the sensors will be protected from environmental conditions and the threat of vandalism, and will be able to be evaluated on a regular basis during routine vehicle maintenance intervals, whereas direct measurement devices require costly onsite sensor maintenance and are subject to harsh environmental conditions and vandalism.

The indirect monitoring approach will not cause traffic interruption, nor require the use of artificial loading devices, such as shakers or controlled load trucks. We consider the many moving vehicles on the bridge as both the excitation and sensing source for the sensing system. The basic idea is that the vehicles collect information about the dynamic vehicle-bridge interaction as they drive on the bridge.

2.5 *Multiresolution Classification Approach*

The task of classification is a standard signal-processing task that involves assigning one of the possible classes to a given input signal. This is typically done by computing certain numerical descriptors, called features, on the given input, in the hope that these descriptors will be sufficient to discriminate among classes. For example, some of the commonly used features (and those we use in this work) are the Haralick texture features (Haralick et al. 1973). Thus, a generic classification system

has a feature extraction block followed by a classifier block (see Figure 1).

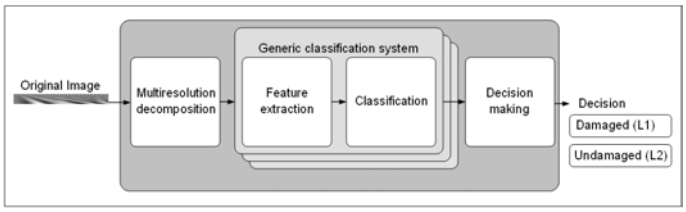

Figure 1. Multiresolution classification system for "Damaged" and "Undamaged" Labels.

Kovačević's lab has developed a new multiresolution classifier (Chebira et al. 2007c), that, instead of working on the original signal, passes it first through a multiresolution decomposition block generating a number of smaller-size signals, called subbands, in different subspaces. These subbands then each undergo separate classification, generating their own local classification decisions. In other words, each of the subbands is classified, each representing a possible classification of the actual signal. To reconcile these different possible classifications, the decision making block arbitrates and decides on the final label. This arbitration can be in closed form (that is, a solution to a least squares problem is found) or open form (where a reward-punishment system is iteratively applied onto subband local decisions). There can be as many labels as desired to take into account the different levels of damage identification. Considering the idea of *Existence*, Level 1 of damage identification, classifying signals as being "Damaged" and "Undamaged" (Figure 1) is the first experiment presented later in this paper.

In the research presented, synthetic data was created that would enable testing of the classification capability of the multiresolution algorithm for our application. The vehicle response vertical acceleration data was generated using a simplified numerical oscillator-beam interaction model that we describe below.

3 OSCILLATOR-BEAM INTERACTION MODEL

To explore the feasibility of this approach, an oscillator-beam interaction (OBI) model was implemented. It considers the coupling of the oscillator and the beam at a regular interval, ΔT. The algorithm iterates until the deflection of the beam at the point of interaction (z_b) and the base degree of freedom of the oscillator (z_v) converge. Only a few iterations are needed. Figure 2 shows a scheme of the implemented model.

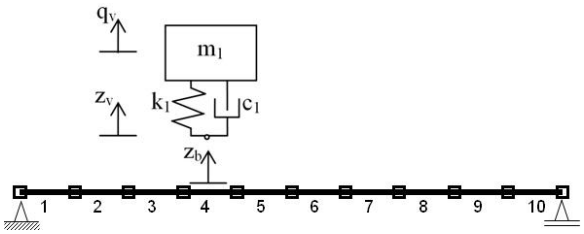

Figure 2. Oscillator beam model.

The model was validated against the results reported by Yang et al. (2004). For the experiments we chose oscillators of three different vertical fundamental frequencies (1.1, 1.7, and 3.4 Hz) to represent a small vehicle, large family vehicle and a truck, respectively. In addition, 30% critical damping was assumed for the different oscillators for modeling the shock absorbers and other energy dissipation in the vehicle. The beam is an idealization of a 40-m bridge, I=0.219 m^4 (Kim & Kawatani 2008) and it is discretized into 10 finite elements.

With some small changes to this numerical model, as described in the following list, we were able to produce synthetic response data for different damage conditions, such as:

a. *distributed section loss*, such as that caused by corrosion, and modeled by a percent reduction of the moment of inertia of a beam finite element;
b. *a section crack*, such as that caused by fatigue in steel elements, and modeled as a rotational spring, where the spring stiffness depends on the crack depth being modeled; and
c. *frozen bearings of bridge supports*, modeled using rotational spring elements linked to the rotational degrees of freedom at the ends of the beam.

The location and severity of these three damage conditions can be easily altered within the modeled beam. This allows a large number of different damage cases and associated response signals to be generated. In the research reported here, we first tested the ability of the approach to detect the absence or presence of section loss in different elements, and then considered an example in which different levels of damage were present.

4 MULTIRESOLUTION CLASSIFICATION

We now give a brief overview of the multiresolution techniques and wavelet-based approach we are using. Multiresolution techniques have been extensively studied and used in signal and image processing over the past two decades (Daubechies 1992, Vetterli & Kovačević 1995, Mallat 1999). We

call multiresolution techniques those signal processing tools that analyze and process signals across different frequency resolutions and scales. They have arisen in response to the inability of some standard techniques, such as Fourier analysis, to deal with nonstationary signals. For example, abrupt transitions in time cannot be captured using Fourier methods. An easy analogy is that with map programs on the Web, such as Google Maps. If we are looking into how to get to New York from Boston, the initial route will be at the scale/resolution of 50 km/1 in. Once close to New York, we will want more detailed directions, say to the Museum of Natural History, and would thus move to a scale/resolution of 2000 ft/1 in, which is the street level. In other words, we first investigated the global behavior of our signal, followed by its local behavior at a certain scale. This approach can be used in any situation where the signal is nonstationary. For example, we may assume that the data collected from a vehicle on a bridge will differ depending on time of day, day of the year, season, and many other factors. This is one of the main advantages of using this approach, as it enables the classification of new data based on a baseline provided by previous records. It takes into account modifications of the response by various factors, such as seasonal changes or daily temperature variations, as long as sufficient data is available.

The multiresolution techniques achieve their goal by decomposing a signal into zooming spaces (e.g., coarse subspaces and detailed subspaces) and are implemented by a signal-processing device called a filter bank. This filter bank then implements a specific multiresolution transform. Some well-known transforms that fit within this framework are the discrete Fourier transform (DFT) and the discrete cosine transform (DCT). Others, originating in the multiresolution literature, are the discrete wavelet transform (DWT) and a family of wavelet packet (WP) transforms. Which one of these to use depends on the specific application at hand. For more details, see, for example, Vetterli & Kovačević (1995).

One possible characterization of multiresolution transforms is in terms of whether they represent the signal in a nonredundant or a redundant fashion. Redundancy often leads to increased accuracy, as has been found in a host of bioimaging problems (see Chebira & Kovačević 2008a and references therein). One possible example of the power of multiresolution techniques in pattern classification is that developed for the classification of Drosophila embryo development (Kellogg et al. 2007). Using a highly accurate multiresolution classification algorithm developed by Kovačević and her group, the process is now automated and reproducible, with accuracy greater than 98% (Kellogg et al. 2007).

The use of wavelets in structural damage identification is relatively new. Melhem & Kim (2003) used continuous wavelet transform and Fourier analysis to detect damage on two full-scale concrete structures (a prestressed beam and pavement on grade) subjected to fatigue loads. Acceleration and deflection measurements were taken directly from the beam. Differences between initial and final damage states were significant and the wavelet analysis allowed for the identification of damage progression on both of the studied structures. Another study by Sun & Chang (2004) used a statistical wavelet-based method for structural health monitoring, which considered progressive damage on a steel cantilever beam. Sun concluded that indicators from the Wavelet Packet Signature (WPS) are excellent indicators for monitoring structural health condition. They are sensitive to structural damage and insensitive to measurement noise.

A recent paper by Law et al. (2008) makes use of wavelet transforms for identifying a moving load over a beam and the prestress condition. In Law's work, the measuring points are located at the bottom of the beam. The forces of two moving axles and the prestress levels are identified successfully over time with high accuracy.

4.1 *Multiresolution Algorithm for Classifying Signals from the OBI Model*

The multiresolution classifier we use here was originally developed for images. For the purpose of testing whether it would make sense to use it for classifying the signals taken from a vehicle moving over a bridge, we had to produce images from the collected vehicle data. This requires the images to be at the same scale, considering the maximum and minimum values of the whole set of data as the scale limits. As an example, Figure 3 shows an image before scaling. It corresponds to a 2% inertia reduction on an element adjacent to the midspan of the beam. The ordinate represents the different velocities of the oscillator, and the abscissa refers to the relative position of the oscillator with respect to the beam. The colors represent the acceleration value of the oscillator. The scale and frequency content of the image represent the dependency of the response on the position and velocity of the oscillator.

At Level 1, the multiresolution decomposition takes an image and produces a number of smaller

images from which the original one can be reconstructed, if needed (see Figure 4). At Level 2, the same, or different, multiresolution decomposition is applied to a subset, or all, of the images from Level 1. The process can be repeated many times, at each level producing subbands at a different resolutions/scales. For example, Figure 4 shows a preprocessed scaled input image, and four subbands at Level 1. In our experiments we used a 2-level full decomposition (meaning at each level, each subband is split into four subbands at the next level). The left-most subband is typically the one that carries the global characteristics of the signal (so it very much looks like the higher-level image, but blurred), while the other three carry the necessary details to reconstruct the original (these are the local changes, or, edges).

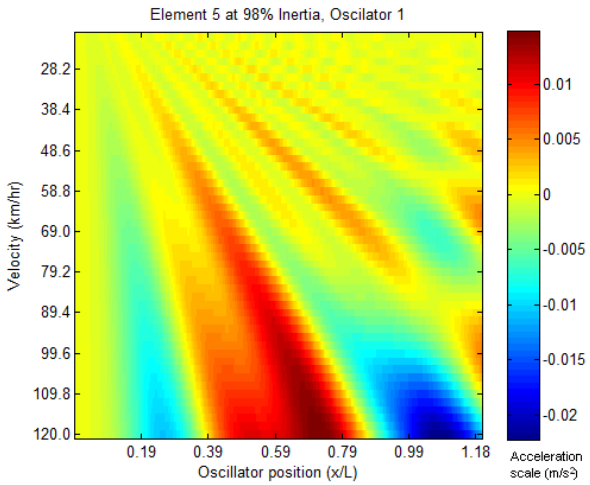

Figure 3. Raw image representation of acceleration response of traveling oscillator.

When small amounts of data are available, as is the case here, a technique called *leave-one-out cross validation (LOOCV)* is used. LOOCV attempts to estimate the generalization error of the classifier, which is effectively the capability of the classifier to correctly classify unseen data (i.e., data that has not been trained on). Let N be the number of data samples, $[x_1\ x_2\ \ldots\ x_N]$, for a particular class of images to be classified. For a particular data sample x_i, the classifier is trained using samples $[x_1\ \ldots\ x_{i-1}\ x_{i+1}\ \ldots\ x_N]$ and is tested over the sample x_i. This is repeated for each available data sample that results in N separate classifiers being trained and tested. The overall accuracy of the classifier is the average accuracy over these N results.

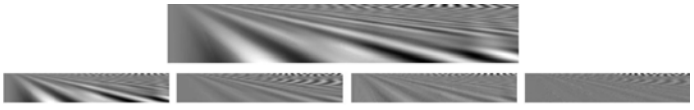

Figure 4. Multiresolution decomposition (Level 1). Preprocessed original image (top); four subbands (bottom).

4.2 *Experimental Setup and Results*

We conducted two experiments with the acceleration response of the oscillators obtained from the numerical model described in Section 3.

Experiment 1: Existence of damage classification. In this experiment, we attempted to classify the bridge into one of two categories: "Undamaged" (5 cases, Table 1); or "Damaged" (Levels 1-4 lumped together, 25 cases, Table 1). Thus, with N=30, there were 30 two-class classification experiments for each of the three oscillators and each of the 10 elements.

Table 1. Damage cases considered.

Label	Damage range [%]	# of cases
Undamaged	0-4	5
Level 1	5-9	5
Level 2	10-14	5
Level 3	15-29	8
Level 4	30-55	7

Table 2 shows the overall classification accuracy for the "Undamaged" and "Damaged" cases over each element. These results indicate that this approach is able to achieve very high accuracy in classifying a damage condition occurring over an element near the midspan for the three oscillators.

Table 2. Two-class damage classification accuracy (in %).

Element	Oscillators			**Accuracy**
	1.1Hz (small veh.)	1.7 Hz (family veh.)	3.2Hz (truck)	
1	96.7	83.3	76.7	**85.6**
2	76.7	86.7	73.3	**78.9**
3	96.7	86.7	86.7	**90.0**
4	93.3	93.3	96.7	**94.4**
5	93.3	93.3	96.7	**94.4**
6	100.0	90.0	100.0	**96.7**
7	93.3	93.3	80.0	**88.9**
8	83.3	90.0	90.0	**87.8**
9	83.3	86.7	73.3	**81.1**
10	90.0	83.3	83.3	**85.5**
Accuracy	**90.7**	**88.7**	**85.7**	**88.4**

Note that the first oscillator achieves the highest average accuracy, followed by the second and then the third. One might infer that the larger the load, the lower the accuracy. These results are preliminary because they use a highly idealized beam to represent a bridge. We must run more experiments involving structures with different natural frequencies and different vehicles to validate this conjecture. Also, in this experiment, it was possible to determine the damage in some elements very accurately (for example, Element 6), while not for others (for example, Element 2 near the end of the beam). Overall, the multiresolution classification approach achieved an average accuracy of 88.4% in

determining the *existence* of damage. We are thus encouraged that a more comprehensive investigation of this approach will improve these accuracies.

Experiment 2: Severity of damage classification. This experiment considers an undamaged class and four damage levels for a total of five classes/labels as described in Table 1. With these five labels, there were again N=30 classification experiments, but there were five classes to be distinguished for each oscillator and element. Table 3 shows the overall classification accuracy for the five classes ("Undamaged" and "Damaged" levels 1-4) for each oscillator and beam element. The results shown in Table 3 can be read as: the percentage of accuracy to classify an image representation (see Figure 3) of the vehicle response as a specific damage level (Table 1), on a particular element and considering a particular vehicle. Note that what is reported in the table is the aggregate accuracy of the methodology for distinguishing each of the five levels of damage.

First, observe that, as expected, the accuracies in Experiment 2 are lower compared to those from Experiment 1. This is because it is harder, using the same number of actual collected signals, to distinguish between five different levels of damage, as opposed to just the undamaged/damaged situations. In contrast to Experiment 1, the middle oscillator gives the most accurate results in Experiment 2. Overall, the multiresolution approach achieved an average accuracy of 71.2% in determining the severity of the damage.

Table 3 Five-class severity of damage classification accuracy (in %).

Element	Oscillators			**Accuracy**
	1.1Hz (small veh.)	1.7 Hz (family veh.)	3.2Hz (truck)	
1	66.7	80.0	83.3	**76.7**
2	66.7	76.7	66.7	**70.0**
3	80.0	73.3	40.0	**64.4**
4	60.0	66.7	70.0	**65.6**
5	76.7	80.0	63.3	**73.3**
6	76.7	80.0	73.3	**76.7**
7	76.7	73.3	70.0	**73.3**
8	73.3	66.7	50.0	**63.3**
9	66.7	83.3	66.7	**72.2**
10	73.3	80.0	76.7	**76.7**
Accuracy	**71.7**	**76.0**	**66.0**	**71.2**

5 CONCLUSIONS

We have presented an alternative approach for indirect monitoring of the structural health of bridges through data collected from vehicles passing over a bridge. To test the validity of this approach, we created a numerical model of the interaction between a simple oscillator and a simple beam and subjected the beam to different levels of section loss at different locations. We then subjected the simulated responses to these damaged states to a multiresolution classification system in order to determine how accurately the damage level could be classified. The results of these two experiments, while limited and very preliminary, seem promising. We are encouraged to further pursue the refinement and evaluation of this approach.

6 AKNOWLEDGEMENTS

Financial support for this research from the following sources is gratefully acknowledged: 1) The Fulbright Foundation-MECESUP2 scholarship, Gobierno de Chile; 2) Professor Pradeep Khosla, Dean of the College of Engineering, Carnegie Mellon; 3) Dr. Richard McCullough, Vice President of Research, Carnegie Mellon University; 4) the Hillman Foundation.

7 REFERENCES

Alippi, C., and Galperti C. 2008. An Adaptive System for Optimal Solar Energy Harvesting in Wireless Sensor Network Nodes. *Circuits and Systems* I: Regular Papers, IEEE Transactions on 55, no. 6: 1742-1750.

Beeby, S. P., Tudor, M. J., & White N. M. 2006. Energy harvesting vibration sources for microsystems applications. *Measurement Science and Technology* 17, no. 12: R175-R195.

Bilello, C., & Bergman L. A. 2004. Vibration of damaged beams under a moving mass: theory and experimental validation. *Journal of Sound Vibration* 274 (Jul 1): 567-582.

Brownjohn, J.M.W. 2007. Structural health monitoring of civil infrastructure. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365, no. 1851 (February 15): 589-622.

Carden, E. P., & Fanning P. 2004. Vibration Based Condition Monitoring: A Review. *Structural Health Monitoring* 3, no. 4 (December 1): 355-377.

Chebira, A. & Kovačević, J. 2007. Adaptive multiresolution frame classification of biological and biometric images. *Proc. of SPIE Conf. on Wavelet Applications in Signal and Image Proc*., San Diego, USA, Aug.

Chebira, A., Coelho, L. P., Sandryhaila, A., Lin, S., Jenkinson, G. W., MacSleyne, J., Hoffman, C., Cuadra, P., Jackson, C., Püschel, M., & Kovačević, J. 2007a An adaptive multiresolution approach to fingerprint recognition," *Proc.*

IEEE Conf. on Image Proc., San Antonio, TX, Sep. pp. I:457–460.

Chebira, A. & Kovačević, J. 2007b. Lapped tight frame transforms. *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc.*, Honolulu, HI, Apr. pp. III:857-860.

Chebira, A., & Kovačević, J. 2008a. Frames in bioimaging. Proc. CISS, Princeton, NJ, Mar.

Chebira, A., Ozolek, J. A., Castro, C. A., Jenkinson, W. G., Gore, M., Bhagavatula, R., Khaimovich, I., Ormon, S. E., Navara, C. S., Sukhwani, M., Orwig, K. E., Ben-Yehudah, A., Schatten, G., Rohde, G. K., & Kovačević, J. 2008b. Multiresolution identification of germ layer components in teratomas derived from human and nonhuman primate embryonic stem cells. *Proc. IEEE Intl. Symp. Biomed. Imaging*, Paris, France, May. pp. 979-982.

Chebira, A., Barbotin, Y., Jackson, C., Merryman, T., Srinivasa, G., Murphy, R. F., & Kovačević, J. 2007c. A multiresolution approach to automated classification of protein subcellular location images. *BMC Bioinformatics*, vol. 8, no. 210.

Daubechies, I. 1992. Ten lectures on wavelets. Society for Industrial and Applied Mathematics.

Doebling, S. W, Farrar C. R, & Prime M. B.. 1998. A summary review of vibration-based damage identification methods. Identification Methods," *The Shock and Vibration Digest* 30: 91-105.

Farrar, C. R., & Jauregui, D. A. 1998. Comparative study of damage identification algorithms applied to a bridge: I. Experiment.*Smart Materials and Structures* 7, no.5:704-719.

Farrar, C. R., & Worden K. 2007. An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365, no. 1851 (February 15): 303-315.

Federal Highway Administartion 2006. Status of the Nation's Highways, Bridges and Transit: Conditions and Preformance. http://www.fhwa.dot.gov/policy/2006cpr

Friswell, M. I. 2007. Damage identification using inverse methods. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365, no. 1851 (February 15): 393-410.

Fugate, M. L., Sohn H., & Farrar C. R. 2000. Unsupervised learning methods for vibration-based damage detection. *Presented at IMAC*, San Antonio, Texas. February 7-10.

Haralick, R.M, K. Shanmugan, and I. Dinstein. 1973. Textural features for image classification. EEE Trans. on Systems, Man, and Cybernetics 3, no. 6: 610-621.

Kellogg, R. A., Chebira, A., Goyal, A., Cuadra, P. A., Zappe, S. F., Minden, J. S., & Kovačević, J. 2007. Towards an image analysis toolbox for high-throughput Drosophila embryo RNAi screens. *Proc. IEEE Intl. Symp. Biomed. Imaging*, Arlington, VA, Apr. pp. 288-291.

Kim, C. W, & Kawatani M.. 2008. Pseudo-static approach for damage identification of bridges based on coupling vibration with a moving vehicle. *Structure and infrastructure engineering* 4, no. 5: 371-379.

Law, S. S, & Zhu. X. Q. 2005. Nonlinear characteristics of damaged concrete structures under vehicular load. *Journal of Structural Engineering-ASCE* 131, no. 8 (August): 1277-1285.

Law, S. S., Wu S. Q., & Shi Z. Y. 2008. Moving Load and Prestress Identification Using Wavelet-Based Method. *Journal of Applied Mechanics* 75, no. 2 (Mar. 0): 021014-7.

Law, S. S., & Zhu X. Q. 2004. Dynamic behavior of damaged concrete bridge structures under moving vehicular loads. *Engineering Structures* 26, no. 9 (July): 1279-1293.

Lin, C.W., & Yang Y.B. 2005. Use of a passing vehicle to scan the fundamental bridge frequencies: An experimental verification. *Engineering Structures* 27, no. 13 (November): 1865-1878.

Lynch, J. P. 2007. An overview of wireless structural health monitoring for civil structures. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365, no. 1851 (February 15): 345-372.

Mahmoud, M. A., & Abou Zaid M. A.. 2002. Dynamic response of a beam with a crack subject to a moving mass. *Journal of Sound and Vibration* 256, no. 4 (September 26): 591-603.

Majumder L., & Manohar C. S. 2003. A time-domain approach for damage detection in beam structures using vibration data with a moving oscillator as an excitation source. *Journal of Sound and Vibration* 268, no. 4 (December 4): 699-716.

Mallat, S. 1999. A Wavelet Tour of Signal Processing. 2nd ed. Academic Press.

Melhem, H., & Kim H. 2003. Damage Detection in Concrete by Fourier and Wavelet Analyses. *Journal of Engineering Mechanics* 129, no. 5 (May 0): 571-577.

Rytter, A, 1993. Vibration based inspection of civil engineering structures. *Department of Building Technology and Structural engineering, Aalborg University*, Denmmark.

Sun, Z., & Chang, C. C. 2004. Statistical Wavelet-Based Method for Structural Health Monitoring. *Journal of Structural Engineering* 130, no. 7 (July 0): 1055-1062.

Van der Auweraer, H., & Peeters B.. 2003. International Research Projects on Structural Health Monitoring: An Overview. *Structural Health Monitoring* 2, no. 4 (December 1): 341-358.

Vetterli, M. & Kovačević, J. 1995. Wavelets and Subband Coding, Prentice Hall, Signal Processing Series, Englewood Cliffs, NJ.

Worden, K., &. Dulieu-Barton J. M. 2004. An Overview of Intelligent Fault Detection in Systems and Structures. *Structural Health Monitoring* 3, no. 1 (March 1): 85-98.

Worden, K., & Manson G. 2007. The application of machine learning to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365, no. 1851 (Feb. 15): 515-537.

Yang, Y. B., Lin C. W., & Yau J. D. 2004. Extracting bridge frequencies from the dynamic response of a passing vehicle. *Journal of Sound and Vibration* 272, no.3-5(May 6):471-493.

Yang, Y.B., & Lin C.W. 2005. Vehicle-bridge interaction dynamics and potential applications. *Journal of Sound and Vibration* 284, no. 1-2 (June 7): 205-226