Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Strain-based delamination prediction in fatigue loaded CFRP coupon specimens by deep learning and static loading data

28 July 2024

Strain-based delamination prediction in fatigue loaded CFRP coupon specimens by deep learning and static loading data

Demetrio Cristiani[1,3], Francesco Falcetelli[2,3], Nan Yue[3], Claudio Sbarufatti[1], Raffaella Di Sante[2], Dimitrios Zarouchas[3], Marco Giglio[1]

[1] Politecnico di Milano, Dipartimento di Meccanica, Via La Masa 1, 20156, Milano, Italy

[2] University of Bologna, Department of Industrial Engineering, Forlì, Italy

[3] Structural Integrity & Composites Group, Faculty of Aerospace Engineering, Delft University of Technology, Delft 2629 HS, The Netherlands

* Corresponding author. E-mail address: demetriocristiani@gmail.com

Abstract

Machine learning (ML) methods for the structural health monitoring (SHM) of composite structures rely on sufficient domain knowledge as they typically demand to extract damage-sensitive features from raw data before training the ML model. In practice, prior knowledge is not available in most cases. Deep learning (DL) methods, on the other hand, can obtain higher-level features from raw input data and have proven superior in several applications. This paper proposes a Convolutional Neural Network (CNN) based approach for the delamination prediction in CFRP double cantilever beam (DCB) specimens using raw local array strain measurements via distributed optical fiber sensors. The conventional CNN architecture is modified to perform regression, as the delamination size is a continuous value. 1D and 2D CNN architectures are deployed and compared and different techniques are exploited to encode 1D spatial strain pattern series as 2D images. Raw strain patterns collected during static testing are used to train the CNNs, while testing is performed on unseen raw fatigue strain patterns, showing the CNN ability to automatically extract discriminative features from the non-pre-processed static strain pattern-based signals that generalize to raw fatigue signals as well. This strategy has the potential to reduce fatigue testing expenditures while also shortening the time required to gather training data.

Keywords: Convolutional Neural Network, B. delamination, B. fatigue, A. prepreg

1. Introduction

Carbon fibre reinforced plastics (CFRP) are increasingly used in several industries as an alternative to conventional metallic alloys, due to their specific strength and stiffness, making them suitable for a wide range of applications where weight savings are crucial to the overall performance [1]. Although CFRP have numerous advantages over conventional alloys, they also suffer from a few major disadvantages, as their damage initiation and propagation mechanisms are more difficult to predict compared to metallic materials [2]. These failure mechanisms are known to initiate at the level of the constituents (e.g., matrix microcracks, delaminations, fiber breakage, etc.) and can grow to an extent that compromises the structure integrity.

Among failure modes, delamination is possibly the major and most frequent, crucially affecting the strength, stiffness, stability, and usable service life of laminated composites, eventually leading to catastrophic failure of the composite structures [3]. For these reasons, the delamination monitoring of composite laminates is of prime importance and structural health monitoring (SHM) approaches have been developed in order to replace the traditional nondestructive testing (NDT) based maintenance strategies [4,5], fostering the transition to condition-based maintenance philosophies. Although SHM methods have been successfully applied in both metallic and composite structures, the latter present additional challenges, as the material heterogeneity, together with the incomplete knowledge about the principles underlying the multivariate damage evolution and interaction processes, leads to uncertainty in the assessment of current and future material properties.

SHM methods are typically grouped into two categories, i.e., model-based approaches and data-driven approaches [6]. Model-based approaches tend to be more accurate if the damaged system response can be modeled precisely, requiring prior knowledge about the physics of the system, which is often not available in practice. On the other hand, data-driven approaches rely on historical data and AI methods to quantify the damage state of the structure, and do not require prior expertise about the system underlying physics. The abundance of data engendered by the advances in information and sensing technology has promoted the development of data-driven algorithms, and good results have been obtained.

Among AI techniques, machine learning (ML) is the most acknowledged [7]. Conventional ML techniques, however, require considerable domain knowledge as they demand to extract damage-sensitive features from the raw data before training the ML model. Deep learning (DL) methods have overcome this shortcoming, as they allow data to be used in their raw form, automatically learning from data using a general-purpose learning procedure [8]. Recently, DL methods such as convolutional neural networks (CNN) and recurrent neural networks (RNN) have been successfully applied to the SHM of civil structures [9,10], especially to automatically process image data, as structural damage is typically visible. The same concept has been applied to composite materials; for instance, Fotouhi et al. [11] used a comprehensive image-based data set including common microscale damage mechanisms, such as matrix cracking and fiber breakage, and macroscale damage mechanisms such as impact and erosion to train a CNN, achieving a classification accuracy level of 87% - 96% for identifying the damage severity and types. In addition to image data, time-series data are widely used for damage detection, often transforming the raw time-series data into frequency spectra or spatial time-frequency spectra [12–14]; however, 1D CNNs have been recently proposed to operate directly on time series and immediately achieved the state-of-the-art performance levels in several applications [15].

Despite DL methods have proven superior to conventional ML methods in several applications, they have only been used in a few studies with composite materials for SHM applications [16]. For instance, Tabian et al. developed a passive sensing CNN based framework for impact detection [17]; time domain signals corresponding to distinct sensors are gathered together and transformed into 2D images prior being fed to the CNNs, performing impact localization (the structure is divided into subregions) and impact categorization (energy level, 3 energy classes), achieving a 95% accuracy. Similarly, Damm at al. [18] and Jung et al. [19] developed CNN-based impact damage assessment frameworks. Khan et al. [12] proposed a CNN based approach for the classification and prediction of in-plane and through-the-thickness delaminations in smart composite laminates using structural

vibration measurements. Time domain transient signals are transformed into spectrograms via the short time Fourier transform (STFT). The CNN distinguishes between the damaged and undamaged states, and classifies the damage scenario (i.e., the delamination location) with an overall classification accuracy of 90.1%. Yu et al. used an addressable conducting network (ACN) and deep learning-ANN for damage detection in CFRP specimens, achieving a detection accuracy rate of about 95% [20].

In the literature, a vast range of techniques have been proposed for the delamination assessment in laminated composites [21], among which strain-based methods have been proven to be adequate and present some practical advantages over other techniques, also thanks to the recent advances in fiber optic sensing technologies [22–31]. For instance, several studies have shown the possibility of embedding the optical fiber (OF) within the composite laminate [22,24,26,32,33], and the introduction of distributed OF sensors has fostered their deployment in many SHM applications [22,23,34].

To the authors knowledge, published conventional data-driven SHM approaches require some sort of preprocessing of the raw data to extract relevant features for the delamination assessment, implying that expert's knowledge is essential. On the other hand, DL approaches do not require any preprocessing of the raw data; however, DL methods for the delamination assessment, such as CNNs, typically require the delamination to be described by a finite number of states to make its classification possible, rather than allowing the delamination state to be a continuous variable.

One of the requirements for data-driven approaches is the availability of training data; generating a training dataset for the delamination fatigue assessment would imply running several run-to-failure fatigue tests to gather the training data, requiring considerable resources and preventing the use of data-driven approaches in most cases. Therefore, in this paper a DL approach is proposed that can address the above-mentioned challenges and limitations. Strain-based CNN architectures for the delamination prediction are developed that do not require any pre-processing of the input strain patterns. The CNN conventional architecture is modified to perform regression rather than classification, as the delamination size is a continuous variable. Double cantilever beam (DCB) specimens are tested under static and fatigue loading. Distributed OF sensors are placed on the top-surfaces of the specimens to collect local strain patterns during delamination propagation. Raw static test strain field array measurements at fixed times are fed to 1D and 2D CNNs for training, and different techniques are exploited to encode 1D spatial strain pattern series as 2D images (e.g., continuous wavelet transform, Gramian angular field, Markov Transition Field). Finally, the CNNs testing is performed on raw strain field array measurements collected from fatigue tests, showing that the CNNs can extract relevant features from static strain patterns that also generalize to unseen fatigue strain patterns. This approach can be extended to more complex structures, potentially saving fatigue testing related costs and cutting the time needed to generate training data.

The paper is organized as follows: Section 2 (*Experimental details*) provides a description of the experimental set up and testing procedures, including a detailed characterization of the strain patterns. Section 3 (*Delamination sizing by deep learning and distributed strain measurements*) discusses the system-level operation principle of the CNN and passive sensing-based methodology for delamination assessment, especially on how the sensing data can be prepared and used for the planned task, outlining the proposed methodology for the

local assessment of delaminations from strain field response measurements. Section 4 (*Results and discussion*) illustrates the obtained results by comparing the outcomes of the proposed DL approaches. Section 5 (*Conclusions*) outlines the contribution of the current work, including a discussion of potential future work.

2. Experimental details

2.1. Materials and fabrication

The DCB specimens were manufactured by laminating 24 plies of 300x300 mm carbon fiber unidirectional prepreg named HexPly® 8552 (Epoxy matrix and AS4 12K carbon fibres); the stacking sequence is $[0_{24}]$ and the panel was cured inside an autoclave according to recommendation from Hexcel. A 0.012 mm thick Teflon™ film was inserted at the midplane of the laminate for crack initiation so that the initial starting crack length was approximately 50 mm. Based on ASTM D5528 standard [35], strips 25 mm wide were cut from the plate using a water-cooled diamond saw, while piano hinges were bonded to the either side of the specimens; the specimen relevant dimensions are shown in Figure 1. During tests, a 9 Megapixel camera with 50 mm-focal-length lens was placed at the side of the clamped specimen to monitor the delamination length on one of the edges. The edge surfaces of each specimen were covered with thin white paint in order to enhance the white-black contrast of cracked and uncracked regions (Figure 2). One edge of each test article was marked in 1mm intervals to obtain visual edge measurements using a synchronized camera system. The camera was synchronized with the testing machine and with the interrogator (ODiSI-B system from Luna Innovations Inc.). The OF sensor – single mode Ormocer coated low bend loss 125 μm fiber commercialized by FBGS Technologies GmbH (Jena, Germany) with LC/APC connector – was bonded on the top surface of each specimen using a cyanoacrylate adhesive, where three equally spaced fiber passes were used to assess the delamination front (see Figure 1 and Figure 3) offering a spatial resolution of 1.25 mm along the bonded OF.

Figure 1: Experimental procedure setup.

Figure 2: Specimen edge during testing.

Figure 3: Specimens OF sensor placement.

2.2. Static test

Five DCB specimens were clevis-mounted in a Zwick - 20 kN test frame with a 20 kN load cell and loaded at a displacement rate of 1 mm/min. Based on ASTM D5528 standard [35], precracking was performed before running the test. Baseline data were acquired prior to testing to calibrate the interrogator; the sampling frequency was set to 0.5 Hz.

2.3. Fatigue test

Four DCB specimens were clevis-mounted in a MTS - 10 kN Elastomer hydraulic test frame with a 10 kN load cell and load-control fatigue tested. A schematic representation of the applied loading profile, containing the repetitive cyclic loading blocks and the tensile loading-unloading ramps, is shown in Figure 4. Precracking was performed before running the test to accelerate the delamination onset. Constant amplitude of sinusoidal waves, with a maximum load equal to the 80% of the precracking load, load ratio 0.1 and frequency 5 Hz were applied, while the tensile loading and unloading ramps were applied before and after every 500 cycles to trigger the measurements (Figure 4); baseline data were acquired prior to testing to calibrate the interrogator.


Figure 4: Fatigue test loading pattern.

2.4. Strain characterization of the DCB specimen

The strain profile along the length of the DCB specimen can be predicted under the assumptions that the DCB arms act as if they were Euler-Bernoulli beams clamped at the crack tip and the interface material is infinitely stiff and perfectly brittle [36]. These assumptions lead to a triangle-like strain shape where two regions can be defined: i) bending strain region from the arm free-end up to the delamination front and ii) zero-strain region ahead of the delamination front. Region 1 is due to the separation of the DCB arms; the strain increases linearly due to the bending moment reaching its maximum value in correspondence of the crack tip. Beyond the crack tip, the magnitude of the strain jumps immediately to zero, leading to region 2 (null-strain region). In practice, the arms of the DCB, rather than being clamped, actually rotate at the crack tip [37]; this, along with fiber bridging [38,39], explains why a perturbed strain region in correspondence of the crack tip process zone is observed [40], which is defined as region 3, as shown in Figure 5, where the three segments' strain patterns are also illustrated, along with the regions. The experimental strain distribution for three different crack lengths (i.e., at different times) is shown in Figure 6. As expected, the strain peak moves accordingly to the crack front location and can be used for a rough estimation of the crack front position. Region 3, i.e., the region in correspondence of the crack tip process zone, can be approximately 10 mm in length, and sometimes is hard to identify where the boundary between regions locate, making the delamination tip location via the peak detection unreliable (see Figure 5 and Figure 6).

The experimental strain distribution for the fatigue tests is shown in Figure 7 considering different delamination lengths; note that the shape of the strain pattern is generally different with respect to the static strain patterns and, as for the static test strain patterns, it is not obvious to locate the delamination front by peak detection. Even if loading mode is substantially the same (except for the fact that during fatigue testing the load is repeated cyclically), the phenomena driving the delamination propagation are different [41], mainly affecting the process zone. For instance, the amount of fiber bridging is different under quasi-static and fatigue loading [26,38], being less in fatigue loading as compared to quasi-static loading. Therefore, as fiber bridging has been shown to affect the strain field [26,27], we expect to observe different strain patterns depending on the test type.

The real crack length (i.e., the ground truth) is observed by visual inspection of the specimen edge images. However, this procedure has a few shortcomings: i) it assumes that the crack front is perfectly straight, ii) it is often difficult to identify the crack tip, despite the edge has been painted in white, iii) checking all the images is time consuming [42]. To overcome these drawbacks a calibration method has been adopted as shown in the ASTM 5528 standard [35] and as suggested by Hojo et al. [43,44] and illustrated by Sans et al. [45], generating a least squares plot of the cube root of compliance, $C^{1/3}$, as a function of delamination length, a_c. The compliance, C, is the ratio of the load point displacement to the applied load, δ/P. The values used to generate this plot should be the load and displacements corresponding to the visually observed delamination lengths on the edge. Note that only a subset of the observed propagation values is used; the missing values can be extrapolated after the least squares plot has been generated. The delamination front shape has been visually analyzed after testing the specimens, showing that the delamination front is fairly straight; actually, it is possible to theoretically demonstrate that the $[0_n]$ unidirectional layup has the smallest energy release rate (ERR) variation across the delamination front, minimizing the delamination front curvature [46]. Additionally, calculations show that the ERR is constant in the interior of the specimens and rapidly approaches zero at the point where the delamination tip intersects the edge free surface [47]. Therefore, the three OF loops are likely to produce similar strain patterns and that the same true delamination size (observed from the specimen edge) can be associated to all three of them. Equivalently, we can imagine decoupling the OF triplets, virtually tripling the number of specimens.


Figure 5: Strain pattern along the 3 segments for specimen 1 (for a 72mm delamination length).

Figure 6: Static testing strain patterns along the 3 segments for specimen1 at different times (different crack lengths) (100, 250, 400th time steps).


Figure 7: Fatigue testing strain patterns along the 3 segments for specimen1 at different times (different crack lengths) (100, 250, 400th time steps).

3. Delamination sizing by deep learning and distributed strain measurements

Several SHM approaches targeting delamination detection, sizing and location have been published, all requiring domain expertise and preprocessing of the raw data in order to extract relevant features for the delamination assessment [28–30,48]. Delamination sizing for a DCB specimen having distributed OF sensors along its arms, i.e., the identification of the delamination front, could be performed by simply detecting the peak of the strain pattern (i.e., the minimum strain relative to each OF segment), as, in principle, the strain reaches its maximum absolute value in correspondence of the delamination tip. However, as explained above, the process zone makes it difficult to accurately locate the delamination front, as the strain field in correspondence of the crack tip is generally perturbed [32]. The task is made even more difficult if we consider that the goal is to predict the delamination size given the fatigue strain patterns, whereas the training data is generated via static tests, as this requires that some general features are extracted from the static strain patterns (see the methodology

scheme in Figure 8). Therefore, a DL approach is proposed to automatically extract the strain pattern features that can better locate the delamination tip.

The major advantage of DL methods over traditional ML methods is their ability to autonomously extract discriminative features from raw data [8]. Among DL methods, CNNs have proven to be a very efficient and effective artificial neural network (ANN) structure for image recognition and classification [8], as their architecture is specifically developed to take advantage of the fact that the input consists of images. CNNs layers are typically organized in 3 dimensions (width, height, and depth), and four operations are sequentially carried out: 1) convolution, 2) nonlinear transformation, 3) pooling, 4) classification/regression. These operations are performed in three layers: the convolutional layer, the pooling layer and the fully connected layer. The convolutional layer extracts features from the input image by sliding a filter or kernel over the input image. Note that filters in the above example are random filters. In practice, a CNN learns to adjust the weights of filters such that the last layer of the network can predict classes/perform regression with the maximum accuracy [49]. This is done during training procedure. The convolution operation is followed by a non-linear operation, such as Rectified Linear Unit (ReLU), to introduce non-linearity in CNN, as most of the real-world data that is used for the learning of CNN is non-linear. The major goal of a pooling layer is to reduce the dimensionality of feature maps. The convolutional and pooling layers transform the input image into a high-level feature map that is employed by the fully connected layer (output layer) for the classification of the input image/regression to some output variables. Full convolutional neural network architectures are formed by stacking the convolutional, pooling, and fully connected layers together. A detailed description of CNNs can be found in refs. [50,51].

CNNs have been modelled and created specifically for 2D signals and their application is not straightforward for 1D signals, needing a 1D to 2D conversion. Different conversion techniques have been utilized to convert 1D signals into 2D signals, typically posing a high computational complexity and training dataset size requirements [13,14]. To overcome such drawbacks a modified version of the 2D CNN, called 1D CNN, has been recently developed to directly operate on 1D signals such as time series and has become popular with a state-of-the-art performance in various signal processing applications [11].

In the present work, both 2D and 1D CNN architectures are deployed to assess the delamination size based on the collected strain pattern measurements. OF strain patterns are 1D and can be treated as time series, replacing time with space (OF length coordinate); therefore, 1D to 2D conversion techniques are used to feed the 2D CNN, while no preprocessing is required for the 1D CNN architecture. Three different techniques have been used to transform time-series into image-like representations: i) Gramian Angular Field, ii) Markov Transition Field and iii) Scalogram (continuous wavelet transform or CWT). The fundamental building blocks for the 2D CNN architecture are shown in Figure 9, along with the data pipeline; the same is done for the 1D CNN in Figure 10. For instance, in Figure 9 the single strain pattern is collected from the specimen and preprocessed before being fed to the 2D CNN, i.e., scaled and then transformed into a single channel 2D image (matrix of pixels). After being transformed (in Figure 9 the Gramian angular field transform is shown), the 2D image is scaled again and then fed to the CNN, where a series of layers perform sequentially the convolution and pooling operations, transforming the input image into a high-level feature map. This feature map is converted into a 1D array (i.e., flattening operation), becoming the input for the fully connected layers, ending with the single output

node that is responsible for the delamination size prediction. In Figure 10 the strain pattern is scaled and then fed directly to the 1D CNN; the convolution and pooling operations are sequentially carried out along the 1D array, producing arrays of features that are concatenated (flattening) to produce one single array that is fed to the fully connected layers, ending in the single output node.
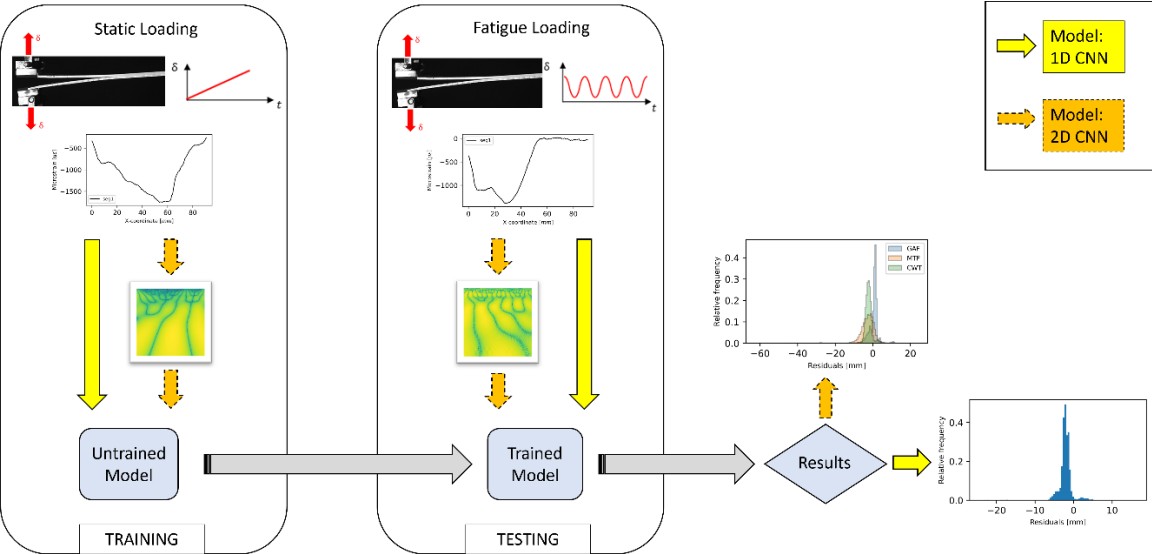

Figure 8: Methodology schematic representation.

Figure 9: 2D CNN pipeline.

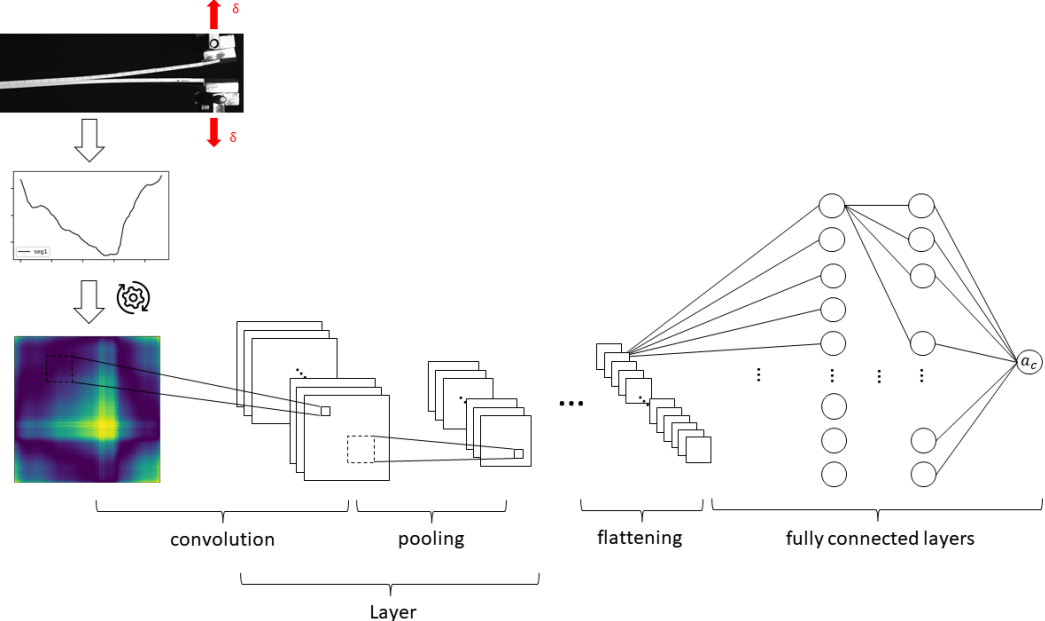Figure 10: 1D CNN pipeline.


Figure 11: Training data pipeline.

3.1. Training and validation of predictive CNN-based models

To prepare the dataset required to train the CNNs, a total of 5 DCB specimens have been statically tested. Note that each specimen has 3 OF loops (Figure 1), each accounting for 140 measuring points along its length, while the number of time observations, i.e., the number of

crack length sizes at which strain measurements are collected, is not fixed and typically depends on the test setup. As mentioned in the previous sections, we can imagine that each OF loop identifies a different specimen, tripling the overall number of specimens, as the delamination front has been acknowledged to be straight in the monitored area. Therefore, for the i-th OF segment ($i = 1,2,3$) corresponding to the j-th specimen ($j = 1, \dots, 5$) a matrix S_i^j is obtained:

$$S_i^j = \begin{bmatrix} \varepsilon_{t=0}^{x=0} & \cdots & \varepsilon_{t=0}^{x=140} \\ \vdots & \ddots & \vdots \\ \varepsilon_{t=T}^{x=0} & \cdots & \varepsilon_{t=T}^{x=140} \end{bmatrix} \tag{1}$$

where rows represent the i-th OF segment strain pattern at the inspection times ($x = 0, \dots, 140$), while columns show the strain time series of the measurement points. Clearly, for each input matrix S_i^j a corresponding output vector a_c^j is associated, where the observed crack lengths at the inspection times are collected:

$$a_c^j = \begin{bmatrix} a_{c,t=0} \\ \vdots \\ a_{c,t=T} \end{bmatrix} \tag{2}$$

Note that the output vector a_c^j is independent of the OF segment, as we assume that the crack front is straight in the observed area, as also shown by tested specimens' post-mortem analysis. The training dataset **T** is obtained by gathering all the input and output data from all the 5 DCB specimens:

$$\mathbf{T} = \begin{bmatrix} S_{i=0}^{j=0}; \ a_c^{j=0} \\ \vdots \\ S_{i=3}^{j=5}; \ a_c^{j=5} \end{bmatrix} \tag{3}$$

In practice, each row of the dataset **T** represents a strain pattern (input) and its relative delamination length (output); in total 6110 examples (i.e., rows) are collected. Note that the rows of the dataset **T** are randomly shuffled before training the CNNs. The CNN model is thus expected to take as input a strain pattern (OF loop) $[\varepsilon_t^{x=0}, \dots, \varepsilon_t^{x=140}] \in \mathbb{R}^{140}$ at the inspection time t, providing as output the estimate for the actual delamination size $a_{c,t} \in \mathbb{R}$. The training dataset is partitioned in order to keep part of the data to validate the model during training; specifically, 20% of the training data is used for validation (Table 3 summarizes the training, validation and test set size). Figure 11 illustrates the training/validation data pipeline (neglecting the 1D to 2D conversion that is needed for 2D CNNs): 3 OF segments (each segment is represented by a 140 elements array) are collected at each inspection time for each specimen (note that the number of inspection times might differ from specimen to specimen). When dealing with 2D CNNs, the arrays collecting the OF strain patterns are transformed into matrices (1D to 2D conversion), as shown in Figure 9.

In order to reduce the variance of the CNNs predictions, an "ensemble" approach is leveraged. In practice, a collection of networks with the same configuration and different initial random weights is trained on different subsets (randomly sampled) of the training dataset. Each model is then used to make a prediction and the actual prediction is calculated as the average of the predictions [53,54]. Each CNN ensemble is constituted of 10 CNN trained models.

3.2.2D CNN architecture

This section discusses the architecture of the 2D CNN. The Pyhton Keras API is used to build the DL models illustrated in the present work (built on top of Tensorflow). Figure 9 shows the 2D CNN architecture fundamental building blocks and data pipeline. The architecture is based on the general architectural principles of the Visual Geometry Group (VGG) model (featured by very small 3x3 convolution filters) [55] and consists of 2 VGG-based blocks, each comprising 2 convolutional layers and one pooling layer, for automatically extracting discriminative features from the static strain patterns, including a dropout stage (fixed dropout rate) [56], followed by a fully connected layer and a ReLU layer (see Table 1 for the detailed architecture of the proposed 2D CNN). The weights are initialized randomly, and the Adam optimizer is used (learning rate is set to 0.0001). A batch size of 64 is used, and the training consists of 30 epochs. A detailed description of the hyperparameters listed in Table 1 can be found in ref. [49]. Tuning the hyperparameters – as well as the architecture – would surely enhance the performance, but it is not the primary objective of the present work (for instance, in ref. [19] the optimal values of the CNN hyperparameters were derived based on Bayesian optimization). Another consideration is computational efficiency; the architecture is somehow limited by available computational resources. The training with the actual CNN architecture can be done in less than 8 hours on a laptop with a 4 cores Intel i7-6700HQ (2.6 GHz) CPU and a 32GB RAM. Exploring the deep network structures and parameters will be addressed in future work.

Layer name	Layer description	Output shape	Trainable Parameters
Input	140x140 single channel image	140x140	
Convolution 1	Kernel size: 3x3, strides: 1, Number of filters: 32, ReLU, batch normalization	140x140x32	320
Convolution 2	Kernel size: 3x3, strides: 1, Number of filters: 32, ReLU, batch normalization	140x140x32	9248
MaxPooling	Max Pooling Filter size: 2×2, strides: 2	70x70x32	0
Dropout	20% dropout	70x70x32	0
Convolution 1	Kernel size: 20x1, strides: 1, Number of filters: 64, ReLU, batch normalization	70x70x64	18496
Convolution 2	Kernel size: 20x1, strides: 1, Number of filters: 64, ReLU, batch normalization	70x70x64	36928
MaxPooling	Max Pooling Filter size: 2×2, strides: 2	35x35x64	0
Dropout	20% dropout	35x35x64	0
Flatten		78400	0
Fully connected ReLU		128	10035328
Dropout	20% dropout	128	0

Fully connected ReLU		1	129

Table 1: 2D CNN architecture and hyperparameters

3.3. 1D CNN architecture

This section discusses the architecture of the 1D CNN. Figure 10 shows the 1D CNN architecture fundamental building blocks and data pipeline. The employed 1D CNN architecture consists of 2 blocks, each comprising 2 convolutional layers and one pooling layer for automatically extracting discriminative features from the static strain patterns. A dropout stage (increasing dropout rate) is added to each block in order to reduce overfitting and improve generalization [56]. Finally, a fully connected layer and a ReLU layer are used to predict the delamination size based on features extracted in the convolutional and pooling layers (see Table 2 for the detailed architecture of the proposed 1D CNN). The weights are initialized randomly, and the Adam optimizer is used (learning rate is set to 0.0001). A batch size of 64 is used, and the training consists of 30 epochs. Details concerning the 1D CNN architecture can be found in ref. [15]. As for the 2D CNN, the hyperparameters have not been tuned, and the architecture is based on the 2D CNN architecture layout (2 blocks where convolution and pooling operations are sequentially performed).

Layer name	Layer description	Output shape	Trainable Parameters
Input	140x1 single channel strain pattern	140x1	
Convolution 1	Kernel size: 20x1, strides: 1, Number of filters: 32, ReLU	121x32	672
Convolution 2	Kernel size: 20x1, strides: 1, Number of filters: 32, ReLU	102x32	20512
MaxPooling	Max Pooling Filter size: 2×1, strides: 2	51x32	0
Dropout	20% dropout	51x32	0
Convolution 1	Kernel size: 20x1, strides: 1, Number of filters: 64, ReLU	32x64	41024
Convolution 2	Kernel size: 20x1, strides: 1, Number of filters: 64, ReLU	13x64	81984
MaxPooling	Max Pooling Filter size: 2×1, strides: 2	6x64	0
Dropout	40% dropout	6x64	0
Flatten		384x1	0
Fully connected ReLU		128	6160
Fully connected ReLU		1	129

Table 2: 1D CNN architecture and hyperparameters

3.4. 1D signal to 2D conversion

Convolutional layers basically mimic the cells in the human visual cortex and are thus developed primarily for 2D signals such as images and video frames [57]. Inspired by the

success of DL methods in computer vision, several studies have proposed to transform time-series into image-like representations, leading to promising results [13,14]. Three different techniques are here leveraged to encode 1D strain patterns into images, namely i) continuous wavelet transform (CWT), ii) Gramian angular field (GAF), iii) Markov transition field (MTF). Details concerning the mentioned techniques can be found in Refs. [13,14], while the hyperparameter settings for the selected encodings are described in the following. The *pyts* Python package for time series classification has been used for the GAF and MTF transforms [58].

The GAF transform is only defined for input arrays $X = (x_1, \dots, x_N)$ that satisfy the following constraint: $x_i \in [-1,1]$, thus requiring the scaling of the strain patterns as proposed by Wang and Oates [13]:

$$\tilde{x}_i = \frac{\left(x_i - max(X) + \left(x_i - min(X)\right)\right)}{max(X) - min(X)} \tag{4}$$

Alternatively, Garcia et al. [14] proposed to scale the training samples based on the full training dataset in order to keep the relationships and differences between the samples scaling, requiring that the $max(X)$ and $min(X)$ values are replaced with an upper (UB) and lower bound (LB) based on the training set distribution. If the test set contains values exceeding these bounds, the values are clipped to the bound. As we assume that it is not necessary to keep the quantitative relationships and differences between the strain pattern samples, we followed the scaling in eq. (4), not requiring the setting of any bound. The strain patterns have thus been scaled accordingly, regardless of the transform technique.

The MTF requires that a discretization of the strain patterns is performed based on $Q + 1$ bin edges, i.e., Q bins. The discretization of the time series plays a crucial role in the amount of information kept or lost by the transformation [13,14]. The number of bins has been set to 5, while their width has been set so that all bins in each sample have the same number of points.

For the scalogram, a CWT with a Morlet wavelet was utilized via the PyWavelets Python package for wavelet analysis [59].

2D encodings have been also scaled prior being fed to the CNNs, so that each element $x_{i,i}$ satisfies the following: $x_{i,i} \in [0,1]$. Encoding examples for a representative strain pattern are shown in Figure 12.

Figure 12: 1D signal (strain pattern) image encodings: Gramian angular field (GAF), Markov transition field (MTF) and continuous wavelet transform (CWT).

4. Results and discussion

The CNNs architectures (1D and 2D) developed in the previous sections are employed to predict the delamination size of fatigue loaded DCB specimens based on their OF strain patterns. The bagging ensemble learning approach is followed to increase the predictive performance by combining the predictions from multiple models; in practice, for each CNN architecture 10 CNNs are trained on different random subsets of the training dataset and the predictions are averaged. Each DCB specimen has 3 OF loops spanning the specimen width and delamination size predictions are based upon one single OF strain pattern, assuming that the delamination front is straight, thus making one strain pattern sufficient for the delamination size estimation. The training dataset, which is generated by storing strain patterns along with the relative observed delamination size from the static delamination growth tests, comprises 6110 examples, and is partitioned into training and validation subsets with a 4 to 1 ratio, respectively (see Table 3). Strain patterns have been resized to comprise 140 strain samples each, discretizing the DCB arms surface strain field, and their image encodings are 140x140 matrices; strain patterns have also been rescaled prior training/testing as shown in eq. (4). The CNNs goal is thus to learn relevant features from the raw static strain patterns (or their 2D counterparts) that could also generalize to unseen fatigue strain patterns for estimation of the delamination size. This strategy could be extended to more complex structures, avoiding the need to run time consuming fatigue test. The test dataset comprises 4 specimens which were fatigue loaded; 3 specimens were tested under load control, while the remaining one was tested under displacement control. In total, 8762 strain patterns were collected during fatigue testing (see Table 3).

The 1D CNN committee predictive performance is shown in Figure 13 (left), where the mean absolute error (MAE) is provided for each OF segment during the fatigue delamination propagation. The delamination size is predicted with a MAE which is lower than 5mm, except for the specimen #4, where the segments 1 and 3 display a much higher MAE. This is due to the poor quality of the relative measurements, which might be due to the imperfect bonding between the OFs and the specimen. Note that the specimen with the lowest MAE is specimen #2, which was loaded in displacement control (slower delamination growth rate and lower noise level). Figure 13 (right) shows the relative frequency of the prediction

residuals, providing an insight in the prediction error distribution; the mean is equal to -2.2 mm while the standard deviation is equal to 2.4 mm.

The 2D CNN committee predictive performance is shown in Figure 14, following the same scheme as for the 1D CNN. The MAE is again provided for each OF segment in Figure 14 (left). The plot follows closely the results of the 1D CNN. The GAF-based CNN generally performs better than the MTF and CWT-based CNNs, also outperforming the 1D CNN in terms of accuracy. The main differences of the different time-series to image encodings are mostly in their ability to better discriminate the features that also generalize to the unseen fatigue strain patterns. Figure 14 (right) shows the residuals relative frequency for each transform technique; overall results, also including the 1D CNN, are summarized in Table 4.

The predictive performance of the illustrated CNN architectures shows that the proposed approach is able to: i) accurately predict the delamination size based on one single strain pattern, ii) automatically extract relevant features from static strain patterns that generalize well to unseen strain patterns. Although in the current work only DCB coupons specimens have been tested, the proposed approach could be extended to assess more complex structures, also considering other types of damages affecting the strain field of the laminate.

Training set size	Validation set size	Test set size
4888	1222	8762

Table 3: CNNs training/validation/test set sizes

	1D CNN	2D CNN		
	-	GAF	MTF	CWT
$\mu\ [mm]$	-2.2	0.47	-3.12	-2.54
$\sigma\ [mm]$	2.3	2.82	5.85	3.16

Table 4: Mean and std. dev. of the prediction error for the CNN models

(a)   (b)

Figure 13: 1D CNN prediction MAE on the test set (a) and residuals relative frequency (b).

Figure 14: 2D CNN prediction MAE on the test set (a) and residuals relative frequency (b).

5. Conclusions

Delamination monitoring is a vital step towards the SHM of laminated composite structures, and the advent of DL methods is fostering the application of data-driven approaches due to their ability to automatically extract relevant features from raw data. In this paper, a CNN-based approach that can perform delamination sizing by analyzing strain patterns collected from composite DCB specimens under fatigue loading is proposed. To obtain the data to train the CNNs, 6110 strain patterns are extracted from delamination growth static tests, while fatigue strain patterns are used only for testing. Strain patterns are obtained via distributed optical fiber sensors bonded to the DCB arm top surfaces; each specimen is equipped with three OF loops spanning the specimen width, sampling the strain field with a resolution of 1.25 mm. Static strain patterns can thus be regarded as surrogate signals replacing the fatigue strain patterns, enabling a more efficient approach to the generation of training data for the delamination monitoring of composite structures. Static strain patterns are observed to qualitatively differ from fatigue strain patterns; however, it is assumed that their information content concerning delamination sizing generalize also to unseen situations. The strain patterns have been scaled prior training/testing to ensure that only the qualitative features of the signals are relevant to the delamination sizing, removing the load magnitude influence. 1D and 2D CNN architectures have been deployed for the delamination size estimation, yielding similar and very reasonable average positional estimation errors of -2.2 mm and 0.47 mm respectively.

The training data set for the delamination sizing has been prepared via 5 static delamination growth tests, while the test data set has been populated by fatigue testing 4 specimens. The 1D CNN has been trained with raw strain patterns, only requiring the scaling of the input, whereas the 2D CNN required the encoding of the strain patterns into images. Specifically, three different techniques have been utilized to transform 1D signals into 2D signals, namely the Gramian angular field, the Markov transition field and the continuous wavelet transform, all delivering reasonable results in terms of accuracy.

It is believed that the CNN-based methodology proposed in this study can be successfully applied to analyze more complex composite structures to monitor the delamination growth based on strain observations, also proving the ability of the CNNs to extract relevant features that generalize to unseen situations. The results presented in this paper are meant to foster the deployment of DL within the SHM of composite structures by enabling real-time structural damage monitoring. Important future work will involve the application of the proposed method to more complex structures, also including the possibility of using different sensing technologies. We are also quite interested in how different loading scenarios might affect the predictions, also investigating different CNNs architectures.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

References

[1] Chung DDL. Composite materials: science and applications. Springer Science \& Business Media; 2010.

[2] Reifsnider KL, editor. Damage in Composite Materials: Basic Mechanisms, Accumulation, Tolerance, and Characterization. West Conshohocken, PA: ASTM International; 1982. https://doi.org/10.1520/STP775-EB.

[3] Wisnom MR. The role of delamination in failure of fibre-reinforced composites. Philos Trans R Soc A Math Phys Eng Sci 2012;370:1850–70. https://doi.org/10.1098/rsta.2011.0441.

[4] Cawley P. Structural health monitoring: Closing the gap between research and industrial deployment. Struct Heal Monit 2018;17:1225–44. https://doi.org/10.1177/1475921717750047.

[5] Falcetelli F, Yue N, Sante R Di, Zarouchas D. Probability of detection, localization, and sizing: The evolution of reliability metrics in Structural Health Monitoring. Struct Heal Monit n.d.;0:14759217211060780. https://doi.org/10.1177/14759217211060780.

[6] Tidriri K, Chatti N, Verron S, Tiplica T. Bridging data-driven and model-based approaches for process fault diagnosis and health monitoring: A review of researches and future challenges. Annu Rev Control 2016;42:63–81. https://doi.org/https://doi.org/10.1016/j.arcontrol.2016.09.008.

[7] Elenchezhian MRP, Vadlamudi V, Raihan R, Reifsnider K, Reifsnider E. Artificial intelligence in real-time diagnostics and prognostics of composite materials and its

uncertainties{\textemdash}a review 2021;30:83001. https://doi.org/10.1088/1361-665x/ac099f.

[8]   LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015;521:436–44. https://doi.org/10.1038/nature14539.

[9]   Ye XW, Tao J, Yun CB. A review on deep learning-based structural health monitoring of civil infrastructures. Smart Struct Syst 2019;24:567–85. https://doi.org/10.12989/sss.2019.24.5.567.

[10]   Bao Y, Tang Z, Li H, Zhang Y. Computer vision and deep learning–based data anomaly detection method for structural health monitoring. Struct Heal Monit 2019;18:401–21. https://doi.org/10.1177/1475921718757405.

[11]   Fotouhi S, Pashmforoush F, Bodaghi M, Fotouhi M. Autonomous damage recognition in visual inspection of laminated composite structures using deep learning. Compos Struct 2021;268:113960. https://doi.org/https://doi.org/10.1016/j.compstruct.2021.113960.

[12]   Khan A, Ko D-K, Lim SC, Kim HS. Structural vibration-based classification and prediction of delamination in smart composite laminates using deep learning neural network. Compos Part B Eng 2019;161:586–94. https://doi.org/https://doi.org/10.1016/j.compositesb.2018.12.118.

[13]   Wang Z, Oates T. Encoding Time Series as Images for Visual Inspection and Classification Using Tiled Convolutional Neural Networks, 2015.

[14]   Garcia G, Michau G, Ducoffe M, Sen Gupta J, Fink O. Time Series to Images: Monitoring the Condition of Industrial Assets with Deep Learning Image Processing Algorithms 2020.

[15]   Kiranyaz S, Avci O, Abdeljaber O, Ince T, Gabbouj M, Inman DJ. 1D convolutional neural networks and applications: A survey. Mech Syst Signal Process 2021;151:107398. https://doi.org/https://doi.org/10.1016/j.ymssp.2020.107398.

[16]   Elenchezhian MRP, Vadlamudi V, Raihan R, Reifsnider K, Reifsnider E. Artificial intelligence in real-time diagnostics and prognostics of composite materials and its uncertainties - a review 2021;30:83001. https://doi.org/10.1088/1361-665x/ac099f.

[17]   Tabian I, Fu H, Sharif Khodaei Z. A Convolutional Neural Network for Impact Detection and Characterization of Complex Composite Structures. Sensors 2019;19. https://doi.org/10.3390/s19224933.

[18]   Damm AM, Spitzmüller C, Raichle ATS, Bühler A, Weißgraeber P, Middendorf P. Deep learning for impact detection in composite plates with sparsely integrated sensors 2020;29:125014. https://doi.org/10.1088/1361-665x/abb644.

[19]   Jung K-C, Chang S-H. Advanced deep learning model-based impact characterization method for composite laminates. Compos Sci Technol 2021;207:108713. https://doi.org/https://doi.org/10.1016/j.compscitech.2021.108713.

[20]   Yu M-H, Kim H-S. Deep-learning based damage sensing of carbon fiber/polypropylene composite via addressable conducting network. Compos Struct

2021;267:113871. https://doi.org/https://doi.org/10.1016/j.compstruct.2021.113871.

[21]   Giurgiutiu V. 16 - Structural health monitoring (SHM) of aerospace composites. In: Irving PE, Soutis C, editors. Polym. Compos. Aerosp. Ind., Woodhead Publishing; 2015, p. 449–507. https://doi.org/https://doi.org/10.1016/B978-0-85709-523-7.00016-5.

[22]   Meadows L, Sullivan RW, Brown K, Ranatunga V, Vehorn K, Olson S. Distributed optical sensing in composite laminates. J Strain Anal Eng Des 2017;52:410–21. https://doi.org/10.1177/0309324717723466.

[23]   Chandarana N, Martinez-Sanchez D, Soutis C, Gresil M. Early Damage Detection in Composites by Distributed Strain and Acoustic Event Monitoring. Procedia Eng 2017;188:88–95. https://doi.org/https://doi.org/10.1016/j.proeng.2017.04.515.

[24]   Batte LK, Sullivan RW, Ranatunga V, Brown K. Impact response in polymer composites from embedded optical fibers. J Compos Mater 2018;52:3415–27. https://doi.org/10.1177/0021998318763274.

[25]   Kuang, KSC ,, WJ C. Use of conventional optical fibers and fiber Bragg gratings for damage detection in advanced composite structures: A review . Appl Mech Rev 2003;56:493–513. https://doi.org/10.1115/1.1582883.

[26]   Stutz S, Cugnoni J, Botsis J. Studies of mode I delamination in monotonic and fatigue loading using FBG wavelength multiplexing and numerical analysis. Compos Sci Technol 2011;71:443–9. https://doi.org/https://doi.org/10.1016/j.compscitech.2010.12.016.

[27]   Sorensen L, Botsis J, Gmür T, Cugnoni J. Delamination detection and characterisation of bridging tractions using long FBG optical sensors. Compos Part A Appl Sci Manuf 2007;38:2087–96. https://doi.org/https://doi.org/10.1016/j.compositesa.2007.07.009.

[28]   Kesavan A, John S, Herszberg I. Strain-based structural health monitoring of complex composite structures. Struct Heal Monit 2008;7:203–13. https://doi.org/10.1177/1475921708090559.

[29]   Milanoski DP, Loutas TH. Strain-based health indicators for the structural health monitoring of stiffened composite panels. J Intell Mater Syst Struct 2021;32:255–66. https://doi.org/10.1177/1045389X20924822.

[30]   Takeda S, Minakuchi S, Okabe Y, Takeda N. Delamination monitoring of laminated composites subjected to low-velocity impact using small-diameter FBG sensors. Compos Part A Appl Sci Manuf 2005;36:903–8. https://doi.org/https://doi.org/10.1016/j.compositesa.2004.12.005.

[31]   Takeda S, Okabe Y, Takeda N. Monitoring of Delamination Growth in CFRP Laminates using Chirped FBG Sensors. J Intell Mater Syst Struct - J INTEL MAT SYST STRUCT 2008;19:437–44. https://doi.org/10.1177/1045389X06074085.

[32]   Drake DA, Sullivan RW. Prediction of delamination propagation in polymer composites. Compos Part A Appl Sci Manuf 2019;124:105467. https://doi.org/10.1016/J.COMPOSITESA.2019.05.035.

[33] Bocherens E, Bourasseau S, Dewynter-Marty V, Py S, Dupont M, Ferdinand P, et al. Damage detection in a radome sandwich material with embedded fiber optic sensors. Smart Mater Struct 2000;9:310–5. https://doi.org/10.1088/0964-1726/9/3/310.

[34] Lu P, Lalam N, Badar M, Liu B, Chorpening B, Buric M, et al. Distributed optical fiber sensing: Review and perspective. Appl Phys Rev 2019;6:41302. https://doi.org/10.1063/1.5113955.

[35] ASTM D5528-01. Standard test method for mode I interlaminar fracture toughness of unidirectional fiber-reinforced polymer matrix composites. Am Stand Test Methods 2014;03:1–12.

[36] Ripling EJ, Mostovoy S, Corten HT. Fracture Mechanics: A Tool for Evaluating Structural Adhesives. J Adhes 1971;3:107–23. https://doi.org/10.1080/00218467108081158.

[37] Škec L, Alfano G, Jelenić G. On Gc, Jc and the characterisation of the mode-I fracture resistance in delamination or adhesive debonding. Int J Solids Struct 2018;144–145:100–22. https://doi.org/https://doi.org/10.1016/j.ijsolstr.2018.04.020.

[38] Khan R. Fiber bridging in composite laminates: A literature review. Compos Struct 2019;229:111418. https://doi.org/https://doi.org/10.1016/j.compstruct.2019.111418.

[39] Spearing SM, Evans AG. The role of fiber bridging in the delamination resistance of fiber-reinforced composites. Acta Metall Mater 1992;40:2191–9. https://doi.org/10.1016/0956-7151(92)90137-4.

[40] Xie J, Waas AM, Rassaian M. Estimating the process zone length of fracture tests used in characterizing composites. Int J Solids Struct 2016;100–101:111–26. https://doi.org/https://doi.org/10.1016/j.ijsolstr.2016.07.018.

[41] Pascoe JA, Alderliesten RC, Benedictus R. Methods for the prediction of fatigue delamination growth in composites and adhesive bonds - A critical review. Eng Fract Mech 2013;112–113:72–96. https://doi.org/10.1016/j.engfracmech.2013.10.003.

[42] Brunner AJ, Murphy N, Pinter G. Development of a standardized procedure for the characterization of interlaminar delamination propagation in advanced composites under fatigue mode I loading conditions. Eng Fract Mech 2009;76:2678–89. https://doi.org/https://doi.org/10.1016/j.engfracmech.2009.07.014.

[43] Hojo M, Tanaka K, Gustafson CG, Hayashi R. Effect of stress ratio on near-threshold propagation of delimination fatigue cracks in unidirectional CFRP. Compos Sci Technol 1987;29:273–92. https://doi.org/10.1016/0266-3538(87)90076-5.

[44] Hojo M, Ando T, Tanaka M, Adachi T, Ochiai S, Endo Y. Modes I and II interlaminar fracture toughness and fatigue delamination of CF/epoxy laminates with self-same epoxy interleaf. Int J Fatigue 2006;28:1154–65. https://doi.org/10.1016/j.ijfatigue.2006.02.004.

[45] Sans D, Renart J, Costa J, Gascons N, Mayugo JA. Assessment of the influence of the crack monitoring method in interlaminar fatigue tests using fiber Bragg grating sensors. Compos Sci Technol 2013;84:44–50. https://doi.org/https://doi.org/10.1016/j.compscitech.2013.04.022.

[46] Sun CT, Zheng S. Delamination characteristics of double-cantilever beam and end-notched flexure composite specimens. Compos Sci Technol 1996;56:451–9. https://doi.org/https://doi.org/10.1016/0266-3538(96)00001-2.

[47] Schön J, Nyman T, Blom A, Ansell H. A numerical and experimental investigation of delamination behaviour in the DCB specimen. Compos Sci Technol 2000;60:173–84. https://doi.org/10.1016/S0266-3538(99)00113-X.

[48] Broer A, Galanopoulos G, Benedictus R, Loutas T, Zarouchas D. Fusion-based damage diagnostics for stiffened composite panels. Struct Heal Monit n.d.;0:14759217211007128. https://doi.org/10.1177/14759217211007127.

[49] Aghdam H, Heravi E. Guide to Convolutional Neural Networks: A Practical Application to Traffic-Sign Detection and Classification. 2017. https://doi.org/10.1007/978-3-319-57550-6.

[50] Kotu V, Deshpande B. Chapter 10 - Deep Learning. In: Kotu V, Deshpande B, editors. Data Sci. (Second Ed. Second Edi, Morgan Kaufmann; 2019, p. 307–42. https://doi.org/https://doi.org/10.1016/B978-0-12-814761-0.00010-1.

[51] Yamashita R, Nishio M, Do RKG, Togashi K. Convolutional neural networks: an overview and application in radiology. Insights Imaging 2018;9:611–29. https://doi.org/10.1007/s13244-018-0639-9.

[52] Kiranyaz S, Ince T, Hamila R, Gabbouj M. Convolutional Neural Networks for patient-specific ECG classification. 2015 37th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc., 2015, p. 2608–11. https://doi.org/10.1109/EMBC.2015.7318926.

[53] Krizhevsky A, Sutskever I, Hinton GE. ImageNet Classification with Deep Convolutional Neural Networks. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ, editors. Adv. Neural Inf. Process. Syst., vol. 25, Curran Associates, Inc.; 2012.

[54] Ripley BD. Pattern recognition and neural networks. 2014. https://doi.org/10.1017/CBO9780511812651.

[55] Simonyan K, Zisserman A. Very Deep Convolutional Networks for Large-Scale Image Recognition. ArXiv 14091556 2014.

[56] Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. J Mach Learn Res 2014;15:1929–58.

[57] Lindsay GW. Convolutional Neural Networks as a Model of the Visual System: Past, Present, and Future. J Cogn Neurosci 2021;33:2017–31. https://doi.org/10.1162/jocn_a_01544.

[58] Faouzi J, Janati H. pyts: A Python Package for Time Series Classification. J Mach Learn Res 2020;21:1–6.

[59] Lee GR, Gommers R, Waselewski F, Wohlfahrt K, O&#8217;Leary A. PyWavelets: A Python package for wavelet analysis. J Open Source Softw 2019;4:1237. https://doi.org/10.21105/joss.01237.