Journal of Sound and Vibration (1998) **211**(1), 133–143

LETTER TO THE EDITOR

NEURAL NETWORK APPROACH TO LOCATING ACOUSTIC EMISSION SOURCES IN NON-DESTRUCTIVE EVALUATION

J. C. Spall, J. L. Maryak and M. S. Asher
Applied Physics Laboratory, The Johns Hopkins University, Laurel, MD 20723-6099, *U.S.A.*

(*Received* 23 *December* 1996, *and in final form* 13 *October* 1997)

1. INTRODUCTION

The need for non-destructive evaluation (NDE) technologies for maintenance of complex systems has long been recognized. NDE typically involves the application of sensors, which may be apart from or temporarily or permanently attached to the structure, possibly in combination with active signal generation procedures. A large number of NDE techniques have been devised and applied, including load deformation, liquid penetrant and dynamic signature analyses, impact hammer and impact echo tests, and ultrasonic, radiographic, microwave, and magnetic signal analyses (see, e.g., Jones and Ellingwood [1] and Plecknik and Henriquez [2]). Teller [3] discusses the advantages and disadvantages of several of these techniques, along with providing an introduction to an emerging technology: acoustic emission (AE) testing, which will be the focus of the rest of this paper for reasons discussed in section 2 below.

AE testing analyzes signals received at a set of transducers installed at several points on the system, with the aim of detecting the presence, location, and intensity of acoustic signals generated by cracking within members of the system. The signals received at a transducer may be generated by an AE cracking event (or events) and by "nuisance" signals caused by environmental factors, such as traffic and weather on a bridge. The transducer outputs also contain measurement noise.

Because of non-linearities and many other complexities in the analysis of such signals (discussed further below), it seems appropriate to use a neural network (NN) approach to AE testing for complex systems. Neural networks represent a class of non-linear function approximators that have proven effective in a broad range of problems for which traditional modelling approaches are infeasible. Our focus will be on the most common type of NN, namely the feedforward multilayer NN (see Narendra and Parthasarathy [4]), but other NNs may be appropriate as well (see brief discussion in section 4). The NN seems considerably easier to apply than standard modelling techniques (in fact, with the inherent difficulty of non-linear signal inversion for structures, there does not appear to have been any application of classical modelling techniques to the NDE of a complex system using AE analysis). Also, the convenient availability of appropriate data, including laser-induced simulations of AE cracking events (discussed below), will facilitate NN training. The idea of using a NN to process AE data in general (i.e., not in a "structures" context) was proposed by Grabec and Sachse [5] and successfully tested by them on a small block of aluminum. A similar project was reported by Yuki and Homma [6]. Of course, there are many issues to be resolved in applying NN to a complex structure such as a bridge; some of these issues are discussed below.

0022–460X/98/110133 + 11 $25.00/0/sv971355

In the following sections of this summary, we expand upon the idea of using AE to monitor cracking in a system or structure, give further motivation for the NN approach, and describe some details of the NN implementation, training, and application. Also included are experimental results on a steel I-beam. We conclude with a summary of the many open issues remaining in traversing from the concept stage of this paper to a practical large-scale test or implementation.

2. BACKGROUND ON NDE THROUGH USE OF AE

This section provides some background on AE-based testing for NDE. This includes a summary of the advantages and disadvantages, and a brief discussion of some related experience with AE. Among the advantages AE offers for monitoring for structural flaws are the following.

Since the AE events will excite Rayleigh (surface) and Lamb (plate) waves† that can propagate for several meters, AE transducers do not have to be co-located with the material fault to be detected (so there is no need to know in advance where an AE event will occur). In structures analysis, an inspection team does not therefore have to crawl over the structure with test equipment looking for flaws.

AE does not require access to both sides of a structural member, as does C-scan ultrasonic testing, where ultrasound is injected on one side of the member and received by a transducer on the other side.

AE does not involve the emission of hazardous X-rays as does radiographic testing.

Analysis by AE is ideally suited to NDE from a remote analysis facility since no human intervention is required and the transducer signals can be transmitted to a remote site (possibly by means of fiberoptic cable). This remote site may be monitoring many systems simultaneously, which has the potential to significantly reduce the inspection costs incurred by the relevant municipal, state, or federal agency.

AE is only emitted from growing cracks—static cracks are silent (except possibly for the noise generated by crack rubbing)—which are indicators of system deterioration. This means that AE testing requires continuous monitoring of the system (unless it is "proof" tested at loads higher than those previously experienced). This is not prohibitive, however, owing to the remote sensing capability of the AE NDE method mentioned above.

The major problems with utilizing AE as a practical NDE method for continuous monitoring of in-service systems are the following.

The difficulty of relating the signals observed at the transducer to the AE events. It is extremely difficult to model the propagation of the elastic waves that carry the AE information through the system: (1) The long-distance modes are guided waves composed of both shear and longitudinal waves that travel at different speeds. For each guided wave mode, the wavespeed is a strong function of frequency. The waveguides in a complex system are thus highly dispersive; an AE signal will appear greatly distorted after it has propagated some distance. (2) The PZT (lead–zirconium–titanate) transducers normally used to measure the elastic wave are highly resonant and therefore may further distort the

† Rayleigh and Lamb are acoustic waves confined to a waveguide. For Rayleigh waves, the waveguide is the air–solid boundary itself while the waveguide for Lamb waves is an infinite plate. The theory of acoustic guided wave modes has much in common with microwave transmission line theory (see Auld [7, Chapter 10]). However, for acoustic waveguides both longitudinal (or compression) waves and transverse shear waves (with vertical and horizontal polarization) are present and propagate with different speeds even in an isotropic solid, whereas isotropic media can support only transverse electromagnetic waves. Guided wave modes are the product of the interaction of shear and longitudinal waves and the conversion of one type to the other at the waveguide boundary.

signal. (3) The AE signals can undergo multiple reflections between source and transducer. Even without reflections, the AE signal can propagate through different paths (multipath) between source and transducer. (4) The difficulty of detecting the signal in the presence of high levels of background noise. This is essentially a composite signal detection problem (see Srinath and Rajasekaran [8, pp. 83–88]): the AE event location must be treated as an embedded parameter to be estimated during signal detection. It is expected that much of the noise will be present at frequencies far below the PZT transducer resonance; this is one reason why these transducers are chosen for field use. It is likely, however, that in an environment with high levels of colored noise, the false alarm rate will be too high without some means of accommodating the noise, or using additional information on when and where AE is likely. This approach was taken in McBride and Maclachlan [9], where AE was used to detect crack growth in a test jig in a CF-100 fighter plane. The "real AE" was flagged from the available candidates (based on arrival time) by using g-loading information provided by accelerometers.

Some aspects of classical signal-processing (deconvolution) techniques for AE analysis are now discussed. The signal response in a complex system is so complicated that modelling it from first principles is impractical. This composite response of a sensor output to an AE event is, however, essentially linear (though non-linearities enter into the "inversion" process of real interest, as discussed below). Structures such as bridges are designed so that deflections should not be large—there is therefore no geometric non-linearity. The whole object of this monitoring is to detect faults before irreversible (i.e., plastic) deformations occur (i.e., while the structural material is still in its linear elastic range). The structural response is also causal and shift invariant. At time t, the net effect of all AE in the system on the output of a transducer located at the position $\mathbf{r}$ can therefore be written as the space–time convolution

$$\int_0^t \int_{structure} h(\mathbf{r}, \mathbf{x}, t-\tau)\beta(\mathbf{x}, \tau)\, \mathrm{d}\mathbf{x}\, \mathrm{d}\tau, \tag{1}$$

where $\mathbf{x}$ is the AE location (dummy) variable, $\beta(\cdot)$ characterizes the distribution of dilatation caused by AE as a function of position and time, $h(\cdot)$ represents the composite structural waveguide-transducr impulse response, and τ and t represent time.

Although equation (1) has a relatively simple form, determining $h(\cdot)$ and $\beta(\cdot)$ in a complex structure such as a bridge would be extremely difficult owing to the irregular shape and varying materials of the components of the structure and the reflection of signals off of boundaries of different components. Further, the effect of ambient temperature, wind speed and direction, and traffic conditions on the signals received at the sensors is extremely complex and possibly non-linear (an example of a non-linear relationship might be the signals induced as a function of wind direction [angle], since wind from certain angles may induce "rattling" noises that are not generated at all from a wind of the same speed but different angle). However, given $h(\cdot)$ and $\beta(\cdot)$, AE event location can, in principle, be determined by deconvolving the transducer outputs to obtain the AE inputs, as in Zgonc *et al*. [10], and (presumably) applying some sort of threshold test to determine location. This location deconvolution step, which is highly non-linear, was not automated in this reference; location was inferred by human observation of the deconvolved signals.

As a practical matter, this deconvolution is quite complicated. The system/transducer composite transfer function $h(\cdot)$ must first be obtained and then "inverted". This problem has been addressed for a relatively simple object by directly calibrating the AE source-to-transducer response [5, 10–12]. In this work, the end-to-end source-to-transducer response was calibrated by simulating an AE event, i.e., breaking a pencil lead or dropping steel balls on a laboratory specimen, and directly observing the transducer output. To obtain the inverse transducer-to-source response, a direct deconvolution approach was used in Zgonc *et al*. [10]. An auto-associative recall technique, (which is essentially a comparison of the observed signal with stored templates via a cross correlation) was used in Zgonc *et al*. [12], who also demonstrated that AE events in a plate could be located using only two transducers, versus the usual three required for triangulating the AE location using first-time-of-arrival techniques (Scott [13, p. 57]). This improvement was obtained by utilizing the information in the complete signal time history, instead of just the ‘‘rising edge’’.

In the above approaches, the AE inputs were obtained by performing linear operations on the transducer data. (In the associative recall techniques, this applies to the operational mode after the learning period was completed.) It should be emphasized that the final stage of the problem, i.e., automatic detection and *location* of the AE, was not addressed in the above studies. (The deconvolved signal strength was high enough to infer the AE location ‘‘by eye’’ in Zgonc *et al*. [12]). This operation, if done ‘‘optimally’’ in a real-world environment (with noise and other effects), would be non-trivial and involve non-linear operations on the sensor data. Thus, although these approaches represent a significant step in AE analysis, they suffer from at least three drawbacks in applications to large structures such as a bridge:

(1) The characteristic wavelengths of the elastic waves that carry the AE signal can be small (a few centimeters) compared to the dimensions of the structure. Techniques that perform linear operations on the raw predetect data will not be able to locate the source of the AE emission by ‘‘interpolation’’ if the test point spacing (where known—or simulated—AE events are introduced to create the above-mentioned template) is larger than the AE wavelengths. The harmonics of the predetect AE signal will have zero correlation with the corresponding harmonics of a stored template if the path length differs by a quarter wavelength.† The number of test points required for a structure such as a bridge (with dimensions on the order of tens of meters) is thus prohibitive.

(2) The Zgonc, Grabec and Sachse auto-associative recall techniques are fundamentally not time invariant. The data must be preshifted to align with the stored templates [12].

(3) If the spectrum of the noise and other non-AE contributions overlap the AE signal spectrum, these techniques may produce an unacceptably high false alarm rate. Zgonc *et al*. [10] explicitly considered the noise in their analysis, but to obtain their Weiner–Hopf integral equation, the noise was assumed to be stationary and Gaussian.

3. NEURAL NETWORK APPROACH

While the advantages of AE analysis seem attractive, the complexities and likely non-linearities involved in applying the above-mentioned classical approaches to AE detection in a complex system are daunting. This leads us to propose the use of an artificial NN to represent the relationships between the sensor data and the AE events instead of more direct deconvolution-based methods. Our fundamental aim with the NN is to maximize the likelihood of correctly locating a significant AE event. NNs are well recognized as appropriate function approximators for complex non-linear (possibly non-continuous) relationships (e.g., Barron [15], Park and Sandberg [16], and Funahashi

† This is analogous to the correlation structure obtained by correlating a predetect radar signal with a stored template, as discussed in Skolnik [14, p. 470]. As shown in Figure 10.11 of this reference, the amplitude of the correlation function has nulls and peaks separated by one-quarter wavelength. The corresponding correlation function for post-detected signals does not have this fine-grain structure.

[17]). As an example of the potential capability of NNs, we expect the NN to be able to use pulse arrival times to estimate the AE location (a non-linear operation consisting of recognizing and characterizing the signal envelope and thresholding). The laser test point spacing in this scenario can be far larger than the wavelengths of the elastic waves (see Cross *et al*. [18, p. 293]). Hence, this non-linear interpolation has considerable practical advantages over the direct linear approach using the raw signal data where the number of laser test points required would be prohibitive (see section 2 above; of course the direct linear approach has other problems as well, as already discussed).

Other possible benefits of the NN versus a deconvolution approach include comparative ease of application (the NN does not require many of the steps described in section 2); convenience of training and easy availability of training data (discussed below); the ability to use environmental data in the training process (discussed below); and the availability of powerful stochastic approximation tools for estimating the weights on a large NN (discussed below).

In practice, the NN will operate by taking recent sensor measurements and predicting the location and intensity of an AE event (if any) that occurred immediately prior to the current time. Likely sensor data include sampled signals from the PZT transducers (with sampling times on the order of half of the period of the AE traveling waves, i.e., sampling times of roughly 1 μs) and environmental data.

4. IDENTIFICATION OF NEURAL NETWORK MODEL

This section outlines the proposed training process for constructing the NN that will take real-time sensor data and infer the most likely location and intensity of an AE event. This training process relies on the use of a ‘‘Q-switched’’ laser that simulates an AE event in the system. The basic idea will be to pulse the system at many locations with the laser and then use these known simulated events together with appropriate sensor data to train a NN (or perhaps to train several NNs, each corresponding to one of the major substructures of the system). Aside from the PZT transducer readings, the vector of sensor data will include measurements related to other relevant effects, such as traffic and weather conditions for a bridge. It is expected that by including such information in the training process, there will be a better chance of isolating an AE signal when the trained NN is put into service (versus simply treating these other effects as ‘‘noise’’). Simulating AE in a system at various locations using a Q-switched laser is similar to what has been done in many laboratory studies (e.g., Scott [13, pp. 225–226]). A Q-switched laser works by manipulating lasing cavity resonance to produce a very short, high-power burst of laser energy. This does not damage the structure, yet it is a faithful simulator of many of the acoustic signals that would be generated by a real AE [13, pp. 225–226]. These simulated AE signals will then be used in the training of a NN. The training data will consist of input–output pairs consisting of sensor (including transducer) signals (input) and the desired NN output, which would be the simulated AE location and the time history of the simulated AE at that location. The energy of the laser pulse has a direct equivalence to the potential energy release of an AE event. The time history would include such intensity information.

For a given NN structure (say, number of layers and nodes), the training process is a non-linear parameter estimation problem involving the determination of the NN weights from a set of training data. The training data for this weight estimation will be collected from non-destructive tests on the system in a representative variety of environmental conditions. To increase predictive ability, it may be desirable to have separate NNs for different categories of environmental conditions, e.g., separate NNs could be used for ambient temperatures below and above 5°C (the exact temperature would still be an input variable to each of the NNs). Consistent with the discussion in section 3, let $y(t)$ represent the vector of sensor data up to time t, from all the relevant sensors on the system. This may represent a "sliding window" of raw sampled data (representing some fixed number of most recently obtained data points) or a summary set of features of the data such as arrival time, signal magnitude, dominant frequencies, etc. (obviously, using just features allows for a large reduction in the dimension of $y(t)$ at the possible expense of a loss of relevant information). From $y(t)$, we wish to train the NN to estimate $z(t)$, a (perhaps) four-dimensional vector of intensity (one component) and position (three components) for any AE event that occurred within (or immediately prior to) the sliding window ending at time t (it does not appear that one would be especially interested in determining the exact time within or before the window that the AE occurred, which eliminates at least one quantity that might otherwise need to be estimated).

The tests used in creating the set of training data will be composed of both AE and non-AE events. For the AE events, the above-mentioned Q-switched laser will be used. In such cases, a simulated AE of known intensity will be imposed at various locations on the structure—providing a known $z(t)$ vector—and corresponding data $y(t)$ will be collected. In cases where there is no simulated AE event, $y(t)$ data will likewise be collected under a variety of traffic and environmental conditions. For the corresponding known $z(t)$, a zero intensity level will be recorded together with a "dummy" location (perhaps not even on the system). It is expected that, after training is complete, the use of the dummy location will enhance the capabilities of the NN to avoid false alarms about possible AE events, since both intensity and location outputs would point toward the "no AE" case. Training will occur by adjusting the NN weights to minimize some loss function related to the prediction error $\hat{z}(t) - z(t)$ over all available $y(t)$, $z(t)$ pairs in the set of test data, where $\hat{z}(t)$ is the NN-predicted AE source signal (the sum of squared errors is a likely loss function). Implicit in this process is the assumption that no significant real (but unknown) AE event occurs during the sliding window in the training set. A conceptual illustration of the training process is shown in Figure 1.

In the case where a sliding window of raw signal data are used for $y(t)$, the dimension of the $y(t)$ vectors will be very large (equal to the number of PZT sensors times the number of data points in the sliding window environmental, etc. information). Hence, it will be necessary to use efficient optimization methods for very high dimensional problems (the

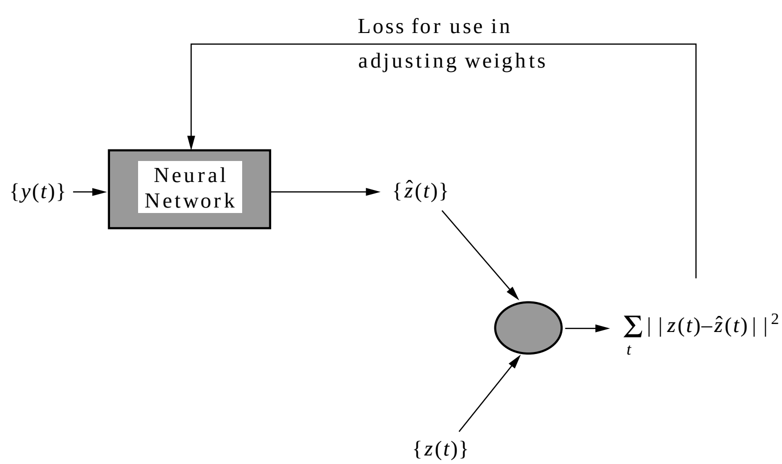

Figure 1. Illustration of NN training process with sum of squared errors loss function.

number of NN weights will significantly exceed the number of elements in $y(t)$). As an indication of the problem dimension, if there is a sliding window length of 100 μs (approximately the "ring down" time of AE signals; see Fowler and Papadakis [19, p. 231]) and a sampling time of 1 μs, each PZT sensor accounts for 100 elements in $y(t)$. A typical bridge substructure, for example, may have a dozen or so PZT sensors and so a modest-sized NN for the substructure with only 50 nodes in the hidden layer following the input layer will typically have many *more* than 60 000 weights!

A powerful technique for high dimensional optimization is the simultaneous perturbation stochastic approximation (SPSA) algorithm in Spall [20]. The SPSA iterative algorithm requires only *two* NN output error evaluations (predicted $z(t)$ versus actual $z(t)$) at each iteration (versus the $2 \times$ problem-dimension evaluations needed for standard finite-difference methods, i.e., more than 120 000 evaluations in the case mentioned above) to construct an approximation to the gradient of the error function. (Spall [20] proves under general conditions that the accuracy of SPSA and finite-difference methods is asymptotically the same for a given number of iterations). This yields enormous savings in computation in such a high dimensional problem. Note that in contrast to steepest-descent-type (backpropagation) algorithms (e.g., Narendra and Parthasarathy [4]), SPSA requires no calculation of the gradient vector for use in optimization, which may be very difficult in such a high dimensional problem, especially for richer NN structures such as the recurrent forms that include internal feedback. SPSA has been thoroughly tested in many estimation problems (although the authors are not aware of applications for problems with dimension as large as mentioned above)—see, e.g., Spall and Cristion [21]. The theory does not indicate any fundamental problems arising as the problem dimension grows very large. Further, a global (versus local) optimization implementation of SPSA is discussed in Chin [22], which may be especially relevant in such high dimensional problems.

5. EXPERIMENT WITH STEEL I-BEAM

The authors have tested the NN-based procedure on a 120-cm steel I-beam mounted on a stable (optical) bench at the Department of Materials Science of Johns Hopkins University. The experiment consisted of pulsing the I-beam with (Q-switched) laser-induced AE signals at various locations and seeing whether a trained NN could accurately recover the location of the AE signal. Both a training and a test set of data were collected.

Figure 2 depicts the location of the six PZT sensors and 11 laser-induced AE sources on the center flange of the I-beam. Seven of the AE sources are used for the NN training data and four sources are used for testing the ability of the NN to predict the location of the source. The sources are located every 6·3 cm along the centerline of the flange as measured from the left edge of the I-beam. For each sensor, a sliding window of 60 000 data points were collected at 1-μs intervals via a digitizing oscilloscope. For simplicity in this experiment, we used only arrival time information from this large collection of data points as the input ($y(t)$) to the NN; other "features"—such as signal magnitude, frequency, and leading edge shape—may be useful in other applications (and still provide for a large reduction in the size of the input vector for the NN relative to using the raw sampled data). The arrival times were the differences between the time that the laser was pulsed and the time that the signal arrived at the sensor; hence the NN is essentially using the *relative* arrival times at the different sensor locations to infer the AE location.

The NN used here was a simple feedforward network with one hidden layer. There were six input nodes (one for each sensor's arrival time value), 25 nodes in the hidden layer,

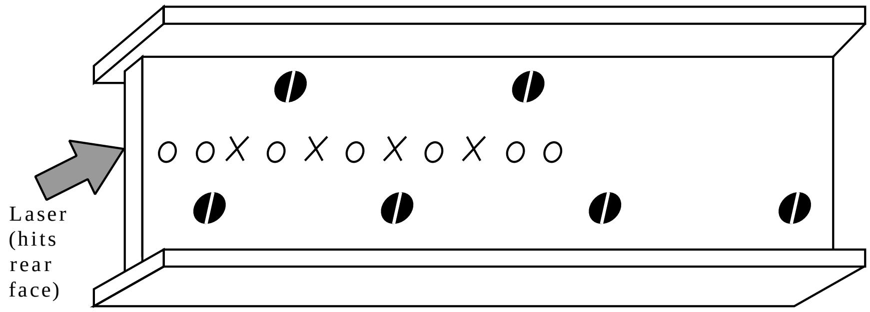

Figure 2. Location of sensors and laser-induced AE sources on steel I-beam; ○, AE source used for training; ×, AE source used for testing; ●, sensors.

and one output node (the location of the AE source along the center line of the I-beam). The NN included bias weights at each hidden and output node, leading to a total of 201 weights to be estimated.

For the NN training, we used the MATLAB Neural Networks Toolbox running on a Pentium-based PC. The standard backpropagation (gradient descent) algorithm was used for the training as given in the LEARNBP routine within the Toolbox. The input data contained a total of 42 elements: the arrival times at each of the six sensors for the seven AE source locations. The NN weights were initialized randomly according to the default mechanism in LEARNBP, and the learning rate (0·01), stopping criterion (sum of squared errors $\leqslant 0{\cdot}01$), and maximum number of passes through the data (500) were also set to their default values.

For testing the trained NN's ability to recover the location of an AE source, the arrival times at the six sensors for the four test points in Figure 2 were fed into the trained network. Table 1 shows the true locations of the AE sources and the corresponding predictions produced by the trained NN. It can be seen that the predictions are quite accurate. One would expect that the accuracy of the approach could be improved if data from additional AE sources were collected for the training process; the authors found little improvement when additional features of the existing seven AE sources were used (such as frequency information, etc., as mentioned above), although in other problems (e.g., those with greater noise effects such as on a bridge) such information may be useful.

TABLE 1

Comparison of actual AE source locations and NN predictions (values are in cm from left edge)

Actual location	Predicted location
18·9	20·3
31·5	31·0
44·1	43·0
56·7	57·3

6. OPEN ISSUES AND CONCLUDING REMARKS

As is clear from the above discussion, work remains in taking this idea from the concept stage to a practical implementation on a large-scale system such as a bridge. Some of the most prominent of these issues are the following.

The need for statistical inference procedures to determine the significance of a predicted AE intensity and location. This is critical for avoiding excess false alarms, (which would impose significant costs on maintenance personnel), while maintaining adequate sensitivity to potentially catastrophic growing cracks.

Determining optimum placement of the transducers and other sensors on the system to maximize the amount of information available relative to a potentially serious AE. This is related to the issue of optimal experimental design. A relatively easy-to-implement approach is described in Sadegh and Spall [23]; this approach uses SPSA to avoid the complex modelling (as in expression (1)) typically needed in optimal sensor placement.

The need for efficient transmission of the sensor data to a central (remote) facility. This involves the movement of a large quantity of data from, e.g., a bridge at frequent intervals, and the ability to process these data as well as the data coming in from other bridges. As mentioned in section 2, it is expected that such data can be efficiently transmitted via fiber optic cable, but this issue needs to be examined more thoroughly in light of practical cost and implementation concerns.

Choice of type of NN (feedforward, recurrent, etc.) and associated structure and training algorithm. This is a standard (but important!) concern for any NN application. Related discussion is in section 4. An additional issue connected to the choice of NN is the development of means for including in the NN prior information on the system's frequency response, when available (e.g., it may be feasible to model localized response under ''quiet'' traffic and environmental conditions (on a bridge) using the convolution form in expression (1)). The inclusion of such prior information in a NN is often not readily automated and needs to be addressed on a case-by-case basis.

Possible problems in the use of the PZT transducers for AE events. One of the potential problems pertains to the long-term adhesion of the sensors to the system. If a bond is weakened, this can reduce the effectiveness of the transducer and/or alter its frequency response from that expected based on the initial training data.

Difficulties in obtaining useful readings for the traffic conditions in a bridge application. Given the relatively small wavelengths of the dominant AE signals (on the order of a few centimeters), it will be necessary to either very accurately know the instantaneous position (and perhaps other characteristics) of vehicles on the bridge or else have auxiliary AE sensors that operate at the less dominant lower frequency range of the AE spectrum. In the latter case, knowing vehicle positions to only within a few meters would allow for the separation of vehicle and AE effects in this lower frequency range (in the higher frequency range of the PZT transducers, the traffic effects would have to be treated as noise).

Recognition of the long-term aging effects of the system on the ability of the NN to make accurate predictions. Periodically, it would seem necessary to recalibrate the NN weights to account for such aging.

Despite the issues mentioned above, it seems that the framework outlined in this summary has promise for practical NDE of complex systems. It largely makes use of existing sensor and communications technology and appears to address some of the seemingly insurmountable inverse modelling problems associated with such complex systems. Further, the approach is aimed at reducing long-term costs of monitoring through the possible reduction of the need for engineering inspection teams to conduct active tests of the system, as required, say, in ultrasound or X-ray inspection (of course, some visual and other inspection will always be necessary to augment such "automatic" NDE schemes).

ACKNOWLEDGEMENTS

This work was partially supported by the JHU/APL IRAD Program and U.S. Navy Contract N00024-97-C-8119. We wish to thank Dr Jane Spicer of JHU/APL for helpful discussions, and Prof. James Spicer of the JHU Materials Science Department, who provided laboratory facilities and assistance.

REFERENCES

1. N. J. Jones and B. R. Ellingwood 1993 *Proceedings of the Conference on Nondestructive Evaluation of Bridges, Arlington, VA*, Chapter 2. NDE of concrete bridges: opportunities and research needs.
2. J. M. Plecknik and O. Henriquez 1993 *Proceedings of the Conference on Nondestructive Evaluation of Bridges, Arlington, VA*, Chapter 4. Composite bridges and NDE applications.
3. C. M. Teller 1993 *Proceedings of the Conference on Nondestructive Evaluation of Bridges, Arlington, VA*, Chapter 5. The state-of-the-art in nondestructive evaluation of steel bridges.
4. K. S. Narendra and K. Parthasarathy 1991 *IEEE Transactions on Neural Nets* **2**, 252–262. Gradient methods for the optimization of dynamic systems containing neural networks.
5. I. Grabec and W. Sachse 1989 *Journal of Acoustic Emission* **8**, S20–S23. Solving AE problems by a neural network.
6. H. Yuki and K. Homma 1992 *Journal of Acoustic Emission* **10**, 35–41. Analysis of artificial emission waveforms using a neural network.
7. B. A. Auld 1990 *Acoustic Fields and Waves in Solids*. Malabar, FL: Krieger Publishing Company; second edition.
8. M. D. Srinath and Rajasekaran 1979 *An Introduction to Statistical Signal Processing with Applications*. New York: Wiley.
9. S. L. McBride and J. W. Maclachlan 1985 *Journal of Acoustic Emission* **4**, S151–S154. AE monitoring of aircraft structures.
10. K. Zgonc, I. Grabec and W. Sachse 1991 *Proceedings of the 4th World Meeting on Acoustic Emission*, 134–140. Acoustic emission analysis by the optimal multidimensional deconvolution.
11. I. Grabec and W. Sachse 1989 *Journal of Applied Physics* **66**, 3993–3999. Experimental characterization of ultrasonic phenomena by a learning system.
12. K. Zgonc, I. Grabec and W. Sachse 1993 *Journal of Acoustic Emission* **11**, 79–84. Solution of simple inverse source characterization problem using associative recall.
13. I. G. Scott 1991 *Basic Acoustic Emission*. New York: Gordon and Breach Science Publishers.
14. M. I. Skolnick 1962 *Introduction to Radar Systems*. New York: McGraw-Hill.
15. A. Barron 1994 *Machine Learning* **14**, 115–133. Approximation and estimation bounds for artificial neural networks.
16. J. Park and I. W. Sandberg 1993 *Neural Computation* **5**, 305–316. Approximation and radial basis functions.
17. K. I. Funahashi 1989 *Neural Networks* **2**, 183–192. On the approximate realization of continuous mappings by neural networks.
18. N. O. Cross, L. L. Loushin and J. L. Thompson 1972 in *Acoustic Emission*, 270–296, ASTM STP 505. Bal Harbour, FL: American Society for Testing and Materials. Acoustic emission testing of pressure vessels for petroleum refineries and chemical plants.
19. K. A. Fowler and E. P. Papadakis 1972 in *Acoustic Emission*, 222–237, ASTM STP 505. Bal Harbour, FL: American Society for Testing and Materials. Observation and analysis of simulated ultrasonic emission waves in plates and complex structures.
20. J. C. Spall 1992 *IEEE Transactions on Auto. Control* **37**, 332–341. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation.
21. J. C. Spall and J. A. Cristion 1994 *Statistica Sinica* **4**, 1–27. Nonlinear adaptive control using neural networks: estimation based on a smoothed form of simultaneous perturbation gradient approximation.
22. D. C. Chin 1994 *Neural Networks* **7**, 573–574. A more efficient global optimization algorithm based in Styblinski and Tang.

23. P. Sadegh and J. C. Spall 1997 submitted (preliminary version in 1996 *Proceedings of the Test Technology Symposium*, sponsored by U.S. Army TECOM). Optimal sensor configuration for complex systems.