

Fault Detection in Steel-Reinforced Concrete Using Echo State Networks

Adam J. Wootton
Foundation year Centre
Keele University
Keele, United Kingdom
a.j.wootton@keele.ac.uk

Charles R. Day
School of Computing and Mathematics
Keele University
Keele, United Kingdom
c.r.day@keele.ac.uk

Peter W. Haycock
Foundation Year Centre
Keele University
Keele, United Kingdom
p.w.haycock@keele.ac.uk

Abstract— The cost of maintaining and repairing the world’s ageing reinforced concrete infrastructure continues to increase, and is expected to cost the United States economy alone \$58 billion by 2020. Consequently, the use of non-destructive testing technologies for the early identification of faults in roads and bridges is becoming increasingly important. One such technology is the Electromagnetic Anomaly Detection (EMAD) technique, which exploits non-destructive magnetic flux leakage to detect defects in steel reinforcing meshes embedded in concrete. Despite the increasing need for such techniques, the data analysis options currently in use are limited. This paper presents an application of Echo State Networks, a recurrent neural network from the field of reservoir computing that features a short-term memory, to data obtained using the EMAD technique. Having been trained to discern real defect signals from other anomalous magnetic features, the performance of the ESNs was then compared to that of an analytical data analysis technique that is currently used to process EMAD data. It was found that average ESN performance was comparable in terms of AUC, while the optimal threshold was more consistent, greatly aiding application in the ‘real-world’. A qualitative analysis of the output of both methods on an unseen testing dataset also demonstrated the superiority of ESNs for practical use as a real time tool for on-site inspections.

Keywords— *Echo State Networks, Reservoir Computing, spatially varying data, structural health monitoring, steel-reinforced concrete, magnetic flux leakage*

I. INTRODUCTION

In the United States alone, the cumulative cost to the economy of deteriorating pavements and bridges is estimated to reach \$58 billion by 2020 [1]. As of December 2014, around 24% of bridges in the United States were structurally deficient or functionally obsolete [2]. One technique that is used to monitor the condition of reinforced concrete roads and bridges is magnetic flux leakage (MFL), which can be used non-invasively to detect defects due to corrosion in steel reinforcing meshes. While methods such as correlation analysis [3], [4] and decision fusion [5] have been used to analyse the large volumes of multi-dimensional data arising from different MFL techniques, the most recent work has seen the data processing limited to a manual expert analysis of the raw data (see, for example, [6–8]). There is therefore a need for the sort of fast, automatic and accurate interpretation of the data that can be provided by computational intelligence (CI) techniques.




Figure 1: A typical ESN topology. The input units (left) are fully connected to the reservoir neurons via randomly weighted connections set at initialisation. The reservoir is sparsely interconnected with randomly weighted connections and there is potential for recurrent loops. Each reservoir node is connected to each output node (right) and these connections are the only ones that are trained.

This paper presents a method for processing MFL data using Echo State Networks (ESN). This builds on previous work by the authors, who have considered the application of CI to different MFL datasets in the past [9, 10]. After the preliminary consideration of the use of ESNs in this problem domain in [10], the work reported here focuses solely on the development of the optimal ESN configuration and training regime for spatially varying MFL data. The efficacy of this approach is demonstrated by a comparison with an analytical technique (AT) that is informed by engineering expertise and is currently used commercially for MFL data analysis. Both methods have been applied to a complex dataset gathered from an MFL scan of a real steel reinforcing mesh embedded in concrete.

The remainder of the paper is organised as follows. Section II introduces ESNs, while Section III explains the theory behind the MFL data acquisition technique used here before detailing the data collection and ESN development regimes used. Section IV presents the results obtained when using the ESN in comparison to a standard AT and Section V concludes the paper.

II. ECHO STATE NETWORKS

ESNs are a type of recurrent artificial neural network in the field of reservoir computing [11]. What sets ESNs apart from many other neural networks is the existence of a sparsely interconnected reservoir, since there are random recurrent

connections within this reservoir that allow components of time-series data stimuli to fadingly resonate throughout the reservoir, effectively giving the network a short term memory. A typical ESN topology can be seen in Fig. 1. The input units are fully connected with the reservoir, whereas the reservoir units themselves are only sparsely interconnected, and all of these weighted connections are randomly generated at the initialisation of the reservoir and kept constant throughout. Only the weights on the connections between each reservoir unit and each output unit are trained, usually by a simple regression technique such as ridge regression [12–14]. The ESNs used here were trained and tested using the reservoir computing toolbox for MATLAB [15].

The ability of ESNs to exhibit a short term memory means that they have often been used to process time-series data in a number of fields, with recent examples including the construction industry [16], medicine [17] and structural health monitoring [18].

In the ESN architecture used here, at any time t , the output of the vector of ESN reservoir neurons, \mathbf{x} , was given by (1).

$$\mathbf{x}(t) = f((1-\delta)\mathbf{x}(t-1) + \delta(\mathbf{W}_{res}^{inp}\mathbf{u}(t-1) + \mathbf{W}_{res}^{res}\mathbf{x}(t-1))) \quad (1)$$

Here, \mathbf{W}_{res}^{inp} is the input to reservoir weight matrix, f is the tanh activation function that was used (other commonly used activation functions include the sigmoid function), $t-1$ is the previous time step, \mathbf{W}_{res}^{res} is the reservoir weight matrix, $\mathbf{u}(t)$ is the vector of the input data at time t and δ is the leak rate.

One advantage of ESNs is that there are several tuneable parameters, allowing networks to be configured for particular tasks. For example, although \mathbf{W}_{res}^{inp} is randomly generated, it can be uniformly scaled according to the task, with a greater input scaling increasing non-linearity and increasing the relative effect of the input at time t compared to past inputs [19]. In this case, the scaled input weights were calculated by (2).

$$\mathbf{W}_{res}^{inp} = \rho \times \mathbf{W}^{inp} \quad (2)$$

where \mathbf{W}^{inp} represents the input to reservoir weight matrix prior to scaling and ρ is the input scaling.

Similarly, the randomly generated internal reservoir weights can be scaled in order to adjust the length of the ESN's short term memory. This can be done by multiplying the initial reservoir weight matrix by a spectral radius scaling factor and dividing by the maximum eigenvalue of the initialized weight matrix [11], as shown in (3).

$$\mathbf{W}_{res}^{res} = \alpha \times \mathbf{W}^{res} / \lambda_{max} \quad (3)$$

Here, α is the spectral radius scaling factor and λ_{max} is the maximum eigenvalue of \mathbf{W}^{res} , which represents the initial reservoir weights. The spectral radius scaling factor is usually limited to a maximum value of 1 in order to ensure the echo state property for the network, and it can be seen in (3) that as the spectral radius decreases, so too will the resonance of past inputs in the reservoir, since the values of the internal reservoir weights will decrease.

Other parameters that can be adjusted include the reservoir size (i.e. the number of neurons in the interconnected reservoir)




Figure 2: Typical triaxial MFL data obtained by the EMAD when mapped to the spatial domain (top) and the orientation of the axes of the EMAD (bottom). The regions of the signal pertaining to defects are highlighted by the black boxes. Note that for both defects, a peak in the z axis data is accompanied by a negative slope in the x axis data.

and the connectivity factor of the reservoir, which controls the sparsity of the connections between reservoir units. The final tuneable parameter, the leak rate (δ), acts as a time constant, speeding up or slowing down the reservoir dynamics. The effect of the leak rate on a reservoir unit can be seen in (1).

III. THE ELECTROMAGNETIC ANOMALY DETECTION TECHNIQUE

The MFL technique used in the work reported here is the Electromagnetic Anomaly Detection (EMAD) technique, developed over the course of several years at Keele University [20–23] and licensed commercially. The EMAD technique is designed to allow the user to scan a two dimensional steel-reinforced concrete surface, such as a road or bridge, and locate accurately any underlying damage due to corrosion of the embedded reinforcing steel bars (rebars).

A. Theory and Statement of Problem

Like some (such as [24]), but not all (see, for example, [4], [7]), MFL methods, the EMAD technique exploits remanent magnetism in order to locate defects. Consequently, a typical survey will see the area to be surveyed split into evenly spaced scan lines that are approximately aligned with either the transverse or longitudinal bars in the underlying mesh (depending on where corrosion is most likely to be seen) before a powerful electromagnet ('the energiser') is passed over each one of these scan lines in turn. This leaves the mesh in a magnetised state. A triaxial electromagnetic probe is then passed over each line, recording the magnetic flux in three components. Data are captured at regular angular rotations of a rotary encoder connected to the rear wheel of the EMAD and can then be readily mapped from the temporal domain to the spatial domain after analysis. The aim of this is to detect anomalous signals that are likely to be indicative of rebar corrosion.

Fig. 2 shows a typical set of data for a triaxial scan of a steel reinforcing bar (rebar) featuring two defects, the signals for which are highlighted. It can be seen that a typical defect signal is distributed over a succession of consecutive data points and is represented by a peak in the z axis data and a corresponding negative gradient in the x axis data. If the polarity is reversed, then a defect could also be represented by the inverse of this: a trough in the z axis data and a corresponding positive gradient in the x axis data. Although the y axis data is recorded, it can be seen to be largely unaffected by the presence of defects.

In theory, the only magnetic flux that would be detected when scanning an ideal rebar with only one defect would be in the region around the defect. However, a real world dataset will always feature magnetic anomalies and noise not caused by defects. In Fig. 2, for example, the ripples in the x and z axis signals seen in the region 0.5 – 1.5 m are simply background noise caused by the presence of transverse bars. The large anomaly to the right of the highlighted defects, located at around 3.5 – 4 m, is an 'end effect', which is a natural consequence of the energisation procedure seen at the end of every scan line, due to a large magnetic pole where the energisation ends, and is not indicative of damage. The challenge for any method of analysing EMAD data is to correctly identify defect signals while ignoring noise and end effects.

The current standard method for processing real world EMAD data is a simple AT that has been specifically derived to exploit the ideal relationship between x axis and z axis data for a clean break. This AT is expressed in (4).

$$O = \frac{P \times G \left(z_0 \times \frac{dx}{dd} \right)}{\bar{O}} \quad (4)$$

In (4), O is the output of the AT for any given data point, P is the output of a peak detection algorithm that calculates the difference in probe reading between consecutive peaks and troughs in the x axis data, z_0 is the value of the z component data, $\frac{dx}{dd}$ is the value of the x axis data differentiated with respect to distance, G is a Gaussian smoothing function and \bar{O} is the mean of the output.




Figure 3: The z and x axis data from Fig. 2, with the output of (4) for each datapoint plotted over the top. Although there is a positive output for each defect signal, the first defect is lost amongst some earlier noise and the largest output is for the end effect located at around 4 m.




Figure 4: The layout of the test bed reinforcing mesh. Red diamonds indicate the position of manually inserted defects.

While the AT is well suited to detecting defects in the sort of ideal rebar mentioned above, or in a dataset with some small levels of noise, it often struggles to discriminate between defects and end effects. An example of this is shown in Fig. 3, which depicts the x and z axis data from Fig. 2, along with the output of the AT at each point. A greater signal should indicate a greater likelihood of corrosion. While the second defect is correctly identified, there is a slight lag due to the Gaussian smoothing. Furthermore, the first defect is lost amongst noise and the largest response comes from the end effect. This is due to the fact that while the AT is designed to detect ideal defect signatures, the defect signatures in this real world dataset are not ideal. While the peak in the z axis data at approximately 2 m is apparent to an expert eye, it is not distinct enough for the AT to pick it up clearly. This is where CI, and ESNs in particular, can offer a solution. While the AT can identify defects but not distinguish between these defects and unwanted noise, the short-

term memory of ESNs enables them to recognise the temporally-extended characteristic signature of defects while learning to ignore noise and end effects.

B. Data Collection

The data used in this study were all taken from the same source: a reinforced concrete test bed designed to simulate corrosion, created as part of a PhD project [20]. The test bed featured a 2.2 m wide by 4.6 m long steel reinforcing mesh, consisting of 5 mm thick rebars separated by 195 mm both longitudinally and transversely, encapsulated in concrete. Prior to encapsulation, a number of defects were manually created in the mesh. These defects took the form of either a clean break in a rebar, accelerated corrosion to a part of the mesh or a combination of both of these. Further corrosion was encouraged by immersing the encapsulated mesh in a sodium chloride solution. A total of twelve defects were initially inserted, but three of the corrosion product defects were lost after encapsulation. A scanning area consisting of ten longitudinal lines, each aligned with a longitudinal rebar, was set up, the layout of which, along with the location of the nine remaining defects and the transverse rebars, can be seen in Fig. 4.

The usual scanning procedure, as outlined in the previous section, was employed to obtain data from the test bed. A total of 17 different datasets (equivalent to 170 separate scan lines, or 153 340 individual time-series data points) were collected over a two year period. No two datasets were the same due to the degradation of the mesh over time and variations in the energisation and scanning procedure. For example, while in most cases the energisation took place from the left to the right of Fig. 4 along each longitudinal bar, in some cases the direction of energisation was reversed, reversing the polarity of the mesh. However, it was manually verified that the relevant defect signals could be seen in each dataset.

C. ESN Approach

In order for ESNs to be used here, a suitable training regime had to be formulated. A total of 82 different individual scan lines from datasets gathered over a number of years were used for training. These were then categorised as either a line featuring a known defect, or a line without any known defects. This was so that during training, the ESNs would be presented with an equal number of lines featuring defects and lines without defects, which would prevent overtraining on either. Each scan line was then labelled so that at points without a defect the ESN had a target output of -1 and at points with a defect the ESN had a target output of +1. During training, each scan line was presented to the ESN in turn and 41-fold cross validation was used in accordance with the number of samples available for training (i.e. two samples per fold). Two datasets, henceforth referred to as dataset A and dataset B, were kept back so that they could be used as unseen testing data.

However, before the ESNs could be compared to the AT, a good set of ESN configuration parameters needed to be found for this problem domain. To this end, a grid search was performed, the results of which are given in Table 1.

TABLE I. THE OPTIMAL VALUES FOUND FOR EACH TUNABLE ESN PARAMETER

Parameter	Range Tested	Value Used
Spectral Radius	0 – 2	0.99
Input Scaling	0 – 2	1.8
Leak Rate	0 – 1	0.05
Adaptation Epochs	0 – 10	1
Reservoir Size	1 – 500	475
Reservoir Connectivity Factor	0.1 – 1	0.65
Activation Function	Tanh, Lorentzian, Triangular Basis, Radial Basis, Fermi	Tanh

500 ESNs with the optimal configuration were then trained, and their performance when presented with the two unseen test datasets compared to the performance of the AT on the same test data. The reason for training 500 ESNs was that the weighted connections between reservoir neurons are randomly generated, and an individual ESN may be particularly well (or poorly) configured, meaning that the performance of that one ESN would not be representative of the typical performance that one might expect to obtain. The results from training 500 ESNs would be more representative of the typical performance in this problem domain.

Since this is a two-class (‘defect’ and ‘no defect’) problem, performance was assessed using a Receiver Operator Characteristic (ROC) curve. Analysis of an ROC curve, where sensitivity is calculated against False Positive Rate at several thresholds, provides a richer measure of classification accuracy due to de-coupling of classifier performance from class skew and error costs, while the area underneath this curve, the Area Under Curve (AUC), is equivalent to the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative instance [25]. An AUC value of unity is indicative of perfect classification, while a value of 0.5 suggests that the results from the classifier are little better than guessing. In the ‘real-world’, the processed EMAD data is normally presented on a thresholded contour plot, so an evaluation method that considers performance at different thresholds is particularly useful. The average AUC for the ESNs was therefore compared to the AUC for the AT on both of the unseen test datasets.

One practical consideration for use of the technique in the ‘real-world’ is the consistency of the optimal threshold, which can be found by determining the point on the ROC curve that had the smallest Euclidean distance to the point that would represent perfect classification, (0,1). When the ground truth is known, as in this case, finding the threshold that produce the clearest contour plots for defect location is trivial. However, in the ‘real-world’, the ground truth is not usually known, meaning that expert analysis is often required when setting the best threshold. If the calculated threshold is consistent across different datasets, then expert analysis is not as important. It was, therefore, considered that for ESNs to be suitable for

analysing data taken with the EMAD, the average optimal threshold should be consistent for both datasets.

Following the ROC analysis, a qualitative analysis of the output of the best performing ESN and AT was also performed.

IV. RESULTS AND DISCUSSION

The two unseen test datasets were presented to the AT and 500 ESNs in turn, with the resulting average AUC and optimal threshold given in Table 2.

TABLE II. AVERAGE NUMBER OF FALSE POSITIVES FOUND BY 500 TRAINED ESNs AND THE AT ON DATASETS A AND B, WITH THE STANDARD DEVIATION FOR THE ESNs GIVEN IN BRACKETS

Classifier and Dataset	AUC	Optimal Threshold
ESNs on Dataset A	0.9369 (0.0187)	0.0683 (0.0076)
ESNs on Dataset B	0.9367 (0.0163)	0.0666(0.0061)
AT on Dataset A	0.9578	4.3468
AT on Dataset B	0.9436	0.9276

The values in Table 2 represent good performance by both techniques on both of the test datasets. On average, the AT marginally outperformed the ESNs in terms of AUC, although it can still be said that the ESNs provided very competitive performance. The fact that the ESNs achieved a high value for the AUC suggests that it was able to learn to model the temporal relationship between the x and z axis data. The ability to do this is key for any approach that seeks to interpret EMAD data.

The results for the optimal threshold, where the ESN can be said to have outperformed the AT, are equally interesting. The difficulty in judging an appropriate threshold when using the AT is made clear by the discrepancy between the optimal threshold for dataset A, 4.3648, and the optimal threshold for dataset B, 0.9276, especially since the two datasets were obtained from the same reinforcing mesh. In direct contrast, the ESNs had, on average, very similar optimal thresholds of 0.0666 for dataset A and 0.0683 for dataset B.

A more qualitative comparison between the output for the AT and the best performing individual ESN reveals some interesting details. It should be noted that this individual ESN actually outperformed the AT on its own, giving an AUC of 0.9697 and an optimal threshold of 0.0757 when presented with dataset A, and an AUC of 0.9632 and optimal threshold of 0.0886 when presented with dataset B. The optimally thresholded output of each technique when applied to dataset A can be seen in the contour plots in Figure 5. The centres of the ground truth defect locations are marked with red diamonds, while the contours give damage locations indicated by the technique, in accordance with the colour bar on the right.

It is immediately apparent from Figures 5 and 6 that neither approach works completely perfectly. The ESN managed to detect all of the ground truth defects, but also erroneously indicated the presence of defects at ‘clean’ locations. The AT indicated slightly fewer false defects, but also failed to detect one of the real defects. This is a significant failing, since this

defect, located at approximately (2, 0.78) was a clean break in the steel. Furthermore, the signal for the defect at (3.6, 1.755) is very weak compared to some of the other signals, while the defect at (3.9, 0.975) is lost amongst a number of false positives.

Had these results been reported as part of a typical ‘real-world’ survey, where the ground truth would not have been known, using the AT output seen in Figure 5 would lead to a very different set of recommendations to using the output of the ESN. Expert analysis of the output of the AT would probably lead to the suggestions that the large anomalous signals at (0.2, 1.755), (1.1, 1.755) and (2.5 - 2.75, 1.755) are the most likely regions for defects. The signals seen at (0.4, 1.17), (1.3, 1.17), (1.4, 1.755) and (3.3, 0.78) would probably also be treated as potentially significant signals, but less likely to indicate defects than the first three areas. The large, anomalous signals in the region (3.6 - 4.1, 0.585 - 1.755) would be treated as ‘end effects’, and would likely be ignored. The small signals at (0.25, 0) and (3.5, 1.755) would probably be treated as low-level noise, and discounted.

Based on this interpretation, even though the AT technically flagged up eight out of the nine defects, defects 1, 3 and 109 would all be missed, while defect 2 could also be missed due to the large offset between the actual defect location and the AT signal. In short, use of the AT would result in most of the defects being investigated, but also in a number of clean areas of rebar being examined and some actual defects being missed, reducing confidence in the efficacy of the EMAD technique.

In contrast, analysis of the output of the ESN would probably lead to the large anomalous signals at (0.4, 1.17), (1.25, 1.755), (1.45, 1.755) and (2.75, 1.755) being identified as potentially significant signals, with the large anomalies around (0-0.15, 0 - 1.755) and (4 - 4.25, 1.56 - 1.755) being discounted as end effects. The anomalous signals at (1.3, 1.17), (1.9, 1.755), (2.3, 1.755), (3, 0.78), (3.25, 1.755), (3.7, 1.755) and (3.75, 0.975) would also be identified as being potentially significant. The regular nature of the anomalies in the range (0 - 4.25, 1.375 - 1.56) suggests that they are probably due to magnetic interference from the transverse rebar on the mesh.

Analysis of Figure 6 leads to similar conclusions. The output of both techniques when applied to dataset B is very similar to their output when applied to dataset A, and the expert analysis for the ESN would not change. For the AT, the analysis may slightly differ, as the presence of a very large end effect signal at (3.8, 1.755) makes the signals for all of the genuine defects appear to be significantly smaller. This could lead to defects 2 and 4 also being lost. While the AT does, on average, marginally outperform ESNs in terms of AUC, the ESN used here would be of more practical use in the ‘real-world’.

V. CONCLUSION

In this paper, it was shown that after suitable training, ESNs are capable of providing performance comparable to that of an AT that has been used for real world industrial applications of the EMAD technique. It was found that the AT marginally outperformed the ESNs in terms of AUC (0.9369 for the ESN compared to 0.9578 for the AT on dataset A, and 0.9367 for the

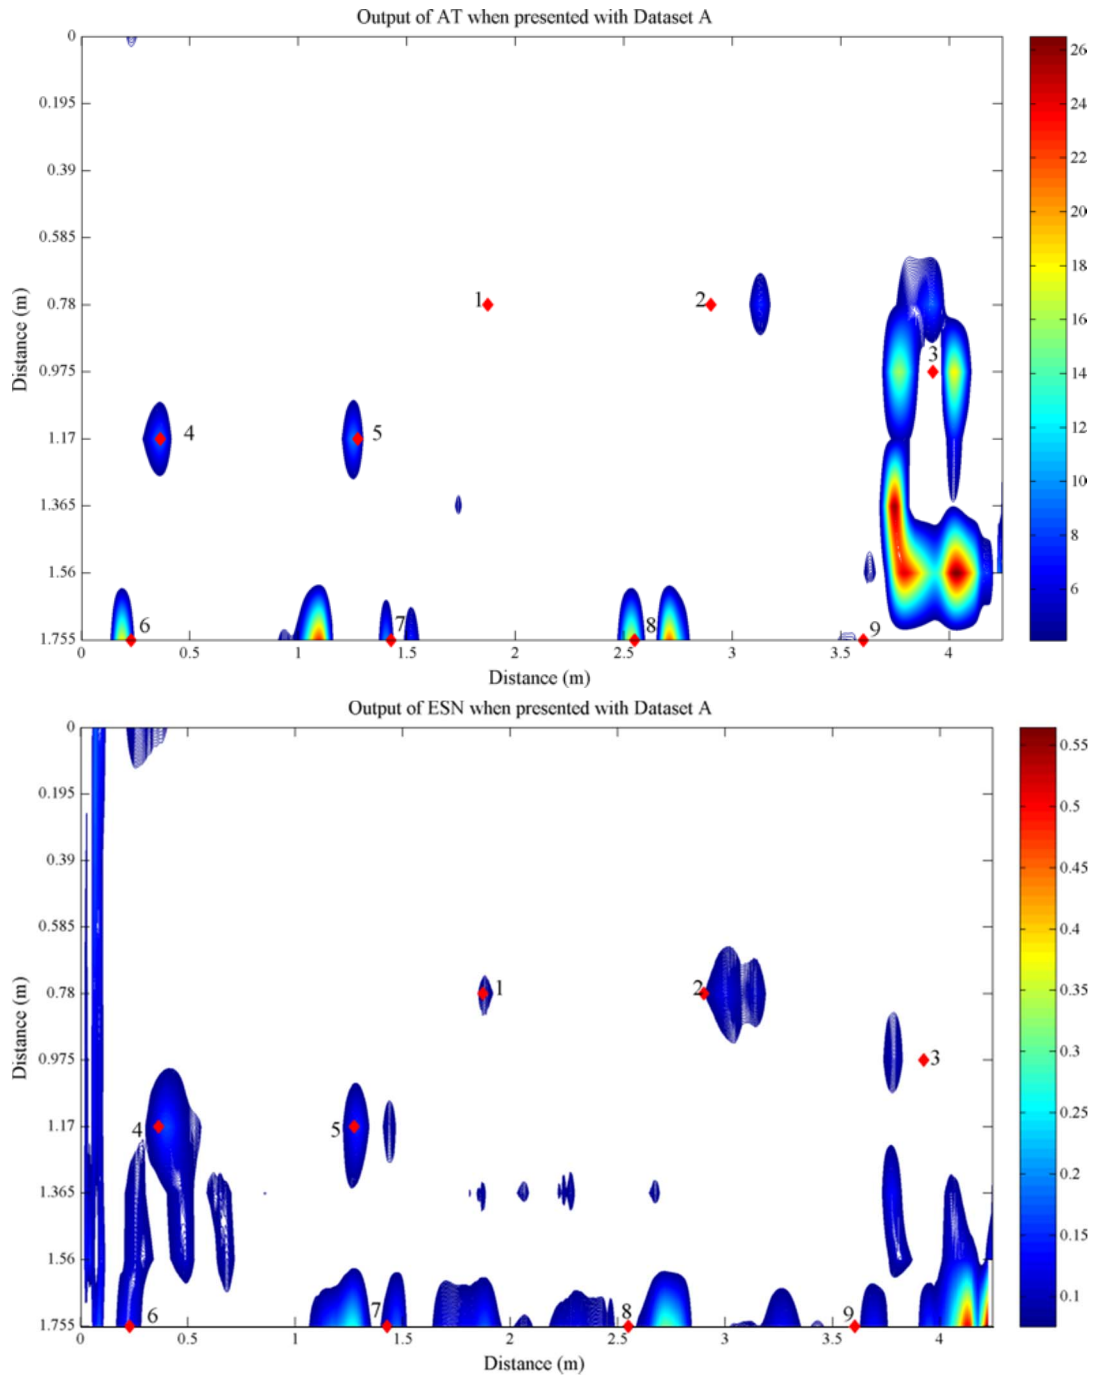


Figure 5: Contour plots of the output of the AT (top) and best performing ESN (bottom) when applied to dataset A. Coloured marks indicate areas on the mesh classified as a defect, while the red diamonds mark the approximate location of known defects (c.f. Fig. 4). The difference in the scale on the colourbar is due to the normalisation of the ESN inputs.

ESN compared to 0.9436 for the AT on dataset B), but that the optimal threshold value was, on average, more consistent for the ESNs than the AT, indicating greater ‘real-world’ applicability. A qualitative evaluation of the performance of the best ESN and the AT showed that while both techniques were successful in locating the known defect signals, the ESN was better able to distinguish between real defect signals and noise, giving a

clearer overall picture of the condition of the mesh and allowing engineers to more accurately locate problem areas. The combination of the EMAD technique and ESNs provides a powerful solution for the accurate spatial localisation of defects using time-series analysis.

Future work now lies in the direction of multi-sensor data fusion, as the fact that this was shown with a relatively simple

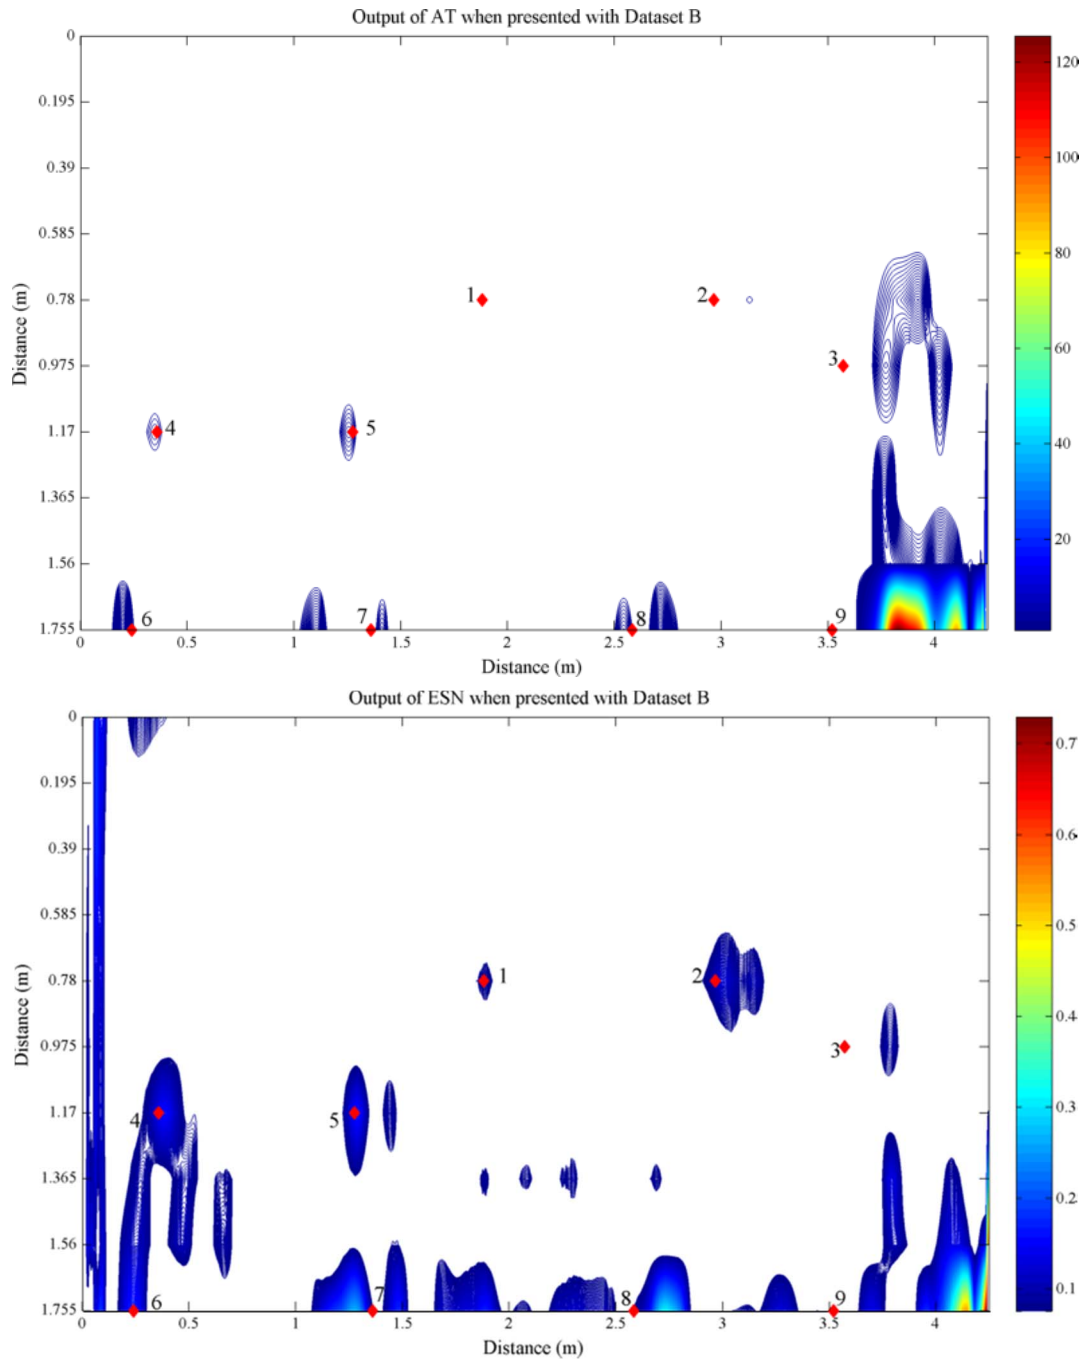


Figure 6: Contour plots of the output of the AT (top) and best performing ESN (bottom) when applied to dataset B. Coloured marks indicate areas on the mesh classified as a defect, while the red diamonds mark the approximate location of known defects (c.f. Fig. 4). The difference in the scale on the colourbar is due to the normalisation of the ESN inputs.

ESN model that had only two input units is interesting, and indicates that a more sophisticated ESN model could potentially outperform the AT. One particularly important result was that the ESN had a more consistent optimal threshold than the AT. Interesting future work could lie in seeing if this holds true when the distance between the EMAD probe and the rebar, representative of greater levels of concrete cover on a road, are

varied. If not, a data fusion approach could be used to integrate cover depth with EMAD data using ESNs. This could potentially account for some of the defects that get lost amongst noisy signals. It would then be interesting to compare the approach presented here with a data fusion-based approach and a number of other neural architectures.

ACKNOWLEDGMENT

The authors are grateful to Dr Jim Austin of Keele University and Emmanuel Mahé from L'École Supérieure de Céramique Industrielle, Limoges, for the development of the EMAD AT.

REFERENCES

- [1] I. Development Research Group, "Failure to Act: The Economic Impact of Current Investment Trends in Surface Transportation Infrastructure," American Society of Civil Engineers, 2013.
- [2] U. S. D. of Transportation Federal Highway Administration, "Deficient Bridges by State and Highway System 2014." U.S. Department of Transportation Federal Highway Administration, 2015.
- [3] A. Ghorbanpoor, "Magnetic-based NDE of steel in prestressed and post-tensioned concrete bridges," in *Proceedings of SPIE 3400, Structural Materials Technology III: An NDT Conference*, 1998, pp. 343–347.
- [4] H.-J. Krause, W. Wolf, W. Glaas, E. Zimmermann, M. I. Faley, G. Sawade, R. Mattheus, G. Neudert, U. Gampe, and J. Krieger, "SQUID array for magnetic inspection of prestressed concrete bridges," *Physica C: Superconductivity*, vol. 368, no. 1–4, pp. 91–95, 2002.
- [5] U. Fiedler, "NDT Data Fusion in Civil Engineering," in *Applications of Ndt Data Fusion*, Kluwer Academic Pub, 2001, p. 193.
- [6] H. Diederich and T. Vogel, "Evaluation of Reinforcing Bars Using the Magnetic Flux Leakage Method," *Journal of Infrastructure Systems*, vol. 23, no. 1, p. B4016001, 2017.
- [7] B. Fernandes, D. Nims, and V. Devabhaktuni, "Comprehensive MMF-MFL inspection for corrosion detection and estimation in embedded prestressing strands," *Journal of Civil Structural Health Monitoring*, vol. 4, no. 1, pp. 43–55, 2014.
- [8] T. C. Jensen, S. E. Wendt, J. N. Gray, C. C. H. Lo, F. J. Margetan, D. J. Eisenmann, N. Nakagawa, and L. H. Brasche, "Feasibility Study for Detection and Quantification of Corrosion in Bridge Barrier Rails," Iowa State University, 2013.
- [9] J. B. Butcher, C. R. Day, P. W. Haycock, D. Verstraeten, and B. Schrauwen, "Defect Detection in Reinforced Concrete using Random Neural Architectures," *Computer-Aided Civil and Infrastructure Engineering*, vol. 29, no. 3, pp. 191–207, 2014.
- [10] A. J. Wootton, C. R. Day, and P. W. Haycock, "Echo State Network Applications in Structural Health Monitoring," in *Proceedings of the 53rd Annual Conference of The British Institute of Non-Destructive Testing (NDT 2014)*, 2014, pp. 289–300.
- [11] H. Jaeger, "The 'echo state' approach to analysing and training recurrent neural networks," Fraunhofer Institute for Autonomous Intelligent Systems, 2001.
- [12] X. Dutoit, B. Schrauwen, J. V. Campenhout, D. Stroobandt, H. V. Brussel, and M. Nuttin, "Pruning and regularization in reservoir computing," *Neurocomputing*, vol. 72, pp. 1534–1546, 2009.
- [13] A. E. Hoerl and R. W. Kennard, "Ridge Regression: Biased Estimation for Nonorthogonal Problems," *Technometrics*, vol. 12, no. 1, pp. 55–67, 1970.
- [14] D. C. Montgomery, E. A. Peck, and C. G. Vining, *Introduction to Linear Regression Analysis*. Wiley, 1982.
- [15] D. Verstraeten, B. Schrauwen, M. D'Haene, and D. Stroobandt, "An experimental unification of reservoir computing methods," *Neural Networks*, vol. 20, no. 3, pp. 391–403, 2007.
- [16] J. Park, B. Lee, S. Kang, P. Y. Kim, and H. J. Kim, "Online Learning Control of Hydraulic Excavators Based on Echo-State Networks," *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 1, pp. 249–259, Jan. 2017.
- [17] Z. H. Khan, N. Hussain, and M. I. Tiwana, "Classification of EEG signals for wrist and grip movements using echo state network," *Biomedical Research*, vol. 28, pp. 1095–1102, 2017.
- [18] E. A. Antonelo, E. Camponogara, and B. Foss, "Echo State Networks for data-driven downhole pressure estimation in gas-lift oil wells," *Neural Networks*, vol. 85, pp. 106–117, 2017.
- [19] M. Lukosevicius and H. Jaeger, "Reservoir computing approach to recurrent neural network training," *Computer Science Review*, vol. 3, no. 3, pp. 127–149, Aug. 2009.
- [20] J. B. Butcher, "Reservoir computing with high non-linear separation and long-term memory for time-series data analysis," Keele University, 2012.
- [21] J. B. Butcher, M. Lion, C. R. Day, P. W. Haycock, M. J. Hocking, and S. Bladon, "A low frequency electromagnetic probe for detection of corrosion in steel-reinforced concrete," in *Concrete Solutions*, CRC Press, 2009, pp. 417–424.
- [22] P. W. Haycock, S. D. Brown, N. Tomlinson, and M. J. Hocking, "An Improved Magnetic Technique For The Non-Destructive Detection Of Faults In Embedded Steel Reinforcement," in *MRS Proceedings*, 2000, vol. 630.
- [23] M. Hocking, L. North, A. Wright, and P. Haycock, "Non-destructive detection of corrosion in steel reinforcement," *Concrete Engineering*, vol. 10, pp. 39–41, 2006.
- [24] H. Scheel and B. Hillemeier, "Capacity of the remanent magnetism method to detect fractures of steel in tendons embedded in prestressed concrete," *NDT & E International*, vol. 30, no. 4, pp. 211–216, 1997.
- [25] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006.