Towards automated characterisation of fatigue damage in composites using thermoelastic stress analysis

W. J. R. Christian[1]*, P. Lambert[1], I. E. Tabrizi[1], C. A. Middleton[1], C. Przybyla[2], E. A. Patterson[1]

1. School of Engineering, University of Liverpool, Liverpool, UK

2. Air Force Research Laboratory, Ohio, USA

*Email: W.J.R.Christian@liverpool.ac.uk

Abstract

Composite materials demonstrate complicated fatigue behaviour due to their microstructure and the varied types of defects that can occur during loading. This necessitates experimentation to determine their performance under loading. In this study an algorithm is introduced for identifying and categorising different defects forming during fatigue tests. Thermoelastic stress analysis was used to obtain high spatial and temporal resolution stress information from the surface of notched composite specimens. Specimens with three different geometries were loaded in tension-tension fatigue to failure. An algorithm was used to identify when and where matrix cracking and delaminations formed within the specimens as well as quantify how this changed over time. By improving how damage events are identified and characterised, the algorithm reduces the amount of time needed to process experimental fatigue data and helps to provide greater understanding of fatigue processes in new materials from early small-scale cracking all the way to final failure.

Keywords

Fatigue, Damage mechanics, Mechanical testing, Thermal analysis

1. Introduction

The growth of fatigue damage in composite materials is complex in comparison to that seen in metals. This stems from the microstructure of composite materials, which means cracks grow preferentially in some directions, whilst being rapidly arrested in others. Other defects lend further complexity to the process, with delaminations between plies being easy to form and inhomogeneities in the material such as voids or inclusions being common [1]. This all leads to damage mechanics that are very difficult to predict and thus composite structures are designed with significant levels of conservatism, leading to excessive weight and pessimistic estimates of lifetime [2]. As many civil aircraft structures are now predominantly made from composites, this means that significant questions will start to arise regarding the strength of these materials and their failure modes as they near retirement over the coming decades [3]. Composite materials are also being used in ever more challenging environments, such as high temperature composites for space applications [4] where similar questions regarding fatigue life arise. In order to answer these questions, greater understanding is needed on how composites degrade over time, and thus experiments are performed to obtain data to inform damage models.

Fatigue of composites is a widely studied phenomenon due to its inherent complexity and importance to industry, thus in this introduction the focus is on the fundamental processes involved and approaches for experimentally characterising fatigue. For consideration of the wider topic and the associated literature the reader is encouraged to refer to a recent review article, such as [5]. The process by which fatigue occurs in composites is highly varied, however, typically matrix cracking initially occurs, before networks of cracks start to nucleate and interact with each other [5]. In the final stages, delaminations and dominant cracks are initiated resulting in the laminate breaking up into sub-laminates and rapidly losing stiffness [3]. The amount these defects affect the material can be observed in how stiffness decreases during tests, with stiffness initially reducing quickly, before slowing to a steady rate of reduction for the majority of the fatigue life, which accelerates in the final stage [6]. These reductions in stiffness have been one of the ways that predictions of material degradation have been made [7]. This approach typically focuses less on which defects are forming and instead considers empirical functions that represent the effect of all possible defects on the material behaviour. Other studies have produced empirical models that link strength reduction to time, based on experimental measurements [2]. Fatigue processes can also be considered in terms of the defects produced by adopting approaches commonly applied to metals such as S-N curves and Paris laws to represent how cracks form [8]. Delaminations, which typically occur towards the end of fatigue life, can also be considered, with models of their initiation and propagation having been

explored [9]. Underpinning all these models is a need for large amounts of informative test data causing many studies to explore how this can be obtained.

Approaches exist to accurately determine damage in composites right down to the microscale using techniques such as optical microscopy [1], scanning electron microscopy [4], and x-ray computed tomography [10]. However, these techniques typically cannot be applied during fatigue loading and thus usually only provide information on the final state of a material. Pausing a test in order to perform an inspection is one option, however this drastically reduces the rate at which new materials can be characterised and may also change the dynamics of the failure process [11]. Optical measurement techniques enable monitoring of structures as they are loaded, however they typically require translucent specimens to identify cracks growing through the material [1, 8] or provide only information from the surface of the material. Optical strain measurement using digital image correlation is another approach that can be used to monitor fatigue processes in composite materials [12]. However, this requires fast exposure times or pausing of the experiment to obtain suitable images for correlation. Digital image correlation also has limitations for spatial resolution, which can be low as square subsets of pixels must be tracked across images in order to calculate strain [11].

An approach which can be used to collect data during loading, without interrupting that loading is thermoelastic stress analysis (TSA). This technique uses a sensitive infrared camera to monitor small fluctuations in temperature on the surface of the specimen. These temperature fluctuations are caused by the thermoelastic effect and are proportional to the magnitude of the cyclic stresses within the material [13]. The technique has been well used for studying fatigue processes in metals [14, 15], and has also been used to monitor fatigue in composite materials [6, 16, 17]. For an orthotropic material loaded under plane-stress conditions the temperature fluctuations, ΔT, are related to stresses by [13]:

$$\Delta T = -\frac{T_0}{\rho C_p}(\alpha_{11}\Delta\sigma_{11} + \alpha_{22}\Delta\sigma_{22}) \quad (1)$$

where T_0 is the mean temperature of the material, ρ is the density, C_p is the heat capacity at constant pressure, α_{11} and α_{22} are the coefficients of thermal expansion in the principal material directions, and $\Delta\sigma_{11}$ and $\Delta\sigma_{22}$ are the fluctuations in stresses in the principal material directions. In order for the thermoelastic effect to be measurable it is necessary for the stress fluctuations to be applied at a high enough frequency to ensure adiabaticity. This means that heat is not lost due to conduction through the material during load cycles. The frequency at which adiabatic behaviour occurs is dependent on the material, and for composite laminates can be difficult to determine. Thus care must be taken when

choosing surface preparation and loading frequencies for laminates [18]. One of the challenges of using any optical technique to monitor a fatigue process is the amount of data that these systems can record. This often means investigators are forced to monitor statistics that represent changes across large areas of specimens [19] and thus potentially miss the intricacies of the fatigue process. Recent work on processing mechanical test data have demonstrated that orthogonal decomposition can be used to reduce the dimensionality of full-field data such that changes in the shape of strain-fields can be monitored and displayed in the form of damage-time maps [20]. These maps mean investigators do not need to manually analyse entire datasets to understand how damage has propagated. In this study the approach is for the first time applied to the fatigue of composite materials such that delaminations forming during a tensile-tensile fatigue experiment can be detected and their size and location quantified. It is also the first time damage-time maps have been obtained from processing thermoelastic stress analysis data. The approach is further extended such that sub-millimetre cracks forming near the surface of specimens can be detected and quantified.

2. Experimental Method

Carbon-fibre reinforced polymer laminates with a $[-45/45]_{3S}$ layup were manufactured out of unidirectional prepreg (RP542-2, PRF, UK). These laminates had a size of 145 mm by 245 mm and were cured using a hot-press (APV-2525, Meyer, Germany) at a pressure of 5 bar. The pressure was first applied and then the press was heated to 120 °C at a ramp rate of 10 °C/min. The specimens were then held at this temperature for 1 hour. After curing, the laminates were allowed to cool naturally down to room temperature before they were removed from the press. The resulting laminates had a nominal thickness of 3.2 mm. The laminates were then cut into 40 mm by 220 mm long specimens using a water jet cutter. Stress raisers were machined into the specimens by first drilling a hole at the centre of the specimens using a 5 mm diameter silicon carbide coated crown-tipped drill (Karnasch, UK) to minimise the potential for drilling induced delaminations. The stress raiser geometry was then milled using a Computer Numerically Controlled (CNC) milling machine fitted with a 4 mm silicon carbide milling bit. Three different stress raisers were machined into the specimens: a horizontal notch, a 45° notch and a -45° notch. Three specimens were loaded for each geometry resulting in a set of nine specimens in total. All notches had the same dimensions, with the only change being their orientation. The notches were stadium shaped, i.e. with semi-circular ends joined by a rectangle. The width of the notches was 6 mm and their length was 18 mm. Before loading the specimens, 40 mm by 40 mm by 3.5 mm crossply glass fibre tabs were bonded to the ends of the specimens on both sides using a structural adhesive (2216 Scotch-Weld, 3M, USA). The tabs helped spread the force applied by the jaws of the test machine during the fatigue test ensuring failure occurred at the stress-raiser and

not in the gripped region of the specimens. An example specimen with its stress raiser machined and tabs bonded is shown in Figure 1. To check for drilling induced delaminations around the stress-raiser, pulse-echo ultrasound scans were performed using a 2-axis C-scan machine (Midas, UK) using a 10 MHz crystal focused transducer. Finally, a thin coating layer of graphite (Graphit33, Kontakt Chemie, Belgium) was applied to ensure a uniform emissivity for thermal measurements.

Specimens were tested in a servo-hydraulic test machine (8501, Instron, USA) operating under load control. All specimens were loaded using the same load ratio $R = F_{min}/F_{max} = 0.1$, where F_{min} and F_{max} are the minimum and maximum forces applied during a cycle. The 0° notch specimens were loaded with a maximum load of 5.20 kN. The ±45° specimens were loaded at a higher maximum load of 6.05 kN to ensure the average stress in the ligament of the specimens was equal for the different geometries. All specimens were loaded at a frequency of 19 Hz, this frequency was primarily chosen to ensure that failure took place over a period of around one to three hours so that progression could be monitored. It was also confirmed to be high enough to ensure adiabatic loading of the specimens. At the start of the tests the specimens were first loaded to their mean value in 20 s. The cyclic load was then ramped up over 20 s to the desired value. The tests were stopped after the specimens had catastrophically failed.

Specimens were monitored using a commercial TSA system (DeltaTherm, Stress Photonics, USA) consisting of a staring array thermal camera (SC7650, FLIR, USA) with a 320 by 256 pixels InSb sensor. The camera was positioned at a distance of 710 mm from the specimen and a 50 mm lens (Asio-F/2.3-MWIR, Janos Technology, USA) was attached resulting in a nominal spatial resolution of 5.37 pixels/mm. The camera operated at a frame rate of 328 Hz with the resulting frames processed by the TSA system with an integration time of 4 s and thermoelastic data captured every 6 s. The load applied by the test machine was used as the loading signal for the TSA system. The specimens were monitored on the same side, except for the -45° specimens. These specimens were monitored on their opposite side resulting in the -45° specimens appearing to have the same profile as the 45° specimens shown in this study but displaying different damage behaviour during the tests. This was done as the orthotropic and laminate nature of composite materials mean such changes result in different stress distributions and thus different fatigue behaviour. After completion of each test, the specimens were photographed and ultrasonically scanned to determine the final damage morphology.

3. Results

3.1. Fatigue experiments

The specimens were all loaded at a constant amplitude until catastrophic failure occurred. The final appearance of all specimens was of two dominant cracks originating from either side of the stress-raiser and travelling in the fibre direction of the surface ply, shown in Figure 2. Ultrasonic C-scans of the specimens indicated that substantial delamination had occurred around the location of the dominant surface cracks. These delaminations were also visible on the edges of the specimen at every ply-interface through the full thickness of the specimen. Post-failure photographs and ultrasound only provide information about the very final stages of the fatigue process, and thus the TSA measurement system was used to obtain information about how the specimens degraded during the test.

The TSA measurement system provided a series of maps at every 6 s of the test, the most important for this study being the DC map, X map and Y map. All of these maps are outputted from the TSA system in arbitrary units. The DC map shows the uncalibrated temperature of the specimen. The X map shows the amplitude of temperature fluctuations that are in-phase with the loads applied to the specimen by the test machine. The Y map shows the amplitude of the temperature fluctuations that are out-of-phase with the loading signal. The outputs from a TSA system and how they are obtained are described in detail in [13]. When components are loaded at a high enough frequency to be adiabatic, the X map is proportional to the amplitude of the stresses in the specimen and the Y map contains just measurement noise. However, as damage starts to form in specimens the temperature fluctuations start to move out of phase with the loading signal, because the initiation and propagation of damage generates localised heating. Thus, whilst the X maps still contain much of the information about the damage state of the specimen the Y maps also contribute as well. As this study is focused on assessing damage in composites based on changes to the stress field, it was decided to process the R maps, which contain information about the magnitude of the combined X and Y maps by calculating for every pixel:

$$R_{i,j} = \sqrt{{X_{i,j}}^2 + {Y_{i,j}}^2} \tag{2}$$

where $R_{i,j}$ is the magnitude for the pixel at location (i,j) in the map, and $X_{i,j}$ and $Y_{i,j}$ are the corresponding pixels from the X and Y maps respectively. Examples of R maps for a specimen with a 45° notch are shown in Figure 3. These maps can be directly interpreted to identify damage, however this requires substantial time and experience. It is also likely that damage events will be missed when manually reviewing data from tests that last many hours as the test would have resulted in thousands

of R maps. Thus an algorithm was developed to process these maps sequentially based on the strain-based damage monitoring approach developed for DIC data in [20]. The algorithm used to automate the processing of the R maps is described in the next three sub-sections and is also summarised in the form of a flow chart in Figure 4.

3.2. Processing of R maps

Orthogonal decomposition was used to process the datasets such that different defects could be identified. The specimens all had stress raisers machined into their centre and thus before the data could be decomposed it was first necessary to mask the R maps to only include areas that correspond with the specimen surface. This can be done manually; however, to reduce the amount of input from the investigator, an algorithmic method was developed based on examining the DC maps. At the start of the tests, the DC maps show the uncalibrated temperature of the specimens and its surroundings. Typically, the specimen will have a different emissivity and slightly different temperature to the surrounding environment and thus locations in the map where the specimen is not present have a substantially different intensity, as shown in Figure 5. These locations can be easily identified by simply using Otsu's method of thresholding [21] resulting in a specimen mask that was applied to the R-maps. In these experiments, the emissivity of the background was significantly different from the specimen to enable this process; however, if this had not been the case then an object with a different uniform emissivity could have simply been placed behind the specimen to achieve the same results.

When performing TSA the edges of specimens can give erroneous information. This is because as the edge moves into and out of a sensor pixel with each cycle it results in large variations in the infrared signal for the affected pixels. These variation are in phase with the load signal and thus the R map can have very high intensities near the edges that do not correspond with stresses at those locations. These artefacts can be seen clearly in Figure 3 around the notch and on the two vertical edges. As these pixels cannot be used to assess damage in the specimen they were removed by morphological erosion of the mask using a disk structuring element with a radius of 6 pixels, or 1.24 mm, which was chosen as it was just larger than the amount of displacement experienced by any point on the specimen surface during fatigue as measured by the test machine.

Subsequently the R maps were orthogonally decomposed using QR factorization-based orthogonal decomposition [22] in order to extract information about the state of the specimen. The QR-factorization approach works by first finding a set of orthogonal polynomials suitable for decomposing the area of the specimens corresponding to the mask. Once these polynomials have been found they

can be used to efficiently calculate feature vectors, $\boldsymbol{f}$, that represent the data in each R map, but in a dimensionally reduced way as:

$$\boldsymbol{f} = \frac{1}{\sqrt{N}} \boldsymbol{Q}^T \boldsymbol{d} \tag{3}$$

where N is the number of pixels in the masked dataset, and $\boldsymbol{d}$ is an R map converted to a column vector. The matrix $\boldsymbol{Q}$ is the output of the MATLAB function `qr`, which calculates an orthogonal basis using 2D monomials evaluated at the pixel locations for the corresponding rows in vector $\boldsymbol{d}$. A full explanation of the decomposition algorithm used is given in [22].

The maps were decomposed using monomials up to a maximum order of 10. The number of coefficients in a feature vector goes up with the order of monomials according to the triangular number series. This means the resulting feature vectors contained 66 coefficients as it is the eleventh, i.e. $(10 + 1)^{th}$, triangular number. This is a significant reduction in data dimensionality compared to around 43,000 pixels prior to decomposition. This maximum order was chosen to maintain consistency with other studies that have applied orthogonal decomposition to TSA data [23]. These feature vectors contained information about large defects and collections of defects that covered several millimetres or more of the specimen surface. Small individual defects, such as matrix cracking and fibre breakage, typically result in sub-millimetre changes to the R map and thus are filtered out by decomposing the data. These subtle changes in the R map can be explored by reconstructing the data from its feature vector and then calculating the difference between the reconstruction and the original map. This difference between the two maps is called the reconstruction error, an example of these error maps can be seen at the top of Figure 6. Thus, after orthogonal decomposition, the algorithm has generated feature vectors that vary with time to describe gross changes due to large cracks or delaminations across many millimetres of the specimen. The algorithm has also generated maps of reconstruction error that also vary with time and contain information about highly localised changes due to matrix cracking and fibre breakage. These two different forms of data were subsequently processed separately to obtain maps that clearly show the two different defects, as shown as the divide in the flow chart from the "Decompose to feature vector" process in Figure 4. The processing of the reconstruction error maps to explore cracks on or near to the observed surface is described in the next sub-section.

3.3. Tracking crack formation

To obtain maps of crack formation near the monitored surface of the specimens, each reconstruction error map was subtracted from the reconstruction error map obtained from a time period, h, earlier.

The choice of time period is arbitrary, low values result in substantial amounts of noise in the outputs but high values result in reduced temporal resolution and thus rapid events can be lost. As the tests in this study typically lasted for around two hours, 60 seconds was chosen as a suitable value. This resulted in a difference map showing where the stresses had changed substantially during the preceding 60 second period, an example map is shown at the bottom of Figure 6. Sequential difference maps were calculated across overlapping time periods, resulting in a new difference map being obtained at the same rate as R maps were measured. The difference maps were then temporally-smoothed using a Gaussian filter applied to every pixel such that only sustained differences in the maps were retained. The standard deviation of the Gaussian filter was set to the time period used to calculate the difference maps, h. This was done as cracks would be visible for that length of time, equivalent to ten sequential difference maps, whilst measurement noise would change with every map. After temporal filtering the difference map clearly showed the location of individual matrix cracks and fibre breakages; however, the magnitude of these maps changed during the test, making it difficult to compare early maps with later ones. To compensate for this, the difference maps were normalised by subtracting the mean value of the map from every pixel and dividing through by the standard deviation of the map. This resulted in maps of the specimen showing where at any moment in time cracks were formed, an example is shown in the top-left of Figure 7. These are referred to here as near-surface crack maps. Example videos of how these near-surface crack maps change over time are available in the supplementary materials for this study.

To make it easier to locate the moment in time when crack events occurred, the maximum change in the near-surface crack map as a function of time was plotted, as shown in the bottom-left of Figure 7. Due to the normalisation of the difference maps, the noise in the maps had a mean of zero and standard deviation of unity. This means that when no damage formed the maximum value of any pixel in the map was around 5, which is the expected maximum if the value of every pixel was randomly generated. When near-surface cracks occurred, the maximum peaked at values of around 10 to 60. Thus, periods during the test when substantial amounts of cracking occurred can be identified using these plots.

The near-surface crack maps only showed where cracks formed at a particular time. Some effort is still required on the part of the investigator to determine the crack propagation and the evolution of the crack shape during a test. Thus, the near-surface crack maps were further processed to obtain damage-time maps showing crack propagation. As the near-surface crack maps were all normalised, the noise in each map when a crack was static was Gaussian, when a crack formed the map ceased to be Gaussian and pixels at the location of the cracks were outliers with significantly higher values. Thus,

simple statistics were used to determine a threshold for identifying regions where these significant changes were occurring. The inverse normal function was used to determine a threshold such that there is a 1% chance of a pixel being randomly above the threshold in any test. For the typical number of pixels and R maps captured during the tests in this study that threshold was determined as:

$$Threshold = Z^{-1}\left(\frac{1}{100 \times N_m \times N_p}\right) = 6.3 \quad (4)$$

where $Z^{-1}(\cdot)$ is the inverse normal function, N_m, is the number of maps captured in a test, and N_p, is the number of pixels in the maps. Pixels with values greater than the threshold defined in equation (4) indicated the location of damage in the near-surface crack map. Thus, thresholding each map resulted in binary images showing where in the maps cracks formed or propagated at a particular moment in time. These binary images were then combined to obtain damage-time maps that showed the location and shape of cracks as a function of time. These maps were calculated by first defining a blank damage-time map. As each new binary image showing where damage occurred was obtained, the pixels in the damage-time map at the non-zero locations in the binary image were updated to represent the time at which the damage occurred, i.e., when the binary image was obtained. If the damage-time map already had damage recorded at a particular location then the pixels were not updated. This meant that the damage-time maps show the earliest time in a test at which damage occurred at a particular location. The final damage-time maps showing cracks for three example specimens are shown in Figure 8.

3.4. Obtaining delamination damage-time maps

Damage-time maps showing delaminations were obtained by processing the feature vectors corresponding to each R map. The procedure followed was similar to how damage-time maps were obtained from full-field digital image correlation measurements of quasi-static tests in a previous study [20] and thus is only briefly explained here. Each feature vector was first subtracted from the feature vector from a time period h, earlier resulting in a vector of differences, $\boldsymbol{f}_d(t)$:

$$\boldsymbol{f}_d(t) = \boldsymbol{f}(t) - \boldsymbol{f}(t-h) \quad (5)$$

where $\boldsymbol{f}(t)$, is the feature vector corresponding to the R map captured at time t. For consistency this time period was kept the same as that used for obtaining the near-surface crack maps and their associated damage-time maps. The vector of differences was then reconstructed using the procedure described in [22] to obtain a map showing where on the specimen changes to the stress field had occurred over the preceding time period, h, before the latest feature vector had been captured.

During this reconstruction, the first term in the vector of differences was set to zero to remove the effect of any magnitude fluctuations that are not localised to one particular part of the specimen. To identify changes due to damage propagation, a threshold was defined based on the 99.9th percentile of the reconstruction of the first vector of differences, which was obtained when the time was equal to h. All subsequent reconstructions were binarized using this threshold, resulting in binary images that indicate regions within the reconstruction where changes substantially greater than measurement noise had occurred. These binary images were then combined to form damage-time maps in the same way as described in the previous sub-section. This resulted in damage-time maps which show the locations at which delaminations had first been detected, examples of the resulting damage-time maps are shown in Figure 9.

3.5. Damage progression

The near-surface crack maps allowed observations of damage initiation and propagation during most of the test. It was only in the last 5 to 10% of the test time that damage propagation was so rapid and wide-spread that these maps became difficult to interpret, at which point the delamination damage-time maps began to show changes to the specimens spread across large areas near the notches. These areas likely represent the formation of the delaminations evident in the photographs and ultrasonic C-scan data at test completion. From observing all data, three stages to the tests were identified. The transitions between these stages are indicated for a 0° specimen in Figure 10 by the dashed lines, however similar behaviour was observed for all geometries. The transitions between the stages were identified by observing both the near-surface crack damage-time maps and the delamination damage-time maps shown in Figures 8 and 9 respectively. In the first stage, the near-surface crack maps indicated localised damage appearing near the edges of the notch, with cracks initiating at the notch and propagating outwards. The second stage was more sustained, with the near-surface crack maps and their associated damage-time maps, as shown in Figure 8, indicating damage occurring at random locations within the general vicinity of the notches. During this second stage of the tests, two different modes of damage were observed in the near-surface crack maps. The first mode consisted of small, approximately round, zones of less than 1 mm diameter. The second mode were linear indications, these sometimes increased in length and split apart into two round zones that moved away from each other. An example of this second mode is shown Figure 7. To view these different modes during the tests the reader is encouraged to view the supplementary material for the study which shows videos of the near-surface crack maps throughout the entirety of the tests for the three specimens shown in Figure 2. The frequency of peaks occurring in the maximum change value shown in Figure 10 appeared to increase as the test progressed until the start of the final stage of the test. In the final stage, the

stress in large areas of the specimen started to rapidly change and thus the transition to the final stage is given by the earliest time damage is detected in the delamination damage-time maps, shown in Figure 9. These three stages were observed for all specimens regardless of stress-raiser geometry.

The area of cracked material was also explored by counting the number of pixels indicating the presence of damage as the crack damage-time maps were updated. There was substantial variation in the fatigue life of the different specimens due to inherent variability of composite materials under fatigue loading [24] and the fact that the specimens had different geometries. The area of cracked material was thus plotted against time normalised by the lifespan of each specimen in Figure 11. Broadly, the specimens had a similar behaviour with the cracked area increasing with time through their lifespan at similar rates. The three stages previously mentioned are also visible in Figure 11, with the gradient of the curves being initially higher, before becoming constant for most of the test and finally increasing again as failure neared.

4. Discussion

Optical methods have been well-used in the study of fatigue processes in materials. However, a remaining challenge is how to process the data such that it can be rapidly interpreted and used to inform the design of structures. The tests in this study typically lasted around two hours, resulting in hundreds of datasets captured by the TSA system. To manually analyse such data to reveal changes would take substantial amounts of time. It is also likely that small changes due to fibre breakage or matrix cracking might be missed as they only result in subtle changes to the stress field. Some fatigue tests in industry can last weeks or even months, as many years of service are simulated within the laboratory, and thus there is a clear industrial need for automating the processing of data. By processing the data using the algorithm introduced in this study, it is possible to rapidly identify when and where damage occurs and thus infer the type of damage without manually analysing the entire dataset.

The algorithm shown in Figure 4 provides information about the time-dependent formation of different defects in the specimens. The near-surface crack maps and their associated damage-time maps principally show defects that form in the surface of the ply visible to the camera. For example, Figure 7 shows a matrix crack forming on the surface of the specimen. The crack is initiated at a single point on the specimen and then grows outwards in both directions parallel to the fibres in the ply. This process can be seen repeatedly occurring in all of the specimens tested. The technique can detect very small cracks down to sub-millimetre scales despite the vertical motion of the specimen as the stress concentrations around the crack tips are typically larger than the cracks themselves. From comparing

372 the near-surface crack maps with the layup of the specimens and the post-failure photographs it is
373 possible to determine that the cracks mostly form along the fibre direction in the top two plies and
374 propagated away from the edge of the notch. Thus it can be inferred that the damage detected in the
375 near-surface crack maps corresponds to the top two plies of the specimen. If a specimen was
376 monitored on both sides then four plies could be monitored, equivalent to around a third of the
377 specimen thickness in these tests. This could provide ample data for fitting statistical models of crack
378 location and propagation rates such as the Paris-law-like empirical models developed in [25]. Such
379 models could be used to help inform the design of structures and inform decisions about inspection
380 and maintenance regimes. Another potential way to obtain information about crack formation
381 through the thickness is to monitor the edges of specimens, where the formation of cracks can be
382 observed during both static and fatigue tests due to sudden rises in temperature [17].

383 From observing the near-surface crack maps and the plots of maximum change in these maps as a
384 function of time, it is possible to identify different stages in the fatigue process of the composite
385 specimens, these stages are marked on Figure 10. In the first stage there are many cracks that are
386 detected near the edge of the notch. From observing the uncalibrated R maps from the TSA
387 measurement system, see example at top of Figure 3, these locations correspond to the high stress
388 locations around the notch and thus, would be expected to be the location at which damage initiated
389 first. The fibres at the cut edges also provide less reinforcement as they are no longer continuous in
390 the direction of the loading and thus more prone to fibre pull-out, leading to matrix cracking. Some of
391 the crack initiation near the edge of the specimens will have been missed due to the masking
392 procedure to remove the TSA edge effect described in Section 3.2. Some TSA systems come with
393 motion compensation algorithms that can be used to resolve the thermoelastic data much closer to
394 the edge to observe these cracks; however, this was unavailable in this study. In the second stage, the
395 rate of crack initiation and propagation decreased, such that it was possible to observe discrete events
396 as damage occurred. The damage appears to form randomly along lines running from the notch to the
397 edge of the specimen in directions parallel to the fibres. Many of these damage initiation events have
398 a characteristic pattern consisting of two high intensity spots forming in the near-surface crack maps
399 and propagating away from each other in the fibre direction, see Figure 7. This suggests that these are
400 individual matrix cracks forming, with the high intensity spots being the stress concentrations at the
401 crack tips. In the damage-time maps showing cracks in Figure 8, these events occur along straight lines
402 parallel to fibres running in the ply beneath the observed ply. This suggests these matrix cracks are
403 initiated from a larger, more continuous crack running beneath the surface. Towards the end of the
404 second stage of damage propagation, the rate of damage events increases and the damage-time maps
405 showing delaminations start to indicate changes in the TSA signal spread across a large area. This is

406 the start of the final stage of damage propagation where the near-surface crack maps indicate
407 extensive crack propagation and the damage-time maps shown in Figure 9 indicate delaminations
408 forming and propagating. It is possible that if a higher data capture rate were used then the time
409 period, h, could be reduced such that individual cracking events would still be discernible during this
410 later stage using the near-surface crack maps. These observed stages correspond well to the typical
411 stages of fatigue growth reported in other studies [3, 5, 6]. The key benefits of the approach described
412 in this study compared to others is that it can be applied to opaque laminates [8] and provides good
413 spatial and temporal indications of individual crack and delamination initiation and propagation.
414 Performing and interpreting TSA on composite laminates is known to be challenging as the nature of
415 heat conduction between plies affects the measurement of the thermoelastic signal [26]. This varies
416 with the layup of the laminate, the thickness of the resin rich layer near the surface of the laminate
417 and the thickness of any paint on the surface [18], which results in uncertainties about the depth of
418 measurement in the laminate of the stresses. The algorithm developed in this study uses simple
419 statistics and image processing algorithms to compare R maps with earlier maps, these uncertainties
420 are unlikely to cause issues as the algorithm is only considering relative changes as opposed to
421 absolute values. However, care must be taken when interpreting the damage-time maps for different
422 layups in order to determine the depth of the detected damage.

423 The use of simple image processing techniques enables the algorithm to be computationally fast, with
424 each R map taking milliseconds to process, and thus could be incorporated into software for
425 performing TSA such that real-time information about the state of a specimen undergoing testing
426 could be obtained. This could help inform decisions about when tests should be terminated so that
427 damage can be inspected using higher resolution imaging techniques, for example computed
428 tomography. Whilst damage information can be manually observed in R maps, the benefit of this
429 algorithm is that the process is automated removing subjectivity from decision processes and freeing
430 investigator time. It also makes it much easier to quickly identify subtle changes resulting from the
431 formation and growth of cracks that are just a few millimetres long. The techniques also provide
432 information about the fatigue behaviour of the materials, with the same progression of failure
433 observed for all the specimens in this study suggesting it is inherent to the layup as opposed to the
434 geometry of the specimen. This implies that a single specimen design could be used to obtain the
435 necessary information to understand damage propagation at the mesoscale within a larger composite
436 structure. The level of detail observable in the near-surface crack maps also suggest that empirical
437 models of micro- and meso-scale damage propagation in composite materials could be constructed
438 using data from these maps. These models could be formed by fitting curves to the area of cracked
439 material curves shown in Figure 11, which have a close to constant rate of growth for most of the

fatigue life of a specimen. Final failure also appeared to occur shortly after the area of cracked material had passed a critical value, in the case of these tests the average final cracked material area was 66.2 mm^2 (with a standard deviation of 13.0 mm^2) suggesting a failure criterion could be defined. Previously, most information about damage processes during fatigue have been obtained by post-test inspection using destructive or non-destructive techniques, or by inspecting with visible light during the test to observe surface damage or sub-surface damage in translucent materials. The TSA technique combined with the algorithm described in this study provides more information about damage progression than could be obtained using other approaches and thus can help reduce the conservativeness of designs of composite structures.

Conclusions

A novel approach for obtaining data-rich information about damage propagation within composite material laminates has been developed. This approach is computationally efficient and uses as its input thermoelastic stress analysis (TSA), a well-established experimental stress analysis technique. The approach is based on an algorithm with two main strands, the first processes the TSA data such that sub-millimetre crack initiation away from the specimen edges, and the growth of cracks can be clearly detected and observed, yielding greater insight into the fatigue processes within composite materials than using existing techniques. This resulted in maps showing where and when matrix cracks formed within the laminates. The rate of crack initiation and propagation during the test was found to be similar across the specimens. The second strand of the algorithm focuses on changes spread across larger areas of the specimens and thus detects the formation of delaminations. To demonstrate the capabilities of the algorithm it was applied to nine specimens of three different geometries. The combined information provided by the algorithm yielded larger amounts of information about damage morphology as a function of time than could be obtained from purely optical observations. It also provides a greater understanding of fatigue propagation than can be obtained by microscopy or post-test nano-computed tomography. By improving the understanding of how material systems degrade, it is easier to build reliable fatigue models and thus predict the lifespan of composite structures.

Data accessibility

The experimental data and scripts can be found on Mendeley Data:

https://data.mendeley.com/preview/fv5yzvhw5c?a=bd3d0b1a-b5d3-47d0-b448-5d087ff87d7b

Author contributions

WJRC supervised the project, processed the data and wrote the first draft of the manuscript. PL performed the experiments. IET captured post-failure data and performed initial data processing. CAM, CP, and EAP supervised the work. All authors contributed to the final manuscript.

Competing Interests

We have no competing interests.

Funding

This effort was sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF under grant number FA8655-20-1-7034. The U.S. Government is authorised to reproduce and distribute reprints of Governmental purpose notwithstanding any copyright notation thereon. Maj. David D. Swanson (EOARD) was the program officer for this grant.

Some of the experimental data used in this study was captured as part of the UK Research Centre in Non-Destructive Evaluation (RCNDE) funded project, SteADfASt: Strain-based Assessment of Defects for Aviation Structures.

References

[1] Maragoni L, Carraro PA, Peron M, Quaresimin M. Fatigue behaviour of glass/epoxy laminates in the presence of voids. Int J Fatigue 2017; 95:18-28.

[2] Rajpal D, Mitrotta FMA, Socci CA, Sodja J, Kassapoglou C, De Breuker R. Design and testing of aeroelastically tailored composite wing under fatigue and gust loading including effect of fatigue on aeroelastic performance. Compos Struct 2021; 275:114373.

[3] Zimmermann N, Wang PH. A review of failure modes and fracture analysis of aircraft composite materials. Eng Fail Anal 2020; 115:104692.

[4] Ruggles-Wrenn MB, Lee MD. Fatigue behavior of an advanced SiC/SiC ceramic composite with a self-healing matrix at 1300°C in air and in steam. Mater Sci Eng 2016. 677:438-445.

[5] Alam P, Mamalis D, Robert C, Floreani C, Ó Brádaigh CM. The fatigue of carbon fibre reinforced plastics - A review. Compos Part B 2019; 166:555-579.

[6] De Finis R, Palumbo D, Galietti U. Evaluation of damage in composites by using thermoelastic stress analysis: A promising technique to assess the stiffness degradation. Fatigue Fract Eng M 2020; 43(9):2085-2100.

[7] Gao J, Zhu P, Yuan Y, Wu Z, Xu R. Strength and stiffness degradation modeling and fatigue life prediction of composite materials based on a unified fatigue damage model. Eng Fail Anal 2022; 137:106290.

[8] Glud JA, Dulieu-Barton JM, Thomsen OT, Overgaard LCT. Fatigue damage evolution in GFRP laminates with constrained off-axis plies. Compos Part A 2017; 95:359-369.

[9] Boni L, Fanteria D, Lazzeri L, Mariani U, Rigamonti M. Delamination onset in composite materials due to fatigue loading. J Compos Mater 2022; 56(16):2585-2598.

[10] Yu B, Bradley RS, Soutis C, Hogg PJ, Withers PJ. 2D and 3D imaging of fatigue failure mechanisms of 3D woven composites. Compos Part A; 77:37-49.

[11] Battams GP, Dulieu-Barton JM. Data-rich characterisation of damage propagation in composite materials. Compos Part A 2016; 91:420-435.

[12] Dai S, Cunningham PR, Marshall S, Silva C. Open hole quasi-static and fatigue characterisation of 3D woven composites. Compos Struct 2015; 131:765-774.

[13] Greene RJ, Patterson EA, Rowlands RE. Thermoelastic Stress Analysis. In: Sharpe WN, editor. Springer Handbook of Experimental Solid Mechanics. Boston: Springer, 2008, p.743-768.

[14] Díaz FA, Patterson EA, Tomlinson RA, Yates JR. Measuring stress intensity factors during fatigue crack growth using thermoelasticity. Fatigue Fract Eng M 2004; 27(7):571-583.

[15] Middleton CA, Gaio A, Greene RJ, Patterson EA. Towards Automated Tracking of Initiation and Propagation of Cracks in Aluminium Alloy Coupons Using Thermoelastic Stress Analysis. J Nondestruct Eval 2019; 38(1):18.

[16] Emery TR, Dulieu-Barton JM. Thermoelastic Stress Analysis of damage mechanisms in composite materials. Compos Part A 2010; 41(12):1729-1742.

[17] Pitarresi G, Scalici T, Catalanotti G. Infrared Thermography assisted evaluation of static and fatigue Mode II fracture toughness in FRP composites. Compos Struct 2019; 226:111220.

[18] Jiménez-Fortunato I, Bull DJ, Thomsen OT, Dulieu-Barton JM. On the source of the thermoelastic response from orthotropic fibre reinforced composite laminates. Compos Part A 2021; 149:106515.

[19] Quinlan A, Castro O, Dulieu-Barton JM. Towards assessment of fatigue damage in composite laminates using thermoelastic stress analysis. Compos Part C 2023. 12:100377

[20] Christian WJR, Dvurecenska K, Amjad K, Pierce J, Przybyla C, Patterson EA. Real-time quantification of damage in structural materials during mechanical testing. Roy Soc Open Sci 2020; 7(3):191407.

[21] Otsu N. A Threshold Selection Method from Gray-Level Histograms. IEEE T Syst Man Cyb 1979; 9(1):62-66.

[22] Christian WJR, Dean AD, Dvurecenska K, Middleton CA, Patterson EA. Comparing full-field data from structural components with complicated geometries. Roy Soc Open Sci 2021; 8(9):210916.

[23] Middleton CA, Weihrauch M, Christian WJR, Greene RJ, Patterson EA. Detection and tracking of cracks based on thermoelastic stress analysis. Roy Soc Open Sci 2020; 7(12):200823.

[24] Huang J, Pastor ML, Garnier C, Gong XJ. A new model for fatigue life prediction based on infrared thermography and degradation process for CFRP composite laminates. Int J Fatigue 2019; 120:87-95.

[25] Quaresimin M, Carraro PA, Mikkelsen LP, Lucato N, Vivian L, Brøndsted P, Sørensen BF, Varna J, Talreja R. Damage evolution under cyclic multiaxial stress state: A comparative analysis between glass/epoxy laminates and tubes. Compos Part B 2014; 61:282-290.

[26] Wong AK. A non-adiabatic thermoelastic theory for composite laminates. J Phys Chem Solids 1991; 52(3):483-494.

Figures

Figure 1: A typical 40 mm wide carbon-fibre composite laminate test-specimen with a 45° notch machined at its centre and additional tabs bonded at each end.

Figure 2: Photographs of 45° notch (top), -45° notch (middle) and 0° notch (bottom) specimens of the geometry and type shown in Figure 1, following fatigue testing. As the -45° notch specimens were monitored on their reverse side they have a similar appearance to the 45° notch specimens.

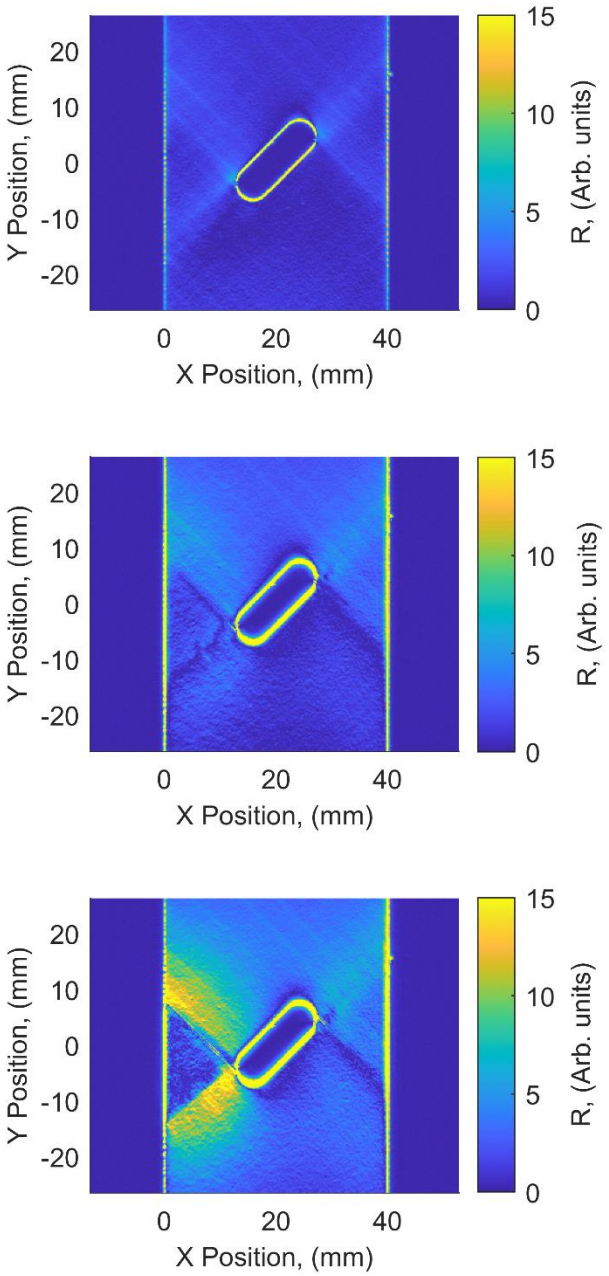

561

562 *Figure 3: Example R maps from a specimen with a 45° notch at 5% (top), 50% (middle) and 95%*
563 *(bottom) of its fatigue life. The units of R are arbitrary (Arb.) measurement system units.*

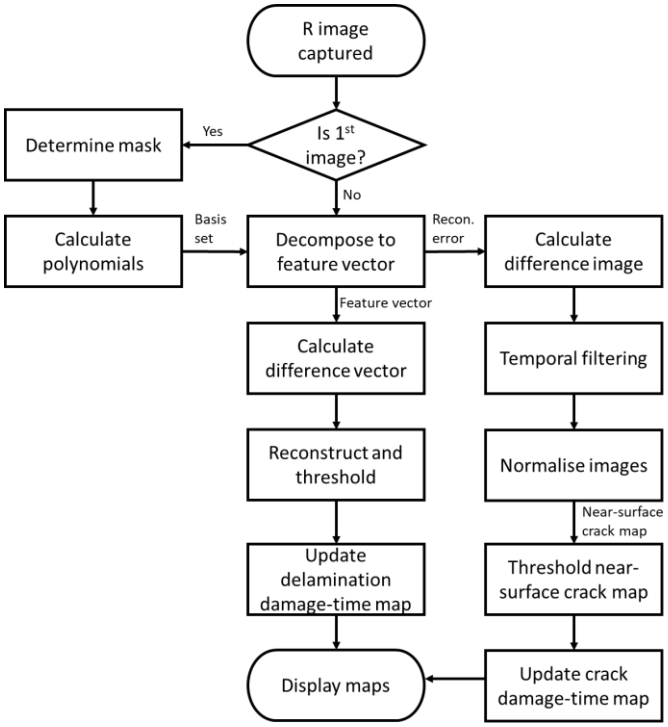

564

565 *Figure 4: Flow chart showing the damage monitoring algorithm.*

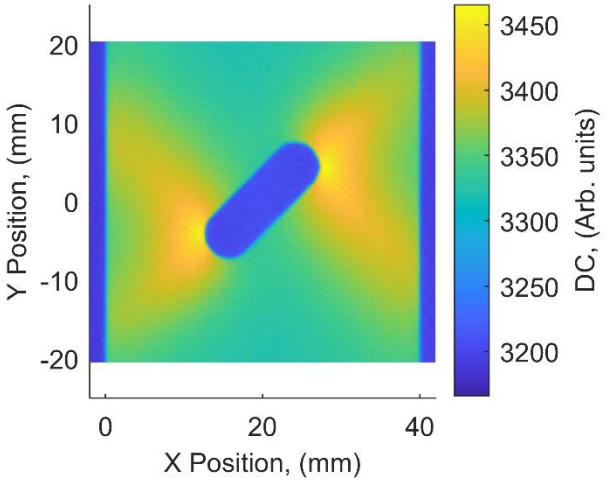

566

567 *Figure 5: Raw DC map from a specimen with a -45° notch near the start of a test.*

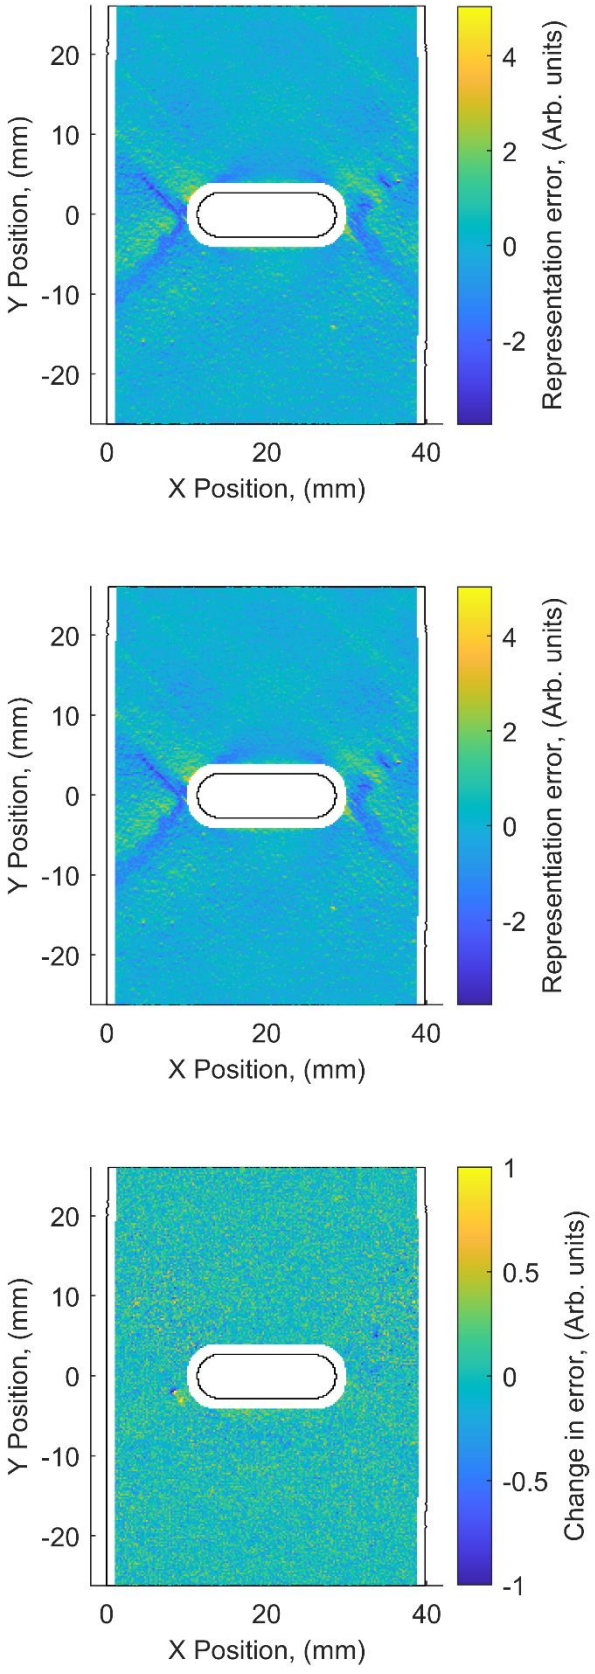

Figure 6: Reconstruction error maps for a specimen from just before (top) and just after (middle) a damage event and the difference between the two maps (bottom).

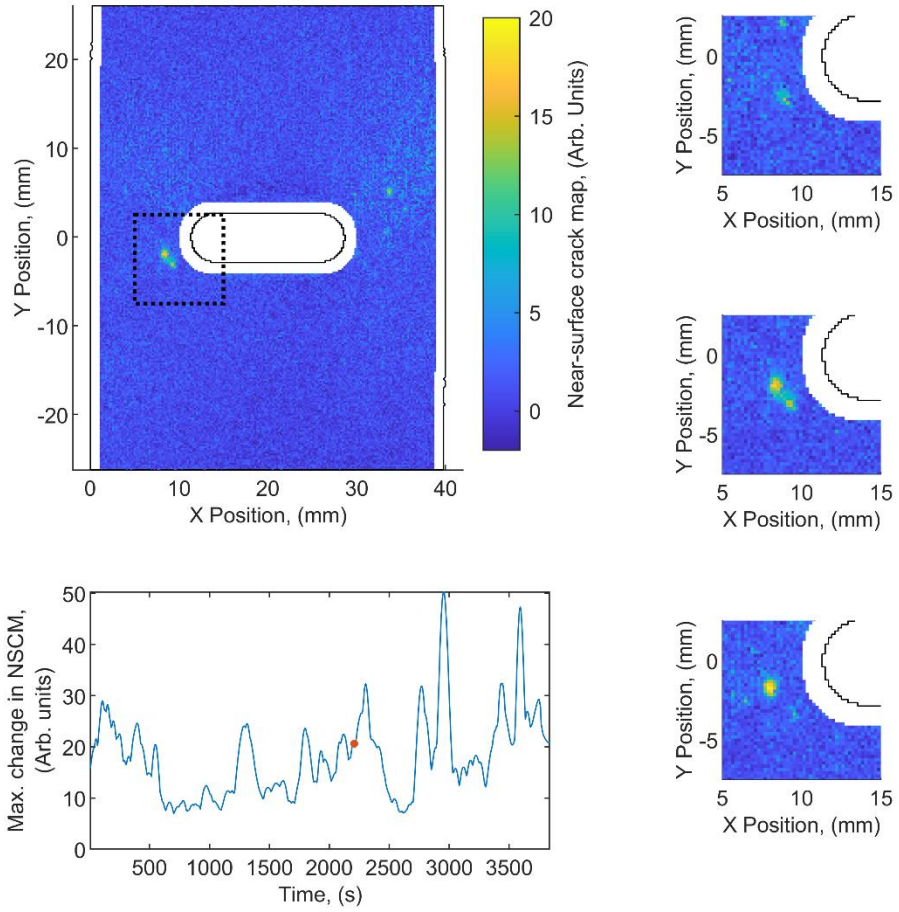

Figure 7: Near-surface crack map (top-left) at time of a matrix crack forming on the top surface of a specimen. The insets show an enlarged version of the near-surface crack map at time, t (middle-right), 90 seconds earlier, i.e., t - 90 (top-right) and 90 seconds later, i.e., t + 90 (bottom-right). The maximum change in the near-surface crack map (NSCM) as a function of time is shown (bottom-left) with a red dot indicating the time, t.

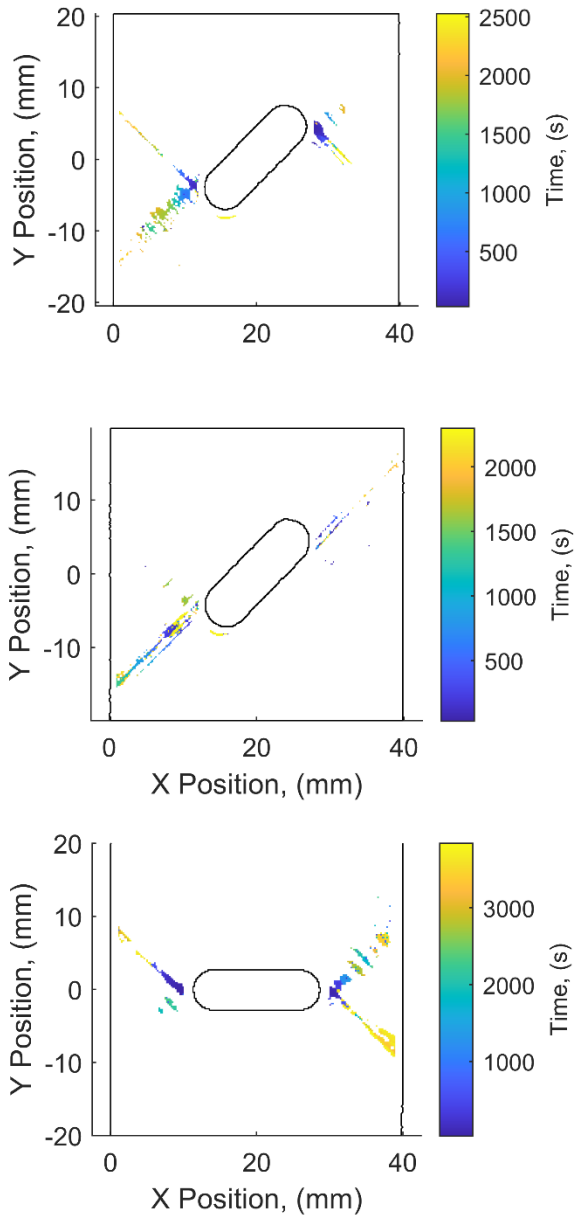

578

579 *Figure 8: Damage-time maps showing near-surface crack growth as a function of time for the*
580 *specimens shown in Figure 2 with a 45° notch (top), -45° notch (middle) and 0° notch (bottom).*

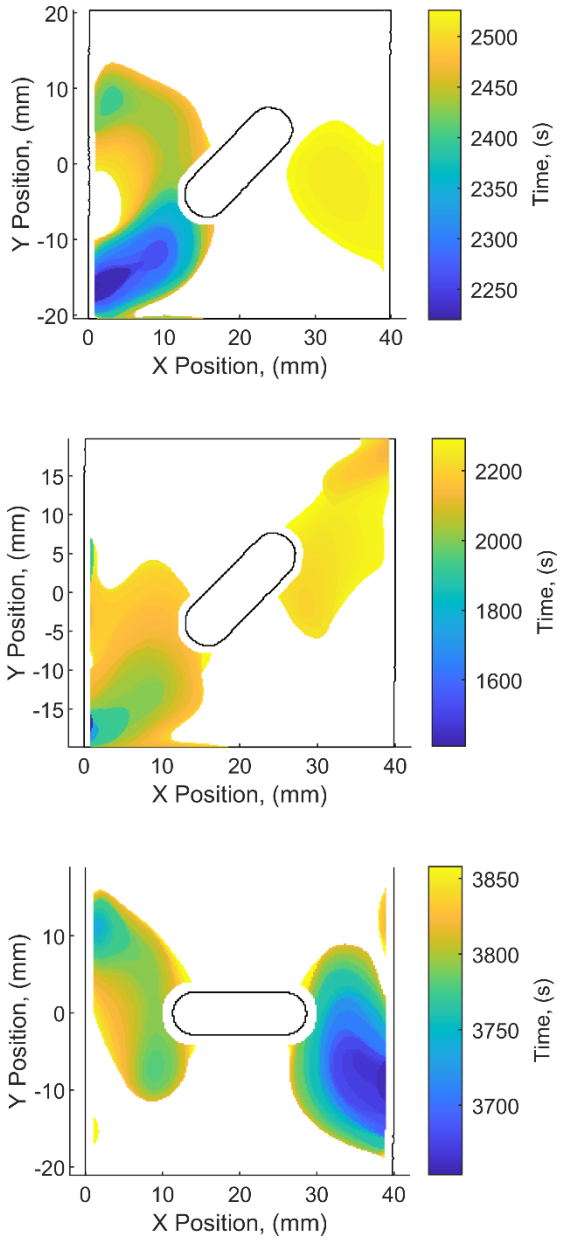

581

582 *Figure 9: Damage-time maps showing delamination propagation as a function of time for the*
583 *specimens shown in Figure 2 with a 45° notch (top), -45° notch (middle) and 0° notch (bottom).*

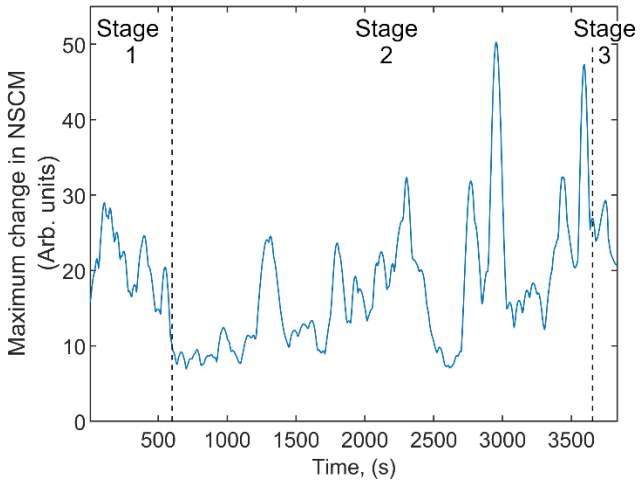

584

585 *Figure 10: The maximum change in the near-surface crack map as a function of time for a specimen*
586 *with a 0° notch showing peaks when cracks were formed. The dashed lines indicate the three stages of*
587 *defect formation during fatigue tests and are identified based on observations from Figures 8 and 9.*

588

589

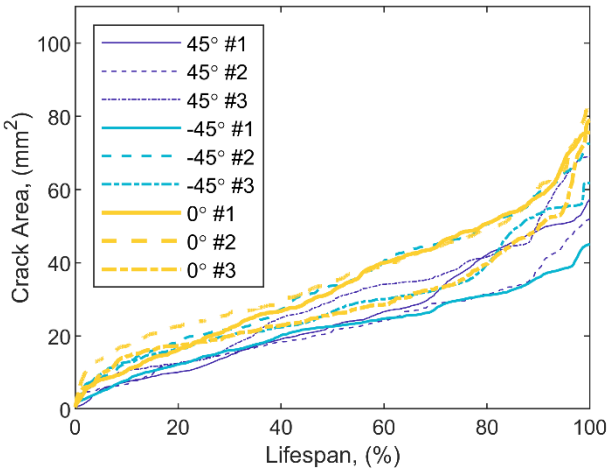

590

591 *Figure 11: Total area of crack damage as a function of time normalised by lifespan from crack damage-*
592 *time maps. The solid lines correspond to the specimens used to produce the crack damage-time maps*
593 *shown in Figure 8.*