Ships and Offshore Structures

ISSN: 1744-5302 (Print) 1754-212X (Online) Journal homepage: https://www.tandfonline.com/loi/tsos20

Deep machine learning for structural health monitoring on ship hulls using acoustic emission method

Petros Karvelis, George Georgoulas, Vassilios Kappatos & Chrysostomos Stylios

SHIPS AND OFFSHORE STRUCTURES
https://doi.org/10.1080/17445302.2020.1735844

Taylor & Francis
Taylor & Francis Group

Check for updates

Deep machine learning for structural health monitoring on ship hulls using acoustic emission method

Petros Karvelis [a], George Georgoulas [b], Vassilios Kappatos [c] and Chrysostomos Stylios [a]

[a]Laboratory of Knowledge and Intelligent Computing, Department of Informatics and Telecommunications, University of Ioannina, Ioannina, Greece; [b]DataWise LLC, Atlanta, GA, USA; [c]Laboratory of Constructural and Infrastructure Research in Maritime and Air Transport, Hellenic Institute of Transport (HIT), Center for Research and Technology Hellas (CERTH), Thessaloniki, Greece

ABSTRACT
Corrosion, fatigue and corrosion-fatigue cracking are the most pervasive types of structural problems experienced by ship structures. These damage modes, can potentially lead to unanticipated out of service time or catastrophic failure. Acoustic Emission is gaining ground as a complementary Structural Health Monitoring (SHM) technique, since it can offer real-time damage detection. Deep learning, on the other hand, has shown great success over the last years for a large number of applications. In this paper, the SHM on ship hulls is treated as a classification problem. Firstly, the AE signals are transformed, using the Discrete Cosine Transform, followed by a dimensionality reduction stage. Afterwards, a Deep Neural Network is employed by the classification module. The proposed approach was validated and the results indicate that our proposed method can be very effective and efficient, selecting the optimum AE sensor positions and providing almost perfect localisation results.

ARTICLE HISTORY
Received 13 September 2019
Accepted 25 February 2020

KEYWORDS
Ship hull; structural health monitoring; deep learning; deep belief networks; acoustic emission; source localisation; Piecewise Aggregate Approximation

1. Introduction

Shipping has been the largest carrier of freight throughout recorded history, serving more than 90% of global trade by carrying huge quantities of cargo, fuel etc., cost-effectively, cleanly and safely IMO (2012). Apart from commodities, more than 400 million sea passengers a year travelling through European ports, and passenger ships and ferry services have a direct impact on the quality of life of citizens in islands and peripheral regions. In the ship industry, structural failure, due to the severe corroding and metal fatiguing environment, is a major cause of the loss of ships, vessels and tankers. A relatively recent study IUMI (2009) on total vessel losses during 1994–2009 showed that the most frequent cause of accidents for vessels greater than 500GT is weather, followed by grounding, fire/explosion, collision/contact etc. The ship hull damage is among the top five causes of vessel sinking, causing huge loss of human life and pollution of the seas and coastal areas. The sinking of 'Erika' in 1999 FMAIO (1999) and 'Prestige' in 2002 ABoS (2002) were caused by the degradation of the ship hull structure and more than 20,000 tons and 35,000 tons of oil respectively were spilled into the sea, polluting the coast of Spain and France, causing fatal consequences for flora and fauna. Each year, over 400 ocean-going ships sink, many as a result of weakened structures due to corrosion and inadequate and poor welding quality LMIU (2020). A ship's life-cycle is estimated to at least 30 years of operational service and maintenance and condition servicing is an essential part of maritime industry.

In order to ensure that the strength of the ship structure is kept safe for operation, regular hull inspections and repairs of paint coatings, excessively corroded plate and fatigue cracks monitoring must be carefully planned and carried out. Dry-dock inspection is mainly done to determine hull plating thickness at key points to extrapolate the extent and rate of corrosion. It is desirable to detect all cracks above a critical size that may propagate, but a complete inspection of an entire hull or just an entire weld is impractical (Figure 1).

Non-Destructive Testing (NDT) methods have been widely used in dry-docking services to evaluate the reliability of ship structures; however, it is not feasible to inspect the entire hull or even all of the major welds for cracks and defects due to time and cost constraints. Thus, the primary purpose of dry-dock inspections is to determine the thickness of the hull plating at strategic points to extrapolate the extent of corrosion. Typically, this is done via an ultrasonic thickness measurement TMFSC (2020). Using current NDT techniques, only 10% of the total welds length is inspected – unrepresentative of its condition ABS (2007). Additionally, these tests are manual, slow and expensive. Besides dry-docking inspections, some NDT techniques are used for underwater inspection of steel welds LA (2020). The working conditions for divers are difficult and hazardous – new tankers are double-hulled and the inside hull is not accessible – a major limitation of manual underwater NDT by divers. It is obviously important to detect, identify and take early corrective action, in case of hull structural failures.

Acoustic Emission (AE) is the phenomenon of radiation of acoustic (elastic) waves in solids that occurs when a material undergoes irreversible changes in its internal structure. Possible causes of the internal structure changes are crack initiation and growth, crack opening and closure, dislocation movement, twinning, and phase transformation in monolithic materials

CONTACT Petros Karvelis pkarvelis@gmail.com

(i) (ii)

Figure 1. (i) The sinking of Maltese Tanker, Erika, in the Bay of Biscay (1999). (ii) The oil tanker Prestige suffered catastrophic mid-ship structural failure (2002). (This figure is available in colour online.)

Miller et al. (2015). Corrosion, fatigue and corrosion-fatigue cracking, which are the most pervasive types of structural problems experienced by ship structures, can be detected using AE method. AE monitoring involves listening to the process of corrosion itself, which causes AE as a result of the fracture and debonding of expansive corrosion structure, localised yielding, or micro-crack formation. The AE method has been developed for monitoring, detection and location of fatigue cracks in a variety of metal structures, including airframes, oil storage, steel bridges, pipelines and pressure vessels Roberts and Talebzadeh (2003), providing several advantages for monitoring and detection of one or multiple fatigue cracks initiation. The different forms of corrosion could be characterised and identified by AE analysis (Rettig and Felsen 1976; Seah et al. 1993; Jomdecha et al. 2007), offering the great advantage of developing real-time continuous monitoring during ship voyage, if reliable and dedicated noise reduction and pattern classification methods are discovered.

The AE method has been widely used for offshore structures such as ships, oil platforms, bridges with marine basements, subsea pipelines etc. Several researches (Parry 1977; Anastasopoulos et al. 2009) present the use of AE as a global, real-time monitoring of the structural integrity in large-scale offshore structures. The first application of AE analysis technology in an undersea environment was carried out by Exxon Nuclear in an offshore platform Parry (1977). An R&D program Thaulow and Berge (1984) has been carried out to establish relationships between corrosion fatigue crack growth in offshore steel qualities and AE activities. Laboratory experiments on small-scale specimens and wide plates have shown that when a certain combination of crack size and crack surface corrosion deposit thickness has been reached, high AE event rates were recorded Thaulow and Berge (1984). In Løvaas (1985), some full-scale trials have been performed in order to detect cracks in offshore structures, using AE signals. In Kappatos and Dermatas (2011), a detailed review in structural health monitoring of offshore structures including ship hull, using AE testing is provided. The ship hull structures can be monitored by AE techniques (Georgoulas et al. 2009; Kappatos et al. 2009; Kappatos and Dermatas 2009; Kappatos and Dermatas 2011; Georgoulas et al. 2016) during the application of an external stress. The external dynamic stress introduced by the sea-waves and cargo movements in the outside and in the inside of shell respectively is an excellent source of AE phenomena.

One of the ways to tackle the AE detection and source location problem, is to treat it as a typical classification problem. The standard way to build a classifier given a signal is often considered as a three step procedure. The first step is to extract useful and robust features from the signal. In this way, the quite high dimensionality of the signals is reduced. Then, a feature selection is employed in order to reduce further the feature space. Finally, the features are fed into a classifier (statistical, neural, rule based, etc.), which has already been trained using a representative set of historic data.

There is a vast number of research works for the AE detection and source location problem. In Kappatos and Dermatas (2009), a Genetic Algorithm (GA) was used to both reduce the dimensionality of the input vector of predefined features, and also set the spread parameter of a Probabilistic Neural Network (PNN) responsible for the prediction of the location. Wavelet features were extracted from the raw AE signals, and then, after dimensionality reduction through Principal Component Analysis (PCA), a simple linear minimum Mahalanobis distance classifier performed the localisation Georgoulas et al. 2009. In Georgoulas et al. (2016), a radial-basis-function neural network was used to localise AE events in ship hulls. The results showed that the location of a single event can be classified efficiently, using a tiny network configuration and a small set of robust features, selected automatically by the K-means algorithm from a superset of 90 signal parameters.

This work expands on our latest findings (Georgoulas et al. 2016) in which the original AE signal is transformed using the Discrete Cosine Transform (DCT) Strang (1999). DCT is a transform often used for compression Keogh et al. (2000) (condensing the relevant information to few non zero coefficients). Then, the dimensionality reduction technique from the field of time series data mining, namely the Piecewise Aggregate Approximation (PAA) Keogh et al. (2000), Yi and Faloutsos (2000) is applied. The resulting representation fed a Deep Belief Network (DBN) Hinton (2010), which belongs to the family of deep learning paradigms.

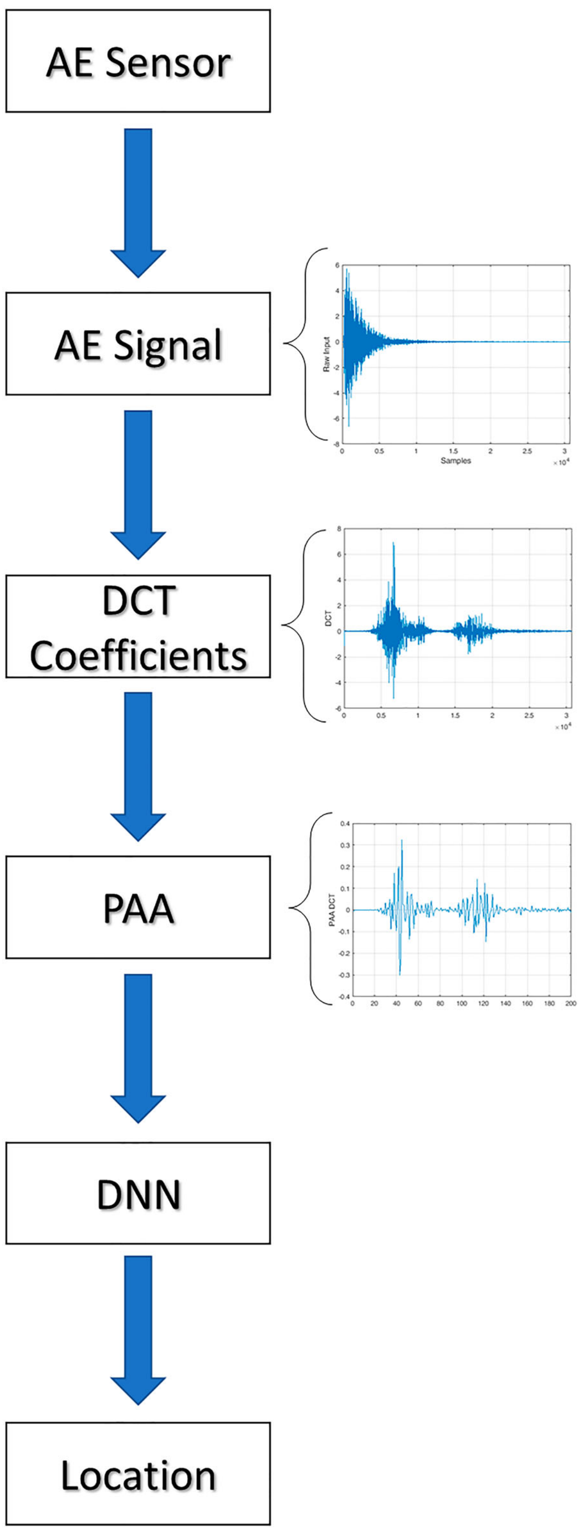

Figure 2. The localisation procedure. AE signal is transformed using the DCT. The produced coefficients undergo a dimensionality reduction stage using PAA. The PAA representation finally feeds the DNN which performs the localisation by assigning one out of three predefined classes. (This figure is available in colour online.)

In the current work, the aforementioned approach is tested using a deep neural network using autoenconders for fast pretraining and different number of input vectors as well as a combination of all sensor measurements. Autoencoders are deep neural networks and were first introduced in the 1980s by Hinton (2010), Rumelhart et al. (1986). The characteristic of these neural networks is the fact that they use the input also as the output of the model reporting great classification results Zhu and Zhang (2019). Our work indicates that high localisation rates can be achieved using various configurations and, more importantly, that the fusion of information can provide almost perfect localisation.

The rest of the paper is structured as follows: Section 2 summarises the proposed procedure with all the involved methods. Section 3 describes the experimental set up, while section 4 presents the achieved results for the various configurations. At the end of paper, section 5 concludes the papers, including some suggestions for future research directions.

2. Proposed procedure

Our proposed procedure relies on a typical classification approach. The raw AE signal is transformed to its DCT coefficients. The number of coefficients is reduced, using the PAA approach and are fed to a Deep Neural Network (DNN) classifier, which predicts one out of three classes as it is described in section 3. Figure 2 depicts a schematic of this approach. The following subsection briefly presents the theoretical background for each one of the aforementioned methods.

2.1. Discrete Cosine Transform

When dealing with time series classification problems, it is common practice to transform the original time series to a more 'informative' space. For example, if different classes have different frequency content, then the Fourier or the wavelet transform are two common choices. In our previous work (Georgoulas et al. 2016), it was shown that DCT can be a very effective transformation for this specific application.

DCT is used to convert the original signal data to a summation of a series of cosine waves of different frequencies. It is quite similar to the Discrete Fourier Transform (DFT) (Ahmed et al. 1974), however DCT involves only the use of *cosine* functions and real coefficients, whereas DFT uses both *cosine* and *sine* functions with complex numbers. DCT is one of the most widely known transforms in signal processing data, and especially in coding for compression Rao and Yip (2014). However, it can also be used for feature extraction, transforming the original time domain signal into DCT coefficients (Rao and Yip 2014).

One of the most common implementations of the DCT is the following Rao and Yip (2014), Proakis and Manolakis (1996).

Suppose a discrete signal x of length L and let s be:

$$s[i] = \begin{cases} x[i], & 0 \le i \le L-1 \\ x[2L-i-1], & L \le i \le 2L-1 \end{cases} \quad (1)$$

Then one can compute the DFT of s as:

$$S[k] = \sum_{n=0}^{2L-1} s[k] \cdot W_{2L}^{nk}, \quad 0 \le k \le 2L-1 \tag{2}$$

and the DCT of x is given by

$$V[k] = W_{2N}^{k/2} V[k], \quad W_L = e^{-j2\pi/L}, \quad 0 \le k \le L-1 \tag{3}$$

The application of the DCT to a raw AE signal is displayed in Figure 3(b). As one can easily notice, the DCT produces L DCT coefficients. In our case, the number of the DCT coefficients was $L = 30,720$ same as the number of the samples of the signal x. In order to compress the signal composing of the DCT components only a fraction of these coefficients is retained, while the other are discarded. In our work, the PAA transform to the DCT coefficients, producing a reduced representation, is applied.

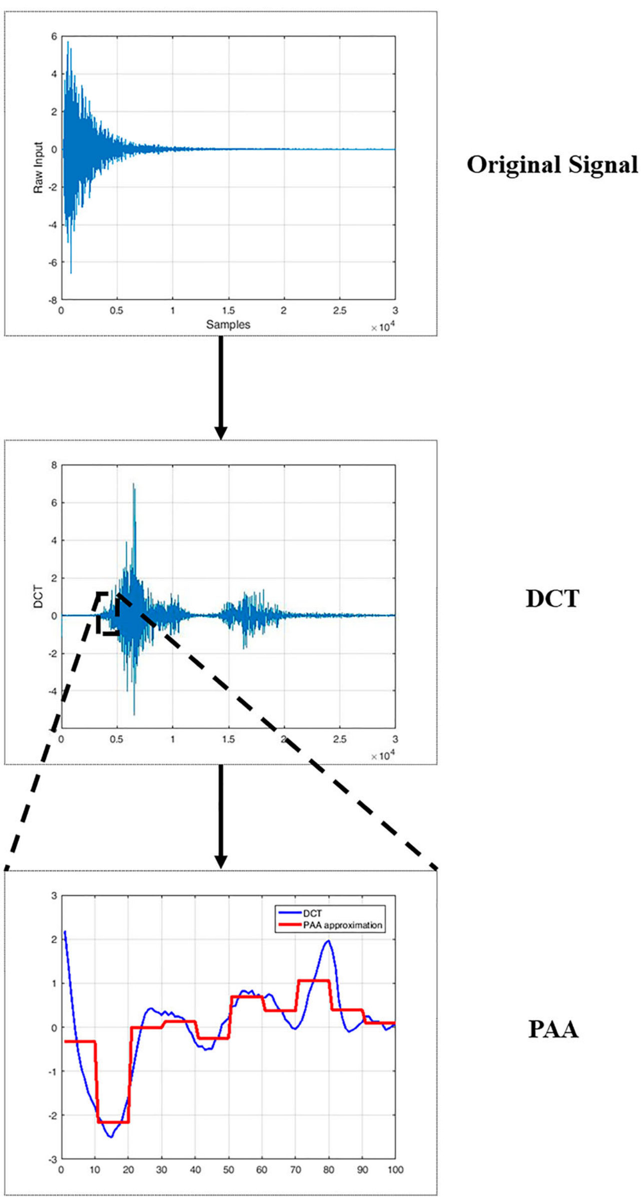

Figure 3. AE signal processing: (a) raw AE signal, (b) the DCT representation of the AE signal and (c) the PAA representation of the DCT coefficients. (This figure is available in colour online.)

2.2. Piecewise Aggregate Approximation

The main role of PAA (Karvelis et al. 2015; SAX 2016), which was introduced by Keogh et al. (2000) and Yi and Faloutsos (2000), is to approximate a discrete signal x of length L into a vector of length w, $x_{PAA} = [\bar{x}[1], \bar{x}[2], \ldots, \bar{x}[w]]$:

$$\bar{x}[i] = \frac{w}{L} \sum_{j=\frac{L}{w}(i-1)+1}^{\frac{L}{w}i} x[j], \quad i = 1, 2, \ldots, w \tag{4}$$

PAA, first divides the signal into w equal sized windows and then computes the mean value for each frame. A slight modified formulation is needed Rao and Yip (2014), in the case where the length of the original signal L is not divided exactly by w. In this case, a signal w times longer than the original signal – a new augmented signal – is created by adding $w-1$ zeros between each consecutive samples of the original signal and then we can easily apply the PAA method Proakis and Manolakis (1996). An illustration of the application of the PAA method to the original DCT coefficients is displayed below in Figure 3.

2.3. Deep Neural Network using pretraining

DNNs have drawn much attention lately, and have been successfully applied in numerous applications and competitions Schmidhuber (2015). On the other hand, very few applications can be found in the available literature in the field of condition monitoring. In Tran et al. (2014) and Tamilselvan and Wang (2013) DBNs were used for the diagnosis of faults in power transformers and reciprocating compressor valves.

In this work, a DNN in used for the classification process. However, training a DNN can be very time consuming and it usually requires a massive amount of data. To alleviate this, an approach was proposed, which involved pre-training the weights of each layer as the weights of an autoencoder Bengio et al. (2007). In other words, the weights of a series of autoencoders can be stuck together to form a DNN.

An autoencoder is composed of two sections representing the encoding half of the net, and the second net that makes up the decoding half. In brief, the autoencoder is a neural network that learns to map the input to an output, which is a copy of the input. Schematically, the architecture of an autoencoder is depicted in Figure 4.

The whole procedure can be better explained by the following Figure 5.

Summing up, the whole procedure consists of:

(1) training of a sequence of shallow autoencoders, greedily one layer at a time, using unsupervised data
(2) training of the last layer of the DNN (created by stacking together the ‘first half’ of the shallow autoencoders of the previous step) using labelled data, and

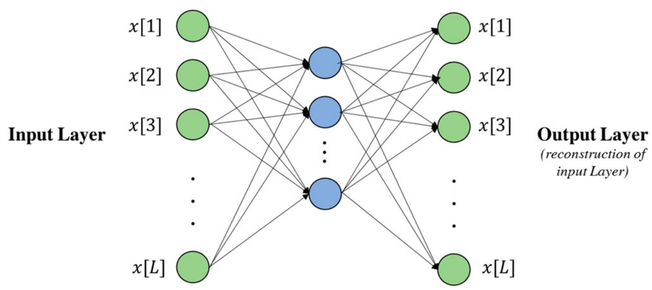

Figure 4. Schematic of a simple (shallow) autoencoder. (This figure is available in colour online.)

(3) use of backpropagation to fine-tune the entire network (end to end) again using labelled data.

This sequential training allows quite deep architectures not only to be trained within reasonable amount of time but can also lead to increased performance.

3. Experimental set-up

A Stiffened Plate Model (SPM) was used to model the side shell of a ship structure. In order to detect the AE signal from all possible positions, four sensors were set in symmetrical positions, taking into account the structure of a ship hull.

The outside side shell is dyed with oil paint in order to simulate as much as possible the outer surface of a real ship's side. Reflections at the end of the SPM were reduced by wrapping the ends in putty. In order to investigate the influence of water, the SPM and its supports were fixed in a water tank. The putty at the edges of the SPM prevents the passage of the water on the inside shell of the plate. The putty is dyed with oil paint for better water-tightness and to avoid putty corrosion. The fixed boundary condition of the model was obtained by clamping the side shell to three heavy bases. Insulation material was placed between the SPM and the three bases, between the three bases and the bottom of the tank, at the bottom of the tank and on the floor to eliminate external noise.

Physical Acoustics Corporation (PAC) R15-Alpha sensors were used to detect waves in the steel stiffened plate. The sensors were stuck on the plate with grease couplant. In real applications, where installation cost is a significant factor, the number of sensors must be reduced and, simultaneously, the sensors must be configured at maximum sensitivity to cover the greatest possible ship hull area. In our experiments, the pulser amplitude was decreased, and the gain was set within the specification limits of the amplifiers used. This setting was close to a real-life application, where the received signal is propagated through several adjacent stiffened plates and the signal reflection at the plate edges is minimised. The sensor signals were amplified at 58 dB (two sensors), 70.7 and 80 dB, digitised at 1 MHz, 16 bit accuracy and stored in the AE analysis system PAC Mistras 2001. The fourth channel's AE system triggers, when the signal on any channel exceeded a pre-defined threshold, and the four-channel simultaneous recordings, were stored on a hard disk. The AE source was simulated with a piezoelectric pulse-generator and pencil-lead breakage (PLB), which are in common use Nielsen (1980). PLB is a long-established standard as a reproducible artificial AE source simulated crack initiation and growth due to corrosion, fatigue and corrosion-fatigue cracking. Often this type of source is also referred to as the Hsu-Nielsen source, based on the original works of Nielsen (1980). Using a mechanical pencil, the lead is pressed firmly against the structure under investigation until the lead breaks. During pressure application with the lead, the surface of the structure gets deformed. At the moment of lead breakage, the accumulated stress is suddenly released, which causes a microscopic displacement of the surface and causes an acoustic wave that propagates into the structure. Since this type of source is easy to handle in laboratory environments, as well as in field testing, it became the most common type of test source in AE testing.

AE source was simulated with a piezoelectric pulse-generator and PLBs and in each location ten repeated measurements

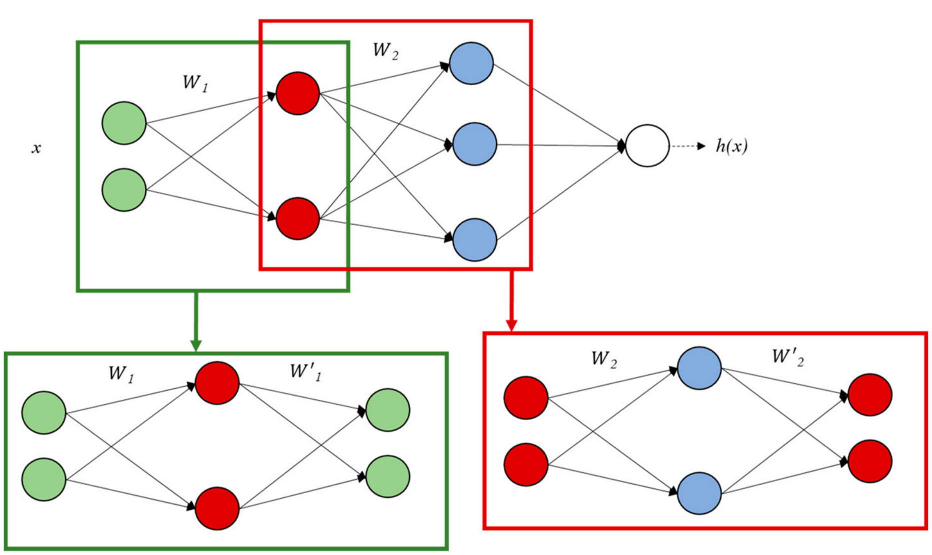

Figure 5. Schematic explanation of the pretraining process. The weights of the input layer of the autoenconder after unsupervised training are stuck to form the neural network. (This figure is available in colour online.)

were recorded. The test signals obtained from PLBs are very reproducible given the handling of the mechanical pencil is repeated accurately. The studies by Strathaus and Bea (1992) and Sucharski (1995) show that the main cause of ship structural damage is fatigue due to wave-induced loads, especially for structures having high stress concentrations at the connection between the longitudinal and the heavy transverse members of the side shell. Three locations/classes were simulated: (i) AE-source in the welding seam between the longitudinal and the heavy transverse member (web), (ii) AE-source in the welding seam between the longitudinal and the side shell, (iii) AE-source in the welding seam between the heavy transverse member (web) and side shell. The collected signals, each one 30,720 samples long, from four different AE sensors were used for the evaluation of the proposed method. A detailed description of the experimental set up and procedure can be found in Kappatos and Dermatas (2009).

4. Dataset description

Ten AE signals were acquired from 15 different locations for each location/class Kappatos and Dermatas (2009) as close as possible to the welding seam, resulting in a total of $3 \times 15 \times 10 = 450$ signals for each one of the four sensors. These recordings were initially tested individually for each sensor and then were also tested all together in the form of a data fusion approach. In each scenario, the 10 fold cross-validation procedure was employed Witten and Frank (2005). The cross-validation procedure is a standard method to assess the generalisation ability of the trained model. More specifically the whole dataset (450) was divided into ten sets of equally distributed classes. The model was trained using the nine sets and the last set was used to access the classification accuracy measure. The above steps were repeated until all sets were tested.

As in Georgoulas et al. (2009), a feature vector of dimension 200 was created after the consecutive application of DCT and PAA. In this work, however, the classification performance of the proposed scheme was evaluated using (all) the 200 features, the first 100 features and the first 50 features (an inspection of the values for the Area Under the Receiver Operating Characteristic Curve (AUC) Wasikowski and Chen (2010), revealed a decreasing trend of the utility of the features).

The second set of experiments features come from all four sensors were put together. To keep the settings as comparable as possible, the first 50 and the first 25 features of each sensor (resulting in $4 \times 50 = 200$ and $4 \times 25 = 100$ features in total), were tested.

As in Georgoulas et al. (2016), using the rule of thumb that smaller layers should be used as we move from the input to the output, to force the neural network to generalise rather than overfit Heaton (2013) three different architectures were employed depending on the size of the input vector: 200-50-10-3, 100-50-10-3 and 50-50-10-3 for the case of individual sensors and 200-50-10-3, 100-50-10-3 for the case of the fusing of all four sensors.

Table 1. Classification performance.

	Features		
	200	100	50
Sensor 1	93.33	94.89	92.44
Sensor 2	64.44	64.89	66.00
Sensor 3	94.67	97.33	96.89
Sensor 4	95.56	98.00	**99.56**
All the sensors combined	**99.78**	98.66	–

5. Results

The results are summarised in Table 1(bold values indicate maximum performance) while all the aggregated confusion matrices can be found in the Appendix.

As in our previous work Georgoulas et al. (2016), sensor no 2 yields very poor results due to the confusion of two out of three classes. This is probably due to its almost equidistance placement from the corresponding welding seams. Regarding the individual sensors, sensor 4 achieves very high results indicating that with the strategic placement of a sensor almost all classes can be successfully covered. The data fusion approach (using measurements from all four sensors) provided the best results. It is also worth noting that in this case, the less informative measurements of sensor 2, were adequately compensated by the information contained in the rest of the sensors.

6. Conclusions and future work

This paper presented a novel data driven method for AE localisation on the extreme complex ship hull structures, combining ideas from the fields of signal processing, time series data mining and deep learning. Our work used a small scale ship hull structure and due to the huge size of typical vessels, a large sensor network will be required to cover the entire or the most critical areas of ship hull structure. The overall network design and development has to ensure reliability and high accuracy detection, location and identification of AE events.

Provided though that a representative data set is available, the proposed tested method seems a viable candidate taking into consideration the generality of the individual methods involved.

An intuitive explanation of its effectiveness as well as the need for their complementary properties can be provided by projecting the multidimensional data (200 features) in a three-dimensional space. The 'projection' was created using a variation of the Stochastic Neighbour Embedding (SNE) method Hinton and Roweis (2002) called t-Distributed Stochastic Neighbour Embedding, t-SNE Maaten and Hinton (2008). Actually, t-SNE is a non linear dimensionality reduction technique which acts in a two-stage process: First, it starts by converting the high-dimensional Euclidean distances between data points into conditional probabilities that represent similarities in the high space. At the second stage, a probability distribution over the points in the low-dimensional map is defined and the method tries to minimise distance between the two distributions.

The implementation of the projection of the data into the 3D space is provided by Figure 6. As it can be seen, the extracted features, using the DCT and the PAA result in compact clusters, belongs to the same class. However, the cluster belonging to a particular class are not forming a global, uniform structure. Therefore, a powerful machine learning algorithm for the classification task is needed.

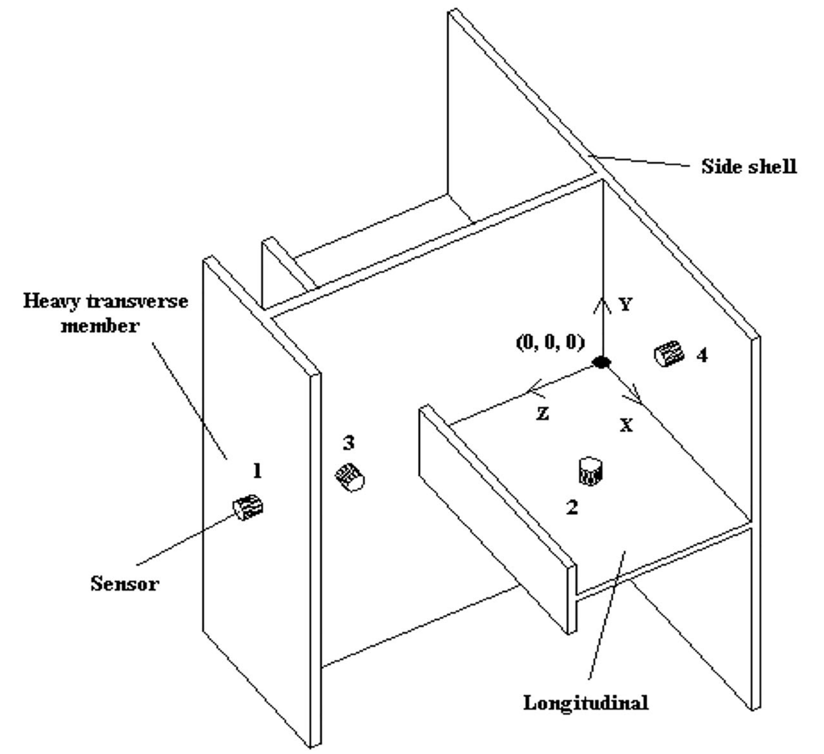

Sensors position (X, Y, Z)		
1. (0, 0, 37)	Side shell	56 x 56 x 1 (cm)
2. (10, 0, 14)	Longitudinal	a. 56 x 1 x15 (cm)
3. (0, 0, 26)		b. 56 x 10 x 1 (cm)
4. (14, 14, 0)	Heavy transverse member	web 1 x 56 x 37 (cm)
		face 20 x 56 x 1 (cm)

Figure 6. The stiffened plate under test and the sensor positions.

On the other hand, it should be noted that this is a data-driven approach. This means that an available set of historic data is needed in order to develop the classification scheme. Moreover, even though the results are promising, it must be kept in mind that the data come from an experimental setting. Therefore, further investigation is needed before the proposed method can actually be used in industrial applications.

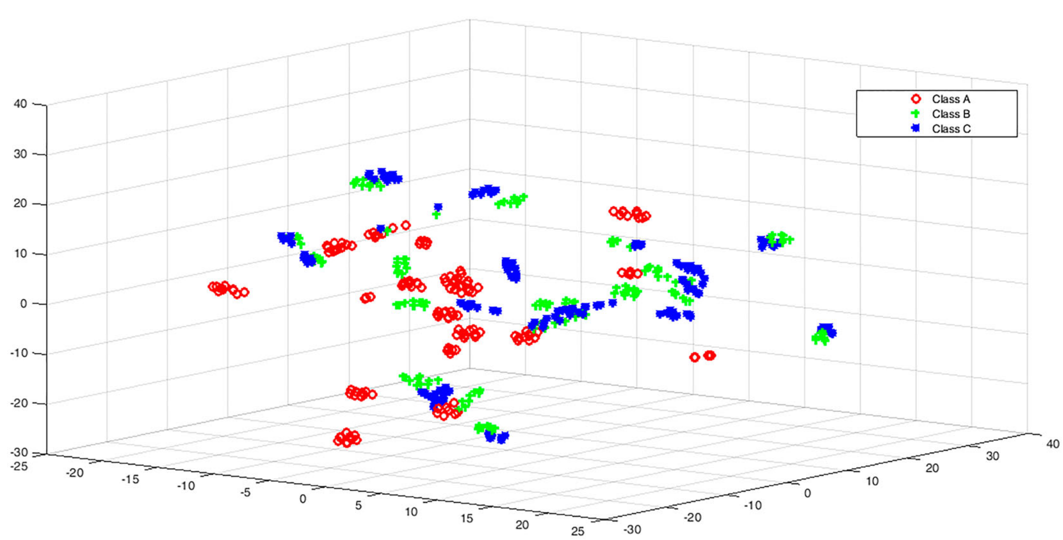

Figure 7. Projection of the original data set using t-SNE. (This figure is available in colour online.)

Starting from this last point, in future work, an attempt will be made to test the method using real-life data. Furthermore, the applicability of the method as a general tool for AE signal analysis will be investigated. Towards that end, other combinations of DCT with methods from the field of time series data mining will be examined in combination with DNN as well as other paradigms from the deep learning family (Figure 7).

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Petros Karvelis http://orcid.org/0000-0002-0483-4868
George Georgoulas http://orcid.org/0000-0001-9701-4203
Vassilios Kappatos http://orcid.org/0000-0002-5314-2622
Chrysostomos Stylios http://orcid.org/0000-0002-2888-6515

References

[ABoS] American Bureau of Shipping. 2002. Technical analysis related to the prestige casualty on 13 November.

[ABS] American Bureau of Shipping. 2007. Hull survey for new construction.

Ahmed N, Natarajan T, Rao KR. 1974. Discrete cosine transform. IEEE Trans Comput. 100(1):90–93.

Anastasopoulos A, Kourousis D, Botten S, Wang G. 2009. Acoustic emission monitoring for detecting structural defects in vessels and offshore structures. J Sh Offshore Struct. 4(4):363–372.

Bengio Y, Lamblin P, Popovici D, Larochelle H. 2007. Greedy layer-wise training of deep networks. In: Advances in neural information processing systems. Vol. 19. Vancouver, Canada: Neural Information Processing Systems Foundation, Inc. (NIPS); p. 153–160.

[FMAIO] French Marine Accident Investigation Office. 1999. Report of the enquiry into the sinking of the Erika off the coasts of Brittany on 12 December 1999.

Georgoulas G, Kappatos V, Nikolakopoulos G. 2016. Acoustic emission localization on ship hull structures using a deep learning approach. 23rd International Conference of Vibroengineering on Modeling, Identification and Fault Detection in Oil and Gas Equipment and Infrastructures, Istanbul, Turkey.

Georgoulas G, Stylios C, Kappatos V, Dermatas E. 2009. Wavelet usage for feature extraction for crack localization. 17th Mediterranean Conference on Control and Automation (MED'09). p. 1540–1545, Thessaloniki, Greece.

Heaton J. 2013. Artificial intelligence for humans. Deep learning and neural networks. Create space independent publishing platform. Chesterfield, Missouri: Heaton Research.

Hinton GE. 2010. A practical guide to training restricted Boltzmann machines. Momentum. 9(1):599–619.

Hinton GE, Roweis ST. 2002. Stochastic neighbor embedding. In: Becker S, Thrun S, Obermayer K, editors. Advances in neural information processing systems. Vol. 15. Vancouver, Canada: MIT Press; p. 833–840.

IMO. 2012. International shipping facts and figures. Information resources on trade, safety, security, environment. The International Maritime Organization (IMO); [accessed 2012 March 6]. www.imo.org.

[IUMI] International Union of Marine Insurance. 2009. Hull casualty statistics.

Jomdecha C, Prateepasen A, Kaewtrakulpong P. 2007. Study on source location using an acoustic emission system for various corrosion types. NDT&E Int. 40:584–593.

Kappatos V, Dermatas E. 2009. Neural localization of acoustic emission sources in ship hulls. J Mar Sci Technol. 14:248–255.

Kappatos V, Dermatas E. 2011. Acoustic emission testing for the monitoring and detection of damage to ship hull structures. Theory and uses of acoustic emissions. Istanbul: Nova Publishers.

Kappatos V, Georgoulas G, Stylios C, Dermatas E. 2009. Evolutionary dimensionality reduction for crack localization in ship structures using a hybrid computational intelligent approach. International Joint Conference on Neural Networks. Atlanta, USA, p. 1531–1538.

Karvelis P, Georgoulas G, Tsoumas IP, Antonino-Daviu JA, Climente-Alarcón V, Stylios C. 2015. A symbolic representation approach for the diagnosis of broken rotor bars in induction motors. IEEE Trans Ind Inform. 11(5):1028–1037.

Keogh E, Chakrabarti K, Pazzani M, Mehrotra S. 2000. Dimensionality reduction for fast similarity search in large time series databases. J Knowled Inform Syst. 3(3):263–286.

LA. 2020. Looking ahead: strong growth for Asia Pacific ports. http://www.porttechnology.org/technical_papers/looking_ahead_strong_growth_for_asia_pacific_ports/.

[LMIU] Lloyd's Marine Intelligence Unit. 2020. http://www.lloydslistintelligence.com/.

Løvaas S. 1985. Acoustic emission of offshore structures, attenuation – noise – crack monitoring. J Acoust Emiss. 4:161–164.

Maaten LVD, Hinton G. 2008. Visualizing data using t-SNE. J Mach Learn Res. 9:2579–2605.

Miller KR, Hill EVK, Patrick O, Moore OP. 2015. Nondestructive testing handbook, third edition. Acoustic emission testing, Columbus.

Nielsen A. 1980. Acoustic emission source based on pencil lead breaking. The Danish welding Institute Publication 80. Copenhagen.

Parry D. 1977. Nondestructive examination of subsea structures using acoustic emission technology. Offshore Technology Conference, Houston (TX).

Proakis JG, Manolakis DG. 1996. Digital signal processing: principles, algorithms, and applications. New Jersey: Prentice Hall.

Rao KR, Yip P. 2014. Discrete cosine transform: algorithms. Advantages, applications. Boston: Academic Press.

Rettig TW, Felsen MJ. 1976. Acoustic emission method for monitoring corrosion reactions. Corrosion. 32(4):121–126.

Roberts TM, Talebzadeh M. 2003. Acoustic emission monitoring of fatigue crack propagation. J Constr Steel Res. 59:695–712.

Rumelhart DE, Hinton GE, Williams RJ. 1986. Learning internal representations by error propagation. In: Rumelhart DE, McClelland, JL, editors. Parallel distributed processing. Vol 1: foundations. Cambridge: MIT Press; p. 318-362.

SAX. 2016. http://www.cs.ucr.edu/~eamonn/SAX.htm.

Schmidhuber J. 2015. Deep learning in neural networks: an overview. Neural Netw. 61:85–117.

Seah KHW, Lim KB, Chew CH, Teoh SH. 1993. The correlation of acoustic emission with rate of corrosion. Corros Sci. 34(10):1707–1713.

Strang G. 1999. The discrete cosine transform. SIAM Rev. 41(1):135–147.

Strathaus R, Bea R. 1992. Fatigue database development and analysis. Technical Report SMP 1-1. Department of Naval Architecture and Offshore Engineering, University of California at Berkeley.

Sucharski D. 1995. Crude oil tanker hull structure fracturing: an operator's perspective. Proceedings of the Prevention of Fracture in Ship Structure, Washington, DC. p. 87–124.

Tamilselvan P, Wang P. 2013. Failure diagnosis using deep belief learning based health state classification. Reliab Eng Syst Saf. 15:124–135.

Thaulow C, Berge T. 1984. Acoustic emission monitoring of corrosion fatigue crack growth in offshore steel. NDT Int. 17(3):147–153.

TMFSC. 2020. Thickness measurement for ship compliance. https://www.lr.org/en/thickness-measurement-for-ship-compliance/.

Tran VT, Althobiani F, Ball A. 2014. An approach to fault diagnosis of reciprocating compressor valves using Teager-Kaiser energy operator and deep belief networks. Expert Syst Appl. 41(9):4113–4122.

Wasikowski M, Chen X. 2010. Combating the small sample class imbalance problem using feature selection. Knowledge and data engineering. IEEE Trans Knowl Data Eng. 22(10):1388–1400.

Witten IH, Frank E. 2005. Data mining: practical machine learning tools and techniques. San Francisco: Morgan Kaufmann.

Yi K, Faloutsos C. 2000. Fast time sequence indexing for arbitrary Lp norms. Proceedings of the 26th International Conference on Very Large Databases, Cairo.

Zhu Q, Zhang R. 2019. A classification supervised auto-encoder based on predefined evenly-distributed class centroids. ArXiv, abs/1902.00220.

Appendix

50 Features

Table A1. Aggregate confusion matrix for sensor 1, accuracy = 92.44%.

		Predicted class		
		A	B	C
TRUE class	A	138	5	7
	B	2	144	4
	C	10	6	134

Table A2. Aggregate confusion matrix for sensor 2, accuracy = 66.00%.

		Predicted class		
		A	B	C
TRUE class	A	148	1	1
	B	1	76	73
	C	1	76	73

Table A3. Aggregate confusion matrix for sensor 3, accuracy = 96.89%.

		Predicted class		
		A	B	C
TRUE class	A	142	6	2
	B	2	146	2
	C	2	0	148

Table A4. Aggregate confusion matrix for sensor 4, accuracy = 99.56%.

		Predicted class		
		A	B	C
TRUE class	A	150	0	0
	B	1	149	0
	C	1	0	149

100 Features

Table A5. Aggregate confusion matrix for sensor 1, accuracy = 94.89%.

		Predicted class		
		A	B	C
TRUE class	A	144	1	5
	B	2	144	4
	C	3	8	139

Table A6. Aggregate confusion matrix for sensor 2, accuracy = 64.89%.

		Predicted class		
		A	B	C
TRUE class	A	148	2	0
	B	6	83	61
	C	6	83	61

Table A7. Aggregate confusion matrix for sensor 3, accuracy = 97.33%.

		Predicted class		
		A	B	C
TRUE class	A	144	4	2
	B	0	148	2
	C	1	3	146

Table A8. Aggregate confusion matrix for sensor 4, accuracy = 98.00%.

		Predicted class		
		A	B	C
TRUE class	A	149	1	0
	B	4	146	0
	C	4	0	146

200 Features

Table A9. Aggregate confusion matrix for sensor 1, accuracy = 93.33%.

		Predicted class		
		A	B	C
TRUE class	A	142	3	5
	B	2	143	5
	C	10	5	135

Table A10. Aggregate confusion matrix for sensor 2, accuracy = 64.44%.

		Predicted class		
		A	B	C
TRUE class	A	145	3	2
	B	5	85	60
	C	5	85	60

Table A11. Aggregate confusion matrix for sensor 3, accuracy = 94.67%.

		Predicted class		
		A	B	C
TRUE class	A	136	11	3
	B	1	149	0
	C	3	6	141

Table A12. Aggregate confusion matrix for sensor 4, accuracy = 95.56%.

		Predicted class		
		A	B	C
TRUE class	A	147	3	0
	B	4	141	5
	C	3	5	142

200 Features, All the Sensors combined

Table 14. Aggregate confusion matrix, accuracy = 99.78%.

		Predicted class		
		A	B	C
TRUE class	A	149	1	0
	B	0	150	0
	C	0	0	150

100 Features, All the Sensors combined

Table A14. Aggregate confusion matrix, accuracy = 98.67%.

		Predicted class		
		A	B	C
TRUE class	A	150	0	0
	B	0	149	1
	C	0	5	145