Pointwise Damage Mapping using Guided Wave Signals for Structural Health Monitoring

KARTHIK GOPALAKRISHNAN and V. JOHN MATHEWS

ABSTRACT

This paper presents a point-wise damage mapping algorithm for composite aerospace structures wherein we classify each location of interest on a structure as belonging to a damaged part of the structure or otherwise. The system employs an array of transducers attached to the structure for inducing guided waves in the structure and recording resultant waveforms arriving at the sensors. We present a baseline signal-based supervised learning algorithm in which baseline measurements on the structure are made prior to deploying it in the field and signals recorded during inspection are compared with their baseline counterparts. For each location of interest on the structure, difference metrics and other characteristics for M closest actuator-sensor paths to the location (M is typically much smaller than the total number of actuator-sensor paths on the structure.) are used as input features to a machine learning classifier. The classifier outputs for all locations of interest are further processed through a series of morphological filters to produce the damage map for the structure. The algorithm was experimentally validated on data recorded from a unidirectional composite panel impacted at five locations. Experimental results show that, for sufficiently high sensor densities, the algorithm presented in this paper is capable of detecting and characterizing multiple damages accurately and with high resolution and Sørensen-Dice index greater than 0.9 for all test cases, and is comparable to results obtained from traditional C-scan ultrasound non-destructive inspection systems. The point-wise damage mapping algorithm offers an attractive high-resolution alternative to traditional grid-based tomographic reconstruction approaches.
*

INTRODUCTION

Guided Wave-based Structural Health Monitoring (GW-SHM) has been widely utilized to inspect and monitor the integrity of structures. A good review of GW-SHM techniques and their applications can be found in [1]. In a standard GW-SHM setup, the

Karthik Gopalakrishnan, PhD Student, Email: gopalakk@oregonstate.edu, Information Processing Group (IPG), School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
V J*ohn Mathews, Professor, Email: mathews@oregonstate.edu, Information Processing Group (IPG), School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.

structure of interest is excited using a pre-defined wave using actuators located on the structure, and the resultant guided waves are recorded at different locations on the structure. The acquired signals are then analyzed to extract damage-related signatures, which are then used to develop algorithms to detect, characterize and localize damage, predict the remaining useful life of the structure, and recommend maintenance strategies. The aerospace industry has evolved significantly in recent years to include different types of fiber-reinforced composite materials. Composites have high stiffness-to-weight properties which are exploited to fabricate large aircraft parts [2]. Inspection of large and complex structures *in-situ* using traditional non-destructive inspection (NDI) systems is time-consuming and difficult. We consider a Lamb wave-based SHM system, in which an array of transducers mounted on a structure is used to induce vibrations in the structure and record the resulting waveforms arriving at different locations on the structure. The signals arriving at the sensors are further processed to reconstruct a map of the structure, which provides information about the location as well as the extent of anomalies on it.

Several researchers have reported on algorithms for reconstructing damage maps to locate damage in structures [3] [4] [5]. Since image reconstruction is typically a sparse problem (as damaged area $\ll$ area of the structure), iterative Algebraic Reconstruction Techniques (ART) based on the Kaczmarcz algorithm have been used for reconstructing damage maps [6] [7]. Though there have been different implementations of the ART algorithm, they are all grid-based algorithms where the inspected structure is divided into pixels or voxels such that the impact of damage on the data recorded at a sensor is modeled as a function of the damage in the voxels that lie on the path between the actuator and the sensor. More details about the ART can be found in [8]. Although ART is a powerful technique for image reconstruction, it suffers from several drawbacks. Since ART is typically implemented in an iterative fashion, the convergence rate can be slow if we want to achieve high-resolution damage maps. Reconstructing high-resolution damage maps requires the image to be divided into a large grid structure, which increases the memory and computational requirements. Moreover, as we increase the number of pixels in the formulation, the optimization problem becomes ill-posed and performs poorly. In this paper, we present a point-wise damage mapping algorithm that offers a high-resolution alternative to traditional grid-based tomographic reconstruction approaches, wherein we classify each sampled location on a structure as belonging to a damaged part of the structure or otherwise. Since the algorithm presented does not depend directly upon the number of points we sample, it is possible to achieve high-resolution damage maps. To the best of the author's knowledge, our approach is novel. Further, experimental results presented in this paper demonstrate the ability of the method to accurately detect and characterize impact damage in carbon fiber-reinforced composite structures.

The remainder of this paper is organized as follows. In the subsequent section, we present the point-wise damage mapping algorithm along with its individual components. This is followed by a section that describes the experimental setup for this work. The section after summarizes the performance of the point-wise damage mapping algorithm. Finally, the paper is concluded in the last section.

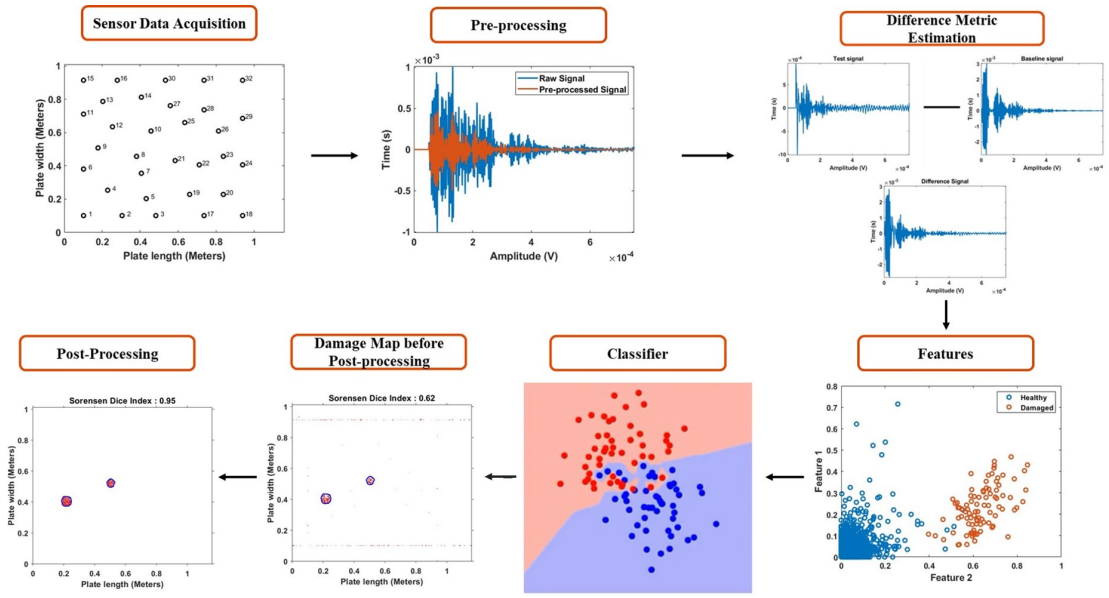

Figure 1. Overall framework of the point-wise damage mapping algorithm

THE POINT-WISE DAMAGE MAPPING ALGORITHM

Figure 1 displays the overall framework of the point-wise damage mapping algorithm. The algorithm is designed to locate damage and to estimate the boundaries of damage present in the structure using a set of transducers that act both as actuators and sensors as needed. The algorithm assumes that the anomaly does not create a through-thickness hole. Lamb waves traveling from an actuator to a sensor may propagate through damaged areas of the structure, resulting in changes in the characteristics of the recorded sensor signal. Characterizing the differences between the baseline and test signals will allow us to interpret the state of the structure. In point-wise damage mapping, for each location of interest, we use damage indices associated with the actuator-sensor paths and the shortest distance of the location to the actuator-sensor paths. For each location, we use only the M closest paths to it, where M is much smaller than the total number of actuator-sensor pairs on the structure. The classifier uses these M sets of features to determine whether the location of interest belongs to a damaged part of the structure or otherwise. By doing this for a large number of points on the structure, a high-resolution map can be obtained. To mitigate the effects of isolated errors in the classifier outputs, we use a post-processing framework inspired by morphological processing. The morphological processing assumes that the damage profile is sparse and clustered, and removes isolated damage locations identified in areas otherwise estimated as damage free as well as isolated healthy locations identified in areas otherwise determined to be damaged.

Our strategy for locating damage, and constructing a representative high-resolution damage map consists of four main steps as presented in Figure 1. The first step is to pre-process the acquired raw signals and prepare them for further analysis. More details about the pre-processing steps are given in the section on ”Experimental Setup”. The next step involves estimating one or more difference metrics between each test signal and the corresponding baseline signal. The difference metric may be computed in

a number of different ways. For example, we could estimate the mean-square value of the difference between the two signals involved. Another option is to compute the correlation coefficient ρ_{tb} between the test and baseline signals.

A feature set that includes the difference metric and other parameters of the structure is established for each location inspected on the structure. The additional parameters are described later in this paper. The third step is to use a machine learning system to predict the probability that each location of interest on the structure belongs to a damaged part of the structure. In the final step, the predicted probabilities are fed to a post-processing framework to mitigate problems due to isolated classification errors. The rest of this section will provide more fundamental insights into the algorithmic components of the system.

Feature Selection

In any damage detection/characterization problem, it is important to extract features from the sensor signals that are sensitive to damage on the structure. Feature extraction also reduces the dimensionality of the data and therefore reduces the computational cost of the SHM algorithm. In this work, for each actuator-sensor pair we used the time average of the squared difference of the baseline and test signals as one of the features. This is computed as

$$D_{ij} = \frac{1}{N}\sum_{k=1}^{N}[\tilde{x}_{t,i,j}(k) - \tilde{x}_{b,i,j}(k)]^2 \tag{1}$$

where $\tilde{x}_{t,i,j}(k)$ is the preprocessed test signal and $\tilde{x}_{b,i,j}(k)$ is the preprocessed baseline signal for the $(i,j)^{th}$ actuator-sensor pair and *N* is the total number of samples in the signals. An advantage of the squared distance metric is that we can expect this metric to be large when the actuator-sensor pair passes through a damaged area, and small when it does not. Other difference metrics may also be used as discussed previously. An example is the correlation coefficient-based metric given by $D_{i,j} = 1 - \rho_{i,j}$, where $\rho_{i,j}$ is the correlation coefficient estimated as

$$\rho_{ij} = \frac{\sum(\tilde{x}_{t,i,j}(k) - \bar{x}_{t,i,j}(k))(\tilde{x}_{b,i,j}(k) - \bar{x}_{b,i,j}(k))}{\sqrt{\sum(\tilde{x}_{t,i,j}(k) - \bar{x}_{t,i,j}(k))^2 \sum(\tilde{x}_{b,i,j}(k) - \bar{x}_{b,i,j}(k))^2}} \tag{2}$$

where $\bar{x}_{t,i,j}(k)$ is the mean of the preprocessed test signal, and $\bar{x}_{b,i,j}(k)$ is mean of the preprocessed baseline signal for the $(i,j)^{th}$ actuator-sensor pair. This difference metric has the additional advantage that it is bounded above and below since $0 \leq D_{ij} \leq 1$. The feature set also includes the distance of each location of interest to its *M* nearest actuator-sensor paths. Let (x_i, y_i) be the location of the $\imath^{th}$ actuator, and let (x_j, y_j) the location to the $\jmath^{th}$ sensor. The distance $l_o(i,j)$ for any location given by (x_0, y_0) to the direct path connecting the actuator and sensor is given by

$$l_o(i,j) = \frac{|x_j - x_i|(y_i - y_0) - (x_i - x_0)|y_j - y_i|}{\sqrt{(x_j - x_i)^2 + (y_j - y_i)^2}} \tag{3}$$

Classification Algorithm

A large number of classifier algorithms exist in the literature [9]. For the analysis in this paper, we use the *k*-Nearest Neighbor (k-NN) algorithm. k-NN is a simple, non-parametric algorithm that does not make any assumptions about the underlying distributions of the data, and is popular in many classification and regression tasks [10]. The four main features of a k-NN algorithm are a) it is a supervised learning algorithm; b) it is a lazy learner ($i.e.$ there is no "learning" during training. The system memorizes the training dataset for improved computational efficiency); c) It is non-parametric; and d) it is based on feature similarity. Since the algorithm is heavily based upon similarity in the feature space, it is important to normalize the features before training and testing to achieve good results. Let *D* be a training set $\{(v_1, c_1), (x_2, c_2), ...(c_n, c_n)\}$, where v_i represents the i^{th} input feature, and c_i its respective class. Let v be the test instance whose class is given by c. The first step is to normalize the features in D and v. In the analysis of this paper, we used the Min-Max normalization to normalize features, and is given by

$$v_i^{\text{norm}} = \frac{v_i - v_{\text{min}}}{v_{\text{max}} - v_{\text{min}}} \tag{4}$$

where $v_{min} = \min(v_i; i \in D)$ and $v_{max} = \max(v_i; i \in D)$ and the min and max operations are performed elementwise. Similarly the division in (4) is also performed elementwise. Next, the distance between the test instance v and each instance in D is computed. The most common choice for this is the Minkowski distance, which is given as

$$d_p(v, z) = \left(\sum_{i=1}^{n} |v_i - z_i|^p \right)^{\frac{1}{p}} \tag{5}$$

where parameter p denotes the order of the Minkowski distance. $p = 1$ gives the Manhattan distance (l_1-norm) and $p = 2$ gives the Euclidean distance (l_2-norm). n represents the total number of input samples. Next, k closest instances to v in D are selected, and their respective class labels are determined. The final step involves computing the probability that the test instance x belongs to class 1, and is computed as

$$P(c = 1|X = v) = \frac{1}{k} \sum I(c_i = 1) \tag{6}$$

where I is the indicator function that represents the count or occurrence of an event. The final predicted class c is determined with the help of a user-defined classification threshold t_{knn} and is given as

$$c = P(c = 1|X = v) \underset{c_0}{\overset{c_1}{\gtrless}} t_{knn}. \tag{7}$$

where c_0 and c_1 represent class 0 and class 1 respectively.

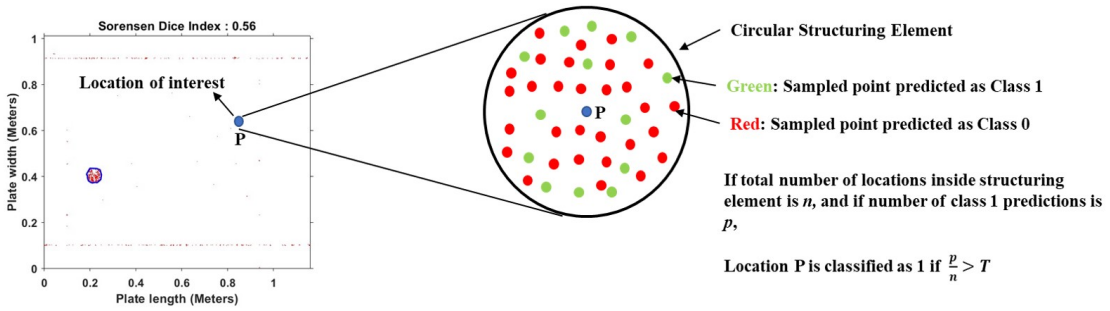

Figure 2. Overview of the post-processing algorithm presented in this paper

Post-Processing

The final part of our framework involves post-processing the output probabilities obtained from the classifier to mitigate problems due to isolated errors at the classifier output. Ideally, the classifier should output high probabilities in and around the damaged area, and relatively low probabilities in other areas of the structure. Therefore, we expect to have high densities of predicted ones around the damaged area, and lower densities of one's otherwise (locations of interest sampled within a damaged area are class 1, while the other areas are labeled class 0). The post-processing framework aims to capture these high-density areas on the structure while simultaneously discounting the lower-density clusters thus enhancing the high-density areas and mitigating the effects of isolated (low-density) classification errors. We accomplish this by designing filters based on ideas from morphological image processing (MIP) [11]. MIP is a collection of non-linear mathematical operations related to the shape or morphology of features in an image. Morphological operations apply a structuring element to an input image, creating an output image of the same size, where the value of each pixel in the output image is based on a comparison of the corresponding pixel in the input image with its neighbors. The number of neighbors it considers is determined by the length or size of the structuring element. The two most widely used operations are *Erosion* and *Dilation*. *Dilation* is an operation that expands the boundaries of regions of interest within an image, while *Erosion* does the opposite by shrinking the boundaries of regions of interest. They can also be thought of as maximum and minimum filters for a region of interest. A combination of these two basic operations forms the basis of many morphological filters. More fundamental theoretical details about MIP can be found in [12].

For our analysis of the data, we designed a post-processing algorithm based on ideas from MIP. Figure 2 illustrates the post-processing algorithm used for a flat panel in the experiments described in the next section. For each sampled location p on the test structure, we considered a circular structuring element of radius r, with p as the center of this circle. Next, the total number of sampled locations on the structure within the circular structuring element is computed, and the predictions made by the classifier for each of these sampled locations are tabulated. If at least T % of the predictions are 1, then the location of interest (*i.e.* center of the structuring element) is classified as 1. This process is applied for each location of interest in the test structure. T is user-defined, and was set empirically to maximize the performance of the algorithm. By repeating this procedure for each sampled location, the damaged area of the structure (high-density

areas) are *dilated*, while isolated errors occurring at the classifier output (low-density areas) are *eroded* thereby enhancing the overall damage map.

EXPERIMENTAL SETUP

The experimental setup and data used were the same as in [4]. We evaluate the performance of the point-wise damage mapping algorithm using experimental data recorded on a unidirectional composite plate 1.16 meters wide, 1 meter long and 2.5 mm thick, made out of eight carbon fiber reinforced polymer (IM7/8552) plies. Thirty-two transducers (Steminc, Steiner Martin Inc., Doral, FL, USA) with approximate bandwidths ranging from 150 kHz to 300 kHz were glued to the upper surface of the plate. The sensors were distributed on an approximately hexagonal grid as shown in Figure 3. All measurements were made by exciting the structure using a 50 µs long linear chirp signal, with linearly increasing instantaneous frequencies from 150 to 300 kHz. This choice was made to match the bandwidth of the transducers and achieve close to the maximum power transfer possible. A tapered cosine window was used to shape the excitation signal with a parameter α_w=0.5 [13]. The parameter α_w is the ratio of cosine-tapered section length to the entire window length. For each measurement, the structure was excited 200 times using the actuator, and the sensor signals were averaged to improve the SNR. Baseline measurements were made on the structure prior to impacting it. The plate was impacted at five different locations using a drop-weight impactor consisting of a 5.6-kg hemispherical impactor and a plastic tube that guided the impactor. A 9.5-mm thick rubber sheet was placed under the plate before impacting to mitigate deformation. Data was recorded each time the structure was impacted. After performing the impact experiments, a C-scan device was used to estimate the extent of damage in the structure to compare with the results of the algorithm of this paper. The damage that resulted from the five experiments was internal to the structure and was not visible on the surface of the plate. Figure 3 also shows the outlines of damage resulting from the five impacts.

Before generating training and testing datasets, the acquired raw signals were preprocessed and prepared for further analysis. The signals were windowed in time such that only the zeroth order symmetric (S0) and anti-symmetric (A0) mode and their first reflections were captured. The S0 and the A0 wave modes are the fundamental guided wave modes that represent the basic or fundamental wave patterns that can propagate through a structure. The window length was determined from the geometry of the structures, the maximum distance between an actuator-sensor pair, and the wave velocities of the two modes as obtained from the dispersion curves for the composite material (IM7/8552). The processed signals were then passed through a bandpass filter to eliminate low-frequency vibrations and any high-frequency electrical noise. The passband of the bandpass filter was 150-300 kHz, which was approximately the same as the bandwidth of the transducers.

To generate training and testing datasets, we randomly sampled damaged and healthy locations. For each location, we extracted features as discussed in the previous section. The baseline data for training was collected prior to impacts, and the inspection data for training was recorded after the first impact. Hence, we evaluated our algorithm by mapping the remaining four impact damages on the composite plate. The baseline and

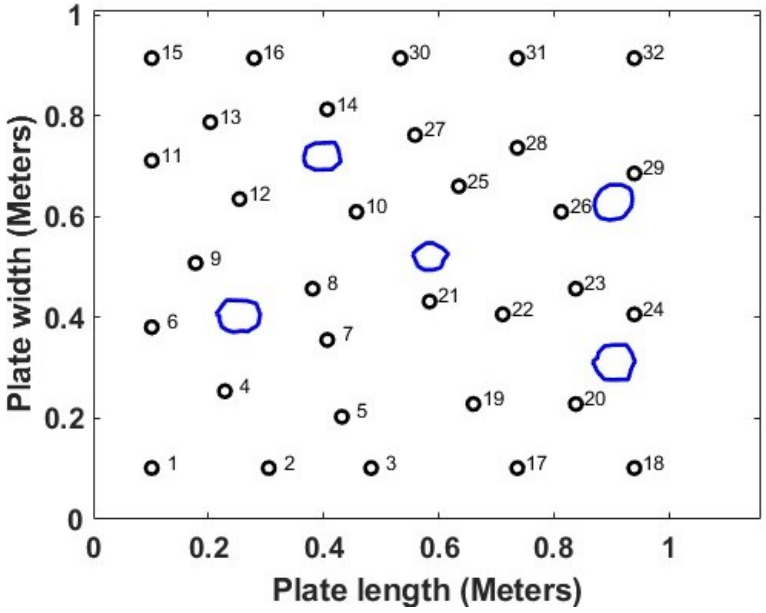

Figure 3. Locations of the transducers and impact damage on the structure. Black circles indicate the sensor locations and damage outlines estimated by C-scan are in blue.

inspection data changed for each text case for which we assessed and mapped the damage from different impacts and their combinations. For example, to assess the damage from the second impact, the baseline signals were recorded from the plate after the first impact, and the inspection data was recorded after the second impact. Similarly, to assess the damage caused by the second and third impacts simultaneously, the baseline signals were recorded from the plate after the first impact, and the inspection data was recorded after the third impact. The machine was trained on signal features and the corresponding labels obtained from 10000 locations sampled randomly on the structure before and after the first impact. The 10000 locations were equally distributed between damaged and undamaged areas. We evaluated our algorithm on 280000 locations sampled for each test case which results in a sampling density of approximately *124 locations/cm*^{2}, which aided in reconstructing high-resolution damage maps.

For each sampled location in the datasets, the damage indices of the 20 closest actuator-sensor pairs, along with the distance of the location to the direct paths connecting the same 20 closest actuator-sensor pairs were used as features to the classifier. In this work, we used a k-NN classifier utilizing the Euclidean distance metric, with the number of neighbors set to $k = 100$. The classification threshold was set to 0.35, and the threshold T for post-processing was set to 0.25. The parameters mentioned above were selected based on a preliminary empirical evaluation. Additional analyses of the effect of individual parameters of the algorithm on the overall performance are planned for the future. The performance of the point-wise damage reconstruction algorithm was quantified using the Sørensen-Dice index which is a metric that compares the similarity between two datasets and is defined for two sets A and B as

$$S(A, B) = 2\frac{|A \cap B|}{|A| + |B|} \tag{8}$$

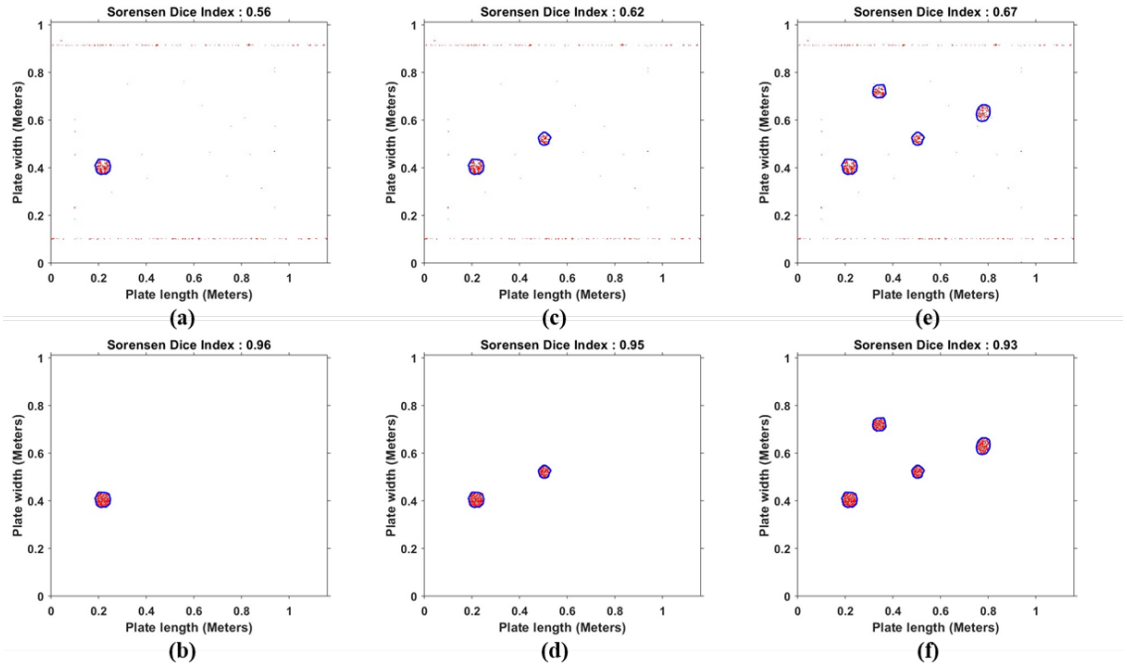

Figure 4. (a-c-e) The reconstructed damage map is shown in red before post-processing for the composite plate with one, two, and four impact damages respectively (b-d-f) The reconstructed damage map in red after post-processing for the same cases. Blue lines indicate the boundaries of the damage estimated by C-scan.

For our problem A may be the set of damage locations on the structure identified using C-scan, and B may be the set that represents the estimated damage locations using the algorithm of this paper. A Sørensen-Dice index value of 1 implies a perfect match between the sets while a value of 0 implies that the sets are disjoint.

RESULTS

Figure 4 shows the reconstructed maps along with their Sørensen-Dice indices for three different test cases. Table I shows the Sørensen-Dice indices for all test cases along with a comparison of the results from [4]. For each test case, the Sørensen-Dice indices were computed for the estimated damage outline obtained by the algorithm of this paper and the impact boundaries obtained using C-scan. It is clear from Figure 4 and Table I that the algorithm presented in this paper is able to produce accurate damage maps with high resolution and Sørensen-Dice index greater than 0.9 for all the test cases. The results further demonstrate the robustness of our algorithm in simultaneously mapping multiple damages with high accuracy as reflected by the improved Sørensen-Dice indices in Figure 4(f). Moreover, the algorithm presented in this paper outperformed the method of [4] for almost all cases. Both [4] and this paper used the same set of data. The differences in the two sets of results were statistically significant with a p-value of 0.0023 (Wilcoxon rank sum test). Figure 4 (a-c) shows the reconstructed damage maps for three test cases before post-processing. The respective Sørensen-Dice indices are displayed

TABLE I. SORENSEN-DICE SIMILARITY INDICES FOR ALL TEST CASES.

Test Case	Algorithm in [4]	Algorithm of this paper
Impact 2	0.92	0.96
Impact 3	0.84	0.94
Impact 4	0.89	0.92
Impact 5	0.94	0.91
Impacts 2,3	0.88	0.95
Impacts 3,4	0.88	0.93
Impacts 4,5	0.76	0.92
Impacts 2,3,4	0.86	0.94
Impacts 3,4,5	0.83	0.92
Impacts 2,3,4,5	0.85	0.93

on the title for each test case. It is seen that the Sørensen-Dice index is relatively lower for the raw classifier outputs prior to post-processing. The errors that contributed to this came from several isolated errors in the classifier output along with a clustered set of incorrect decisions in damaged areas. Furthermore, a sequence of incorrectly predicted damage locations can be observed along the boundaries. This occurrence could potentially be attributed to the placement of some transducers, which aligns approximately along the same boundaries as seen in Figure 3. Post-processing plays a critical role in obtaining highly accurate damage maps by removing most of these errors as discussed above. The current post-processing framework has the ability to detect damage whose dimensions are larger than that of the circular structuring element. If this is not the case, other types of structuring elements may need to be considered.

CONCLUSIONS

This paper presented a new guided-wave-based damage mapping algorithm for active SHM systems. As part of the overall framework, we also designed a post-processing algorithm employing morphological processing. Experimental results showed that for sufficiently high sensor densities, the algorithm presented in this paper was able to detect and characterize multiple impact damages on a unidirectional composite panel accurately, and with comparable results to traditional C-scan-based ultrasound non-destructive inspection systems, while also guaranteeing high-resolution damage maps. Point-wise mapping allows for flexible sensor placement, enabling improved monitoring of critical areas in a structure. Moreover, the algorithm is also relatively simpler to implement compared to many traditional approaches. The current implementation of the point-wise damage mapping algorithm relies on a supervised approach, requiring training on data recorded from damaged structures, which is costly and in many situations, impractical. A more practical approach would involve adopting an unsupervised system that exclusively trains on data from healthy structures. Furthermore, considering the results presented in this paper were obtained for a relatively high sensor density, it would be beneficial to investigate the impact of sensor density on both the algorithm's performance and the resolution of the resulting damage maps. Exploring the relationship

between sensor density and resolution can help determine the best sensor configurations for accurate and cost-effective damage mapping. Overall, the evidence presented in this paper suggests that the point-wise damage mapping algorithm is an attractive high-resolution alternative to traditional guided wave SHM algorithms and has the potential to enhance SHM capabilities for aerospace structures.

REFERENCES

1. Mitra, M. and S. Gopalakrishnan. 2016. "Guided wave based structural health monitoring: A review," *Smart Material Structures*, 25(5):053001, doi:10.1088/0964-1726/25/5/053001.
2. Prasad, S. M., K. Balasubramaniam, and C. V. Krishnamurthy. 2004. "Structural health monitoring of composite structures using Lamb wave tomography," *Smart Materials and Structures*, 13(5):N73, doi:10.1088/0964-1726/13/5/N01.
3. Zhao, X., R. L. Royer, S. E. Owens, and J. L. Rose. 2011. "Ultrasonic Lamb wave tomography in structural health monitoring," *Smart Materials and Structures*, 20(10):105002.
4. Zoubi, A. B., S. Kim, D. O. Adams, and V. J. Mathews. 2019. "Lamb Wave Mode Decomposition Based on Cross-Wigner-Ville Distribution and Its Application to Anomaly Imaging for Structural Health Monitoring," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 66(5):984–997, doi:10.1109/TUFFC.2019.2903006.
5. Mathews, V. 2014. "Damage mapping in structural health monitoring using a multi-grid architecture," in *Proceedings of QNDE 2014*, Boise, Idaho.
6. Kak, A. C. and M. Slaney. 1988. *Principles of Computerized Tomography*, IEEE Press, Piscataway, NJ.
7. Malyarenko, E. V. and M. K. Hinders. 2001. "Comparison of cross-hole and fan beam Lamb wave ultrasonic tomography," *Quantitative Nondestructive Evaluation*, 20:732–739.
8. Gordon, R., R. Bender, and G. T. Herman. 1970. "Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography," *Journal of Theoretical Biology*, 6(2):147–155, doi:10.1016/0021-9991(70)90092-5.
9. Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*, Springer, New York, NY, ISBN 978-0387310732.
10. Fix, E. and J. Hodges Jr. 1951. "Discriminatory analysis, nonparametric discrimination: consistency properties," Project RAND Research Memorandum RM-704, RAND Corporation, Santa Monica, CA.
11. Serra, J. 1982. "Image Analysis and Mathematical Morphology," *Academic Press*, 1(2):87–120.
12. Soille, P. 1999. "Morphological Image Analysis: Principles and Applications," *Springer-Verlag*, doi:10.1007/978-3-662-03882-6.
13. Harris, F. J. 1978. "The Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proceedings of the IEEE*, 66(1):51–83, doi:10.1109/PROC.1978.10837.