The 3rd International Conference on Structural Health Monitoring of Intelligent Infrastructure
Vancouver, British Columbia, Canada
November 13-16, 2007

AUTOMATIC VEHICLE CLASSIFICATION BASED ON MODELLED BRIDGE STRAIN PROFILES

Nishant J. Dhruve
University of Manitoba, Canada

Dean K. McNeill
University of Manitoba, Canada

Abstract

This article reports results obtained in applying neural computing techniques to the problem of vehicle type classification, based upon strain information obtained from a vehicle bridge. In the absence of actual strain measurements from a structure in service, simulated strain responses to the passing of cars, vans, buses and trucks were generated and used for training and evaluation purposes. It was found that a neural classifier is able to accurately identify single isolated vehicles, with success rates up to 100%, based on the simulated strain profile from four sensors.

It was further found that proper identification of one vehicle when two are present within a single observation period was possible, even when the strain responses are overlapping. However, in these cases, the recognition accuracy obtained does not exceed 90% for 100 test vehicles. The vehicle type selected by this system corresponds to the vehicle with the highest magnitude strain signature. Modifications to the system were explored in efforts to improve recognition while removing the emphasis on magnitude alone. In doing so, a classifier was produced, which selected the most consistently identified input pattern over a series of four sensors. Recognition accuracy for this system was found to be 84% for 1000 simulated vehicles.

INTRODUCTION

The owners of large civil infrastructure, such as vehicle and rail bridges, see structural health monitoring (SHM) as an important immerging tool, to better maintain and manage their infrastructure. In the ideal case, an SHM system is permanently deployed on a structure, allowing for continuous oversight of the structure's operation during its daily use. These systems consist of a network of wired or wireless sensors that are placed at key locations throughout the structure in order to monitor important aspects of its operation. A variety of sensor types are typically used, measuring properties such as strain, temperature, acceleration, displacement and vibration of key components.

If the monitoring system is included as an element of the original structural design, and thus is in place from the time the structure is first commissioned, it makes it possible to track the behaviour of components as they undergo regular use. Furthermore, changes in the properties of these components that occur as a result of normal aging of the structure, or in response to external events, can be more readily identified. Thus, SHM provides infrastructure owners with the information needed to continuously assess the state of their infrastructure and, from that knowledge, to respond in a timely manner to concerns or problems as they arise.

A direct consequence of the continuous aspect of structural monitoring is the need to manage the resulting vast collection of measurement data obtained. In previous studies [1,2], we have explored methods of differentiating between measurement events that are typical of a specific structure's normal response to environmental excitation from those which are atypical. This was accomplished by using neural computing techniques to construct a model of the unique response of an individual structure, based solely upon raw sensor measurements obtained over time. Once such a suitable model was obtained, it was then used as a comparative baseline from which to assess further measurements. The degree to which those measurements agreed with previous experience provided an assessment of the novelty of the new measurement. A consequence of this approach is the ability to employ the novelty assessment as a means of filtering collected data to preferentially store noteworthy measurements.

This article extends that previous work by introducing neural computing techniques to classify the nature of an external excitation, which gives rise to a measured sensor profile. In doing so, it is expected that noteworthy or novel events with known causes can be identified and counted. As a consequence, the remainder of events that do not adhere closely to a known class of excitation can be separated from those of known origin, and thus examined in greater detail. The specific nature of excitation we are concerned with in this article relates to strain profiles produced by vehicular traffic on a bridge.

VEHICLE MODELS

Four different vehicle classes are considered in the current study. These are cars, vans, buses and large trucks. In the ideal situation, actual measurements of excitations from these four vehicle types recorded from an in-service bridge structure would serve as the basis for training and evaluating the neural vehicle classifiers. The bridge structure of interest here is the Red River-North Perimeter Bridge in Winnipeg, Manitoba, Canada. However, this structure has only very recently been fitted with a sensor system consisting of 48 strain gauges positioned on one bridge span. Actual measurement data from this bridge was unavailable at the time this study was undertaken. As a result, in order to facilitate investigations into the general suitability of neural techniques for vehicle classification, we instead produced simulated strain profiles intended to mimic, in overall form, those expected to occur in measurements from such a structure. Based on previously observed measurements, profiles for the four vehicle classes were produced using a combination of peaked functions of a Gaussian nature.

Individual Vehicle Data Models

Simulated strain information for car and van classes were produced using the same generating function, consisting of a single peaked response. In the case of these smaller vehicles, the relative spacing of the vehicle axels, along with the low vehicle weight, do not make it possible to distinguish individual vehicle axels in the strain measurements from bridge girders or the bridge deck. Instead, a single strain peak resulting from the entire vehicle passing the area of a sensor should occur. The generating function used for the car and van data is given in Equation

$$StrainData(i) = w_k * e^{\frac{-(t_i - m)^2}{2\sigma^2}} \qquad (1)$$

In order to mimic the expected variations in the weights of vehicles and their speed, the magnitude of the peak, controlled by the term w_k is itself selected from a Gaussian distribution with a fixed mean and a variance. In the case of cars, an average weight of 1450 kg was used, along with a variance of 200 kg. For vans, a mean weight of 2000 kg was used, along with the same 200 kg variance. Vehicle speed and length affects the spread of the peaks as controlled by the variance term σ^2 in Equation 1. As with weight, variations in vehicle length is also modelled using a Gaussian distribution. Here a mean length of 4.5m was used, together with a variance in that length of 0.5m. Finally, the term m in Equation 1 permits the location of the peak along the timeline of the simulated measurement data to be adjusted. Examples of the resulting simulated strain profiles are provided in Figure 1 for a car and van.

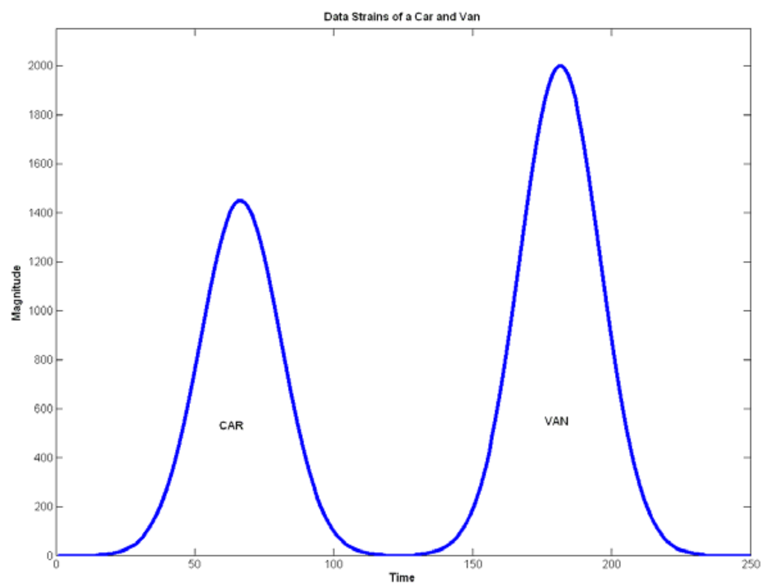

Figure 1: Simulated strain profiles for a car and van.

Unlike cars and vans, which consist of a single peaked response, the greater overall weight of buses and trucks, along with their much larger axel spacing, allows for the distinction of individual axels within a strain profile. Therefore, the simulated strain response for these two classes of heavier vehicles are made up of two prominent peaks. In the case of the bus, the peaks are of similar amplitude, with the height of the rear peak being slightly greater in order to take into account the weight of the engine, which is normally located in the rear of a bus. As well, the space between the two peaks is made slightly shallow by introducing a third Gaussian between the initial two. Equation 2 gives the full formulation of the simulated bus response.

$$StrainData(i) = w_{k1} * e^{-\frac{(t_i - m_1)^2}{2\sigma_2^2}} + w_{k2} * e^{-\frac{(t_i - m_2)^2}{2\sigma_2^2}} + w_{k3} * e^{-\frac{(t_i)^2}{2\sigma_3^2}} \qquad (2)$$

In the case of large trucks, only a two peak response is considered. This is to mimic the truck cab and trailer, each passing separately over the location of a sensor. Equation 3 gives the full functional description of the simulated truck profile.

$$StrainData(i) = w_{t1} * e^{-\frac{(t_i - m_1)^2}{2\sigma_2^2}} + w_{t2} * e^{-\frac{(t_i - m_2)^2}{2\sigma_2^2}} \qquad (3)$$

As was the case with cars and vans, the weight and size of buses and trucks were varied based upon a Gaussian distribution. Specifically, the average weight of a bus is taken as 18000 kg with a variance of 500 kg. For trucks the average is taken to be 20000 kg with a variance of 1000 kg. The length of buses were taken to be fixed at 12m, while that of trucks was nominally 12.5m with a variance of 0.5m. Figure 2 illustrates the bimodal profile of the bus and truck response.

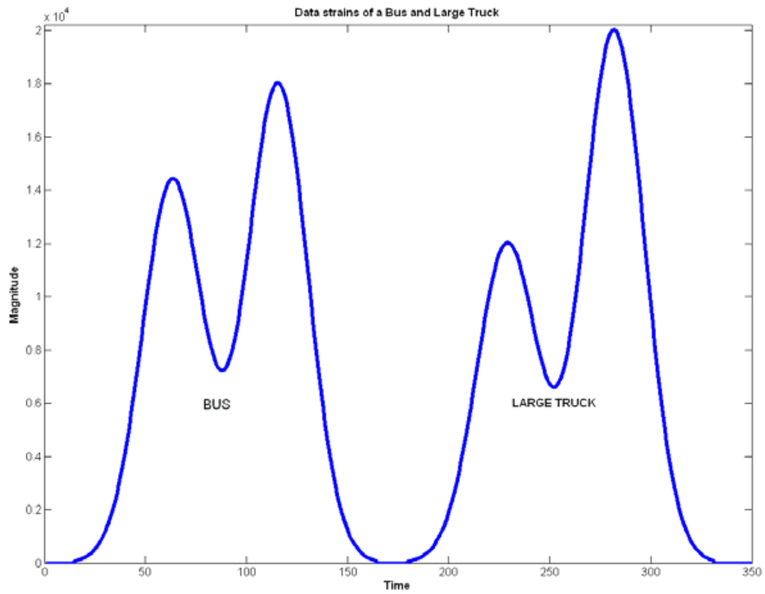

Figure 2: Simulated strain profile for a bus and truck.

Vehicle Datasets

The individual data models were used to generate two types of simulated data for training, testing and validation of the vehicle classifiers being investigated. The first is a *sparse dataset* consisting of single, isolated vehicle events within any four second period in time. While the measurement data itself would be continuous, to facilitate process of the information by the vehicle classifier, it is broken down into four second measurement windows. This window size was selected to insure that the strain profile for a vehicle moving at slow speeds would be fully accommodated within the observation window. The position of a vehicle within this window was selected randomly with a uniform distribution.

The second dataset type is an *overlapping dataset* consisting of at most two vehicles within the same observation window. In many cases, these simulated strain profiles will involve considerable overlap between the two vehicles. This type of response would be typical of two vehicles traveling close to each other in the same traffic lane, or two vehicles traveling together in adjacent lanes.

Each of the datasets is composed of measurements from a series of four simulated sensors positioned 4m apart longitudinally in the centre of a single bridge span. While this sensor count is much lower than the 48 present on the actual bridge, many of the later sensors are recording similar information from adjacent girders and steel strapping all supporting the same deck slab.

NEURAL VEHICLE CLASSIFIER

In this study we wish to explore the ability of neural computing techniques to perform vehicle type classification based on the available sensor data. The neural system we are using is composed primarily of two layers of processing elements, as shown in Figure 3. The lower layer is trained first in an unsupervised manner using the frequency sensitive competitive learning (FSCL) algorithm [3,4]. The algorithm adjusts the layer one interconnection weights in order to model the general statistical distribution of the data within the input space. In this particular application, the input is composed of the full strain information from the four sensors during a single four second observation window. For a simulated sensor sampling frequency of 100Hz this translates into a 1600 dimensional input.

This first layer becomes a feature extractor that is tuned to the specific properties of the input. An exponential output function is applied to the layer one outputs before they are passed on as inputs to layer two. Layer two of the network is trained in a supervised manner [5], using the layer one features along with an output target vector to direct the training. Its output consists of four neurons, each corresponding to one of the four vehicle type classifications of interest.

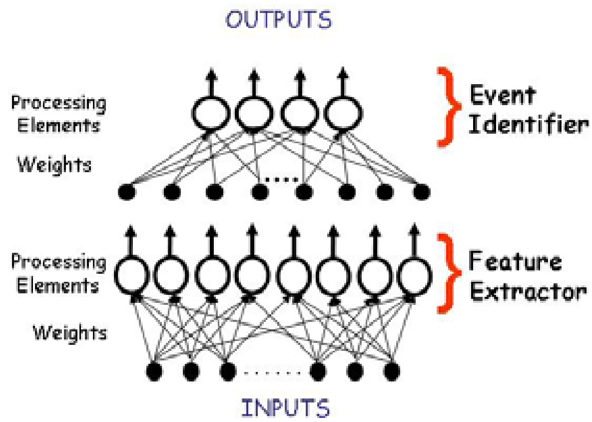

Figure 3: Neural network vehicle classifier module.

The system was trained using a dataset containing only sparse vehicle events. It was subsequently tested on its ability to classify vehicles from a second sparse dataset and an overlapping dataset. When evaluating unseen sparse inputs, the system was able to correctly classify 68% of 3000 input patterns. A significant proportion of the patterns that were labeled incorrectly (14% of the 3000) were misclassified as a direct result of the vehicle strain profile not being completely contained within the observation window. This is analogous to a vehicle entering or leaving the window and thus displaying a clipped strain response.

When tested on overlapping data, classification decisions from five consecutive observation windows were combined using a majority rule in order to formulate a classification consensus based on additional evidence. However, even with the aid of the majority rule, the system was only able to correctly classify 60% of 100 test vehicles. In this context, a window of data was considered to be correctly classified if either one of the two vehicles were properly identified. During these tests, the classifier showed an inability to distinguish between cars and vans, and buses and trucks.

An analysis of the systems operation during these tests showed that the significantly larger magnitude inputs of buses and trucks were dominating the feature extraction process. To address this problem, a fixed data preprocessing stage was introduced, which scaled the data in the current input window to a uniform peak amplitude. By removing the discrepancy in amplitude, this one property would no longer skew the assignment of processing units based on input magnitude. The scale factor used to correct the input magnitude was not discarded, but was provided to a modified second layer as a further input. Figure 4 shows the structure of this modified network.

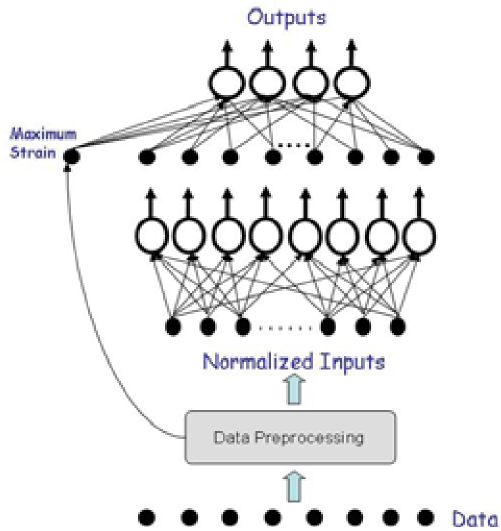

Figure 4: Neural network vehicle classifier with input scaling.

Tests using the modified network architecture with scaling showed a significant increase in recognition accuracy. When evaluated using the sparse dataset, classification accuracy rose to 90% on 1000 vehicles. When a majority rule was used with the sparse dataset, perfect classification of all test inputs were obtained. Tests of the overlapping dataset displayed a similar increase in performance. Successful classification of the dominant vehicle type was achieved in 90% of the tests conducted using 100 vehicles.

One of the shortcomings observed with the scaled classifier is its predisposition to select an input pattern classification that corresponds with the highest magnitude input present in the observation window. The presence of another vehicle with a lower magnitude and appearing for a longer duration is completely lost to a vehicle with a higher response. For example if the observation window contains strains from both a car, which is seen by all the sensors, and a bus, seen by only two of the four sensors, the system classification would be that of a bus. However, given that the car is visible in the response of all sensors, a more consistent classification would be that of a car.

In an effort to correct for this problem, inputs from each of the four sensors were separated and provided to their own individual sub-network. This is illustrated in Figure 5. A single sub-network consists of a feature extractor and an event identifier. Data from each sensor is first normalized and then presented to the feature extractor layer of each subsystem. Outputs from the feature extractor are then presented to the event identifier along with the maximum strain extracted during the preprocessing stage. The classifications from each of the four event identifiers are then forwarded to the final layer for processing. A final classification decision is then made using the maximum number of wins from the four subsystem classifications. Training of the subsystems is accomplished by first training only a single subsystem, as was done previously, and then replicating the same trained sub-networks for each the four sensors. The final output stage does not require training as it is simply a majority winner layer.

The classification accuracy of this system degraded to 84% for 1000 simulated vehicles from the sparse dataset. It is thought that this decrease in performance is due to poor training at the subsystem level. While individual vehicle performance declined, detection of two vehicles present in the same observation window improved significantly. The system was able to detect and classify vehicles with lower magnitudes but with a more consistent presence in the window. The classification decision was given to the vehicle type most represented by the sensors. For example, a car observed by three sensors and a large truck observed by only one will result in the car being selected as the winner.

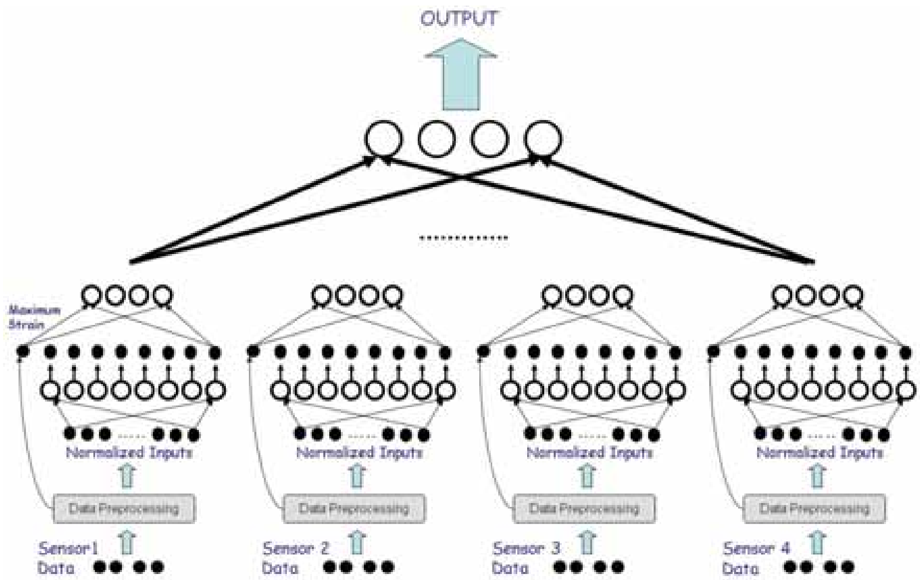

Figure 5: Neural network vehicle classifier with voting.

CONCLUSIONS

This work has explored the use of neural computing techniques for the classification of vehicle type based on modelled strain measurements from a vehicle bridge. Three network architectures were explored. It was found that recognition accuracy of up to 100% was possible for isolated vehicles within a fixed observation window. Accuracy in the presence of two vehicles was found to be 90% in the case of one network architecture favouring overall input magnitude, while an accuracy of 84% was obtained in a second network focusing on the overall consistency of a response.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support of Intelligent Sensing for Innovative Structures (ISIS Canada), and the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

1. Card, L. and McNeill, D. K., “Characterization of System Sensitivity in SHM Event Detection Systems,” Proceedings of the SPIE 10th Annual International Symposium on Nondestructive Evaluation for Health Monitoring and Diagnostics, San Diego, CA, vol. 5767, 2005, 143–154.

2. McNeill, D. K. and Card, L. “Adaptive Event Detection for SHM System Monitoring,” in Sensing Issues in Civil Structural Health Monitoring, Proceedings of the North American Euro-Pacific Workshop on Sensing in Civil Structural Health Monitoring, Honolulu, Hawaii, 2005, 311–319.

3. Ahalt, S. C., Krishnamurthy, A. K., Chen, P. and Milton, D. E., “Vector Quantization using Frequency-Sensitive Competitive-Learning Neural Networks”, IEEE International Conference on Systems Engineering, (1989) 131–134.

4. Ashalt, S. C., Krishnamurthy, A. K., Chen, P. and Milton, D. E., “Competitive Learning Algorithms for Vector Quantization,” Neural Networks, vol. 3 (1990) 277–290.

5. Bishop, C. M., Neural Networks for Pattern Recognition (Oxford University Press, 1995).

6. Hertz, J., Krogh, A., Palmer, R. G., Introduction to the Theory of Neural Computation (Addison-Wesley, 1991).