

A Model-Based, Bayesian Solution for Characterization of Complex Damage Scenarios in Aerospace Composite Structures

H. Reed^{a,*}, Cara A.C. Leckey^b, A. Dick^a, G. Harvey^c, J. Dobson^{d,e}

^aThornton Tomasetti - Weidinger Applied Science, 40 Wall St. 19th Fl., New York, NY 10005

^bNASA Langley Research Center, Hampton VA 23681

^cPZFlex, 19200 Stevens Creek Blvd, Suite 100, Cupertino, CA 95014

^dPZFlex, 6th Floor South Suite, 39 St Vincent Place, Glasgow, UK, G1 2ER

^eCentre for Ultrasonic Engineering, University of Strathclyde, Glasgow, UK, G1 1XW

Abstract

Ultrasonic damage detection and characterization is commonly used in nondestructive evaluation (NDE) of aerospace composite components. In recent years there has been an increased development of guided wave based methods. In real materials and structures, these dispersive waves result in complicated behavior in the presence of complex damage scenarios. Model-based characterization methods utilize accurate three dimensional finite element models (FEMs) of guided wave interaction with realistic damage scenarios to aid in defect identification and classification. This work describes an inverse solution for realistic composite damage characterization by comparing the wavenumber-frequency spectra of experimental and simulated ultrasonic inspections. The composite laminate material properties are first verified through a Bayesian solution (Markov chain Monte Carlo), enabling uncertainty quantification surrounding the characterization. A study is undertaken to assess the efficacy of the proposed damage model and comparative metrics between the experimental and simulated output. The FEM is then parameterized with a damage model capable of describing the typical complex damage created by impact events in composites. The damage is characterized through a transdimensional Markov chain Monte Carlo solution, enabling a flexible damage model capable of adapting to the complex damage geometry investigated here. The posterior probability distributions of the individual delamination petals as well as the overall envelope of the damage site are determined.

Keywords: Nondestructive evaluation; Guided waves; Composites; Delamination; Model-based characterization; Bayesian analysis

1. Introduction

In the past decade a number of industries, including aerospace and automotive, have increased the use of lightweight composite materials for structural and non-structural components [1]. Carbon fiber reinforced polymer (CFRP) composites are extensively used in aerospace components. This expanded use of composites leads to unique nondestructive evaluation (NDE) challenges since defects which are specific to layered composites (i.e., did not exist in heritage metallic parts) can occur during the manufacture and curing of components. Additionally, damage to composites can occur post-curing in a manufacturing or in-service setting. Common defects/damage in CFRP components include delaminations, fiber waviness (marcelling) and wrinkling, microcracking and porosity [2, 3, 4]. Delamination damage due to impact and/or in-service loading can occur as a complex three-dimensional (3D) volumetric geometry with individual delaminations existing at different ply depths through the composite thickness. The defects that commonly occur in CFRPs may lead to no visible indication on the composite surface, thus appropriate NDE methods are required for defect detection and characterization.

Ultrasonic inspection is one of the most widely used NDE approaches applied to composites in both manufacturing and in-service settings. Traditional ultrasonic inspection approaches commonly applied to composites

include pulse-echo and through-transmission scanning. Guided wave methods are also highly applicable for inspection of plate-like composite laminates [5, 6, 7, 8]. Simulation of ultrasonic wave behavior in composites can significantly aid in the development and application of damage characterization methods. The simulation of the wave interaction with materials/defects is commonly referred to as the forward modeling problem. The inverse modeling problem entails the use of models/simulation to estimate the defect characteristics which caused a measured experimental signal, and involves making comparisons between simulated and experimental data. NASA has unique needs for light-weight advanced materials for aerospace applications, including manned space-flight applications which have strict safety, durability, and reliability requirements. Advancing the state-of-art in forward and inverse modeling for the purpose of improved defect detection and characterization in advanced materials (such as composites) is highly relevant to NASA’s aerospace missions.

In recent years there is a growing literature on the use of ultrasonic modeling tools (mostly Finite Element (FE) based tools) for simulation of the physics of wave propagation and scattering in composites. Recent authors reporting forward model simulation of ultrasound in composites includes Ng and colleagues who implemented 3D FE, using LS-DYNA as the FE solver, to study guided wave scattering from delaminations in a quasi-isotropic composite laminate [9]. Ng et al. assume homogeneous orthotropic material properties for each ply and accounted for the individual ply rotations of the layup [9]. They simulated scattering from simple circular geometry delaminations located at a single ply depth. Singh et al. report using COMSOL FE software to simulate guided waves in a composite laminate with homogenized material properties through the thickness (i.e., individual ply layers were not simulated). Guided wave interaction with a simulated conical shaped defect representing impact damage was studied [10]. Leckey et al. used a custom 3D finite integration code to study guided wave propagation in composite laminates (simulating each ply layer) and incorporated a realistic damage geometry using x-ray computed tomography data of impact induced delamination damage [11]. More recently, Murat et al. report using a custom FE code to simulate guided wave propagation in a cross-ply composite laminate and study guided wave interaction with a square shaped delamination [12].

There are significantly fewer reports in published literature addressing the ultrasonic inverse problem for defect characterization in composites. Roberts and Holland recently reported use of a Green’s function based analytical approach for model-based inversion of pulse-echo ultrasound data of delamination defects in composites laminates [13]. Fahim et al. considered a cost function (i.e. an optimization problem) as the inverse solution to characterize delamination damage features through the use of a semi-analytic model of the wave propagation forward problem [14]. Bochud et al. model the composite damage in each layer as a reduction of Young’s modulus and studied various signal processing techniques for use in an inverse problem [15].

In this work a finite element simulation is implemented to model guided wave propagation in composite laminates and wave interaction with delamination defects. A model-based inversion approach using Markov Chain Monte Carlo (MCMC) method is combined with the FE modeling and is applied to ultrasonic wavefield data from a composite laminate containing impact induced delamination damage. Prior to the application of the inversion method to the real delamination case, a material characterization study is performed. Numerous factors can lead to differences between simulation and experiment. For composite materials a primary factor can be differences between the ideal ply level material properties (provided by a manufacturer or found through destructive testing) and the properties of a cured as-manufactured specimen [11]. A comparison between the FE simulation and experimental data is performed for a pristine unidirectional laminate to more closely characterize the as-manufactured composite material properties. Section 2 describes the mathematical approach used in the forward and inverse modeling, along with the wavenumber domain method that is used to compare the simulation data to experimental wavefield data. Section 3.2 discusses the results of the material characterization step. In Section 3.3 the MCMC method is then applied to a single delamination case to verify that the approach works for a simple defect scenario. Application of the inverse model method to the wavefield data recorded in experiment for the complex delamination case is then reported in Section 3.4.

2. Problem description

2.1. Forward problem

The wave propagation problems investigated here are implemented in the commercial finite element software PZFlex [16]. The forward problem entails simulating wave propagation and interaction with delamination damage in a composite laminate. The domain of the problem corresponding to the delaminated quasi-isotropic specimen is given in figure 3. The composite laminate is comprised of 26 ply layers with a layup of $[(0/45/-45/90)_3]_s$. To

build up the entire composite laminate in the finite element model, the material tensor of a single ply is computed in a reference orientation (such as along the 0° fiber direction of the top ply shown in figure 3), and the material tensors of the ply layers at differing ply angles are computed by rotating the material tensor by the appropriate rotation transformation operator. In this work, we assume that a single ply layer containing fibers and resin can be accurately modeled as a single anisotropic material (with transverse anisotropy for the cases studied in this work). However, due to variations in fiber volume ratio, overall specimen thickness, and fiber waviness and warping (i.e. where fibers do not run along the intended direction), accurately predicting the homogenized anisotropic material properties for the ply layer is difficult. An approach for addressing as-manufactured laminate property variations from the idealized properties will be discussed in Section 3.2.

The composite laminate in this work is excited by a 12.7 mm diameter piezoelectric transducer. The response of the laminate to the excitation considered here is the out-of-plane velocity on the top surface of the composite laminate. While the entire displacement and stress time history (in-plane and out-of-plane) is modeled within the laminate in the FE simulation, we only consider the response that is captured in the experimental wavefield measurements. In this work, a 1D laser Doppler vibrometer (LDV) attached to an x-y scanning system is used to record the out-of-plane velocity $v(t, \mathbf{x})$, where \mathbf{x} is the in-plane spatial location on the laminate surface, and t is the time. The details on the inspection setup are provided in Section 3.1.

2.1.1. Three-dimensional Fourier transform

In a model-based inversion problem, the response of the simulation model is directly compared to the experimental inspection data. Therefore, we require a means of processing the recorded velocity time signals such that the simulation data and inspection data are readily comparable. Directly comparing the velocity time signals $v(t, \mathbf{x})$ (i.e., A-scan signals) between simulation and experiment is non-ideal since a number of factors would make a quantitative comparison challenging. Factors influencing a direct velocity time signal (A-scan) comparison include: the LDV pixel locations and the finite element nodal coordinates would need to align; the LDV sampling rate and the finite element time steps would need to be coordinated so that the data is compared at appropriate time steps; and significant effort would need to be undertaken to ensure that excitation amplitude is correctly simulated so that the model response amplitude is the same as that generated by the transducer used in the experiment (including any amplitude variations due to the couplant) and recorded by the LDV, $v^{LDV}(t, \mathbf{x})$.

To mitigate the challenges associated with comparing the velocity spatial-time signals directly, we implement a post-processing method that involves wavenumber comparisons between the inspection data and simulated data. The wavenumber analysis is implemented through the 3D Fourier transformation of the time domain wavefield data to generate frequency versus \hat{x} wavenumber, k_x , and \hat{y} wavenumber, k_y .

$$V(f, \mathbf{k}) = \int_0^{\mathbf{x}} \int_0^t v(t, \mathbf{x}) e^{-j(2\pi ft - \mathbf{k} \cdot \mathbf{x})} dt d\mathbf{x} \quad (1)$$

Two-dimensional wavenumber spectra ($k_x - k_y$) are generated over a range of frequencies f . The wavenumber spectrum associated with the excitation frequency is the metric to be used in comparing the simulated and inspection responses.

2.2. Stochastic inverse problem

We assume that the spatio-temporal wavefield data has been transformed via the 3D Fourier transformation and evaluated at frequency f and discrete wavenumbers $k_{x_i}, i = 1, \dots, N$ and $k_{y_j}, j = 1, \dots, M$. We further assume that only experimental noise recorded by the LDV contributes to discrepancies between the finite element response and the recorded inspection response, such that the finite element response may be written

$$V^{LDV} \equiv V_{ij}^{LDV} = V_{ij}^{FEM}(\theta) + \eta_{ij} \quad (2)$$

where V_{ij}^{LDV} is the 3D Fourier transform of the experimental data at wavenumbers k_{x_i} and k_{y_j} and excitation frequency f ; $V_{ij}^{FEM}(\theta)$ is the 3D Fourier transformation of the associated deterministic finite model output, as a function of the unknown and sought-after model parameters θ ; and η_{ij} is the independent and identically distributed (*i.i.d.*) errors taken from a zero-mean, Gaussian distribution with variance σ_η^2 . Given this assumed noise distribution, we may model the probability of observing a wavenumber spectrum amplitude value V_{ij}^{LDV} , given the measured finite element model response $V_{ij}^{FEM}(\theta)$ as:

$$f_{\eta}(V_{ij}^{LDV} | V_{ij}^{FEM}(\theta)) = \frac{1}{\sqrt{2\pi\sigma_{\eta}^2}} e^{-\frac{1}{2\sigma_{\eta}^2}(V_{ij}^{LDV} - V_{ij}^{FEM}(\theta))^2} \quad (3)$$

Eq. (3) is the probability density function (PDF) for each amplitude value in the wavenumber spectrum; and clearly relates the unknown model parameters θ to the inspection data.

Because the noise is assumed i.i.d. and stationary, we may express the PDF for the *entire* wavenumber spectrum as the product of the individual (marginal) distributions of each wavenumber point:

$$\begin{aligned} f_{\eta}(V_{ij}^{LDV} | V_{ij}^{FEM}(\theta)) &= \prod_i^N \prod_j^M \frac{1}{\sqrt{2\pi\sigma_{\eta}^2}} e^{-\frac{1}{2\sigma_{\eta}^2}(V_{ij}^{LDV} - V_{ij}^{FEM}(\theta))^2} \\ &= \frac{1}{(2\pi\sigma_{\eta}^2)^{N \times M/2}} \exp \left[-\frac{1}{2\sigma_{\eta}^2} \sum_{i=1}^N \sum_{j=1}^M (V_{ij}^{LDV} - V_{ij}^{FEM}(\theta))^2 \right] \equiv f_{\eta}(V^{LDV} | \theta). \end{aligned} \quad (4)$$

Eq. (4) is a probabilistic model, predicting the wavenumber spectrum, given model parameters θ , and is referred to as the *likelihood* of the observations. A maximum likelihood estimation (MLE) therefore provides the values θ that maximize Eq. (4). However, rather than focusing on a single “best” value (in the sense of maximizing the likelihood), we treat each model parameter value as a random variable, modeled with a corresponding PDF. Estimates of this PDF can therefore be used to provide a different definition of “best” parameter value (e.g., mean or mode of the PDF). More importantly, by estimating the entire PDF of the parameter we can quantify the uncertainty in the parameter value.

The above described approach (i.e. treat parameters as random variables) is a hallmark of Bayesian analysis. At the core of this analysis is Bayes’ Theorem which allows us to express the desired *posterior* PDF for the parameter vector θ in terms of the likelihood Eq. (4) as

$$f_{\theta}(\theta | V^{LDV}) = C^{-1} f_{\eta}(V_{ij}^{LDV} | \theta) f_{\pi}(\theta) \quad (5)$$

where $f_{\pi}(\theta)$ is the *prior* PDF, which encapsulates any prior information about the parameters θ ; and C^{-1} is defined as the multidimensional integral $\int_{\mathbb{R}^P} f_{\eta}(V_{ij}^{LDV} | \theta) f_{\pi}(\theta) d\theta$, ensuring that the posterior integrates to unity (a requirement of any PDF). As it is that we are interested in the *marginal* distribution associated with each model parameter $\theta_l \in \theta$, $l = 1, \dots, P$, we are required to integrate Eq. (5) over each of the parameters in θ , excluding θ_l . If we define the parameter vector that excludes parameter θ_l as θ_{-l} , then the expression for the marginal posterior distribution for parameter θ_l is given as:

$$f_{\theta}(\theta_l) = \int_{\mathbb{R}^{P-1}} f_{\theta}(\theta | V^{LDV}) d\theta_{-l} \quad (6)$$

Equation (6) is the end goal of this Bayesian analysis; providing a complete probabilistic description of our unknown model parameters. The problem is that Eq. (6) is analytically intractable in most situations and instead must be solved by numerical methods. The numerical approach used in this work is discussed in section (2.4).

2.3. Unknown model parameters

As the end goal of this analysis is to solve Equation (6) to characterize the damage, we must formulate and implement a parameterization model for the damage. Because the laminate damage in this work is a complex site of several delaminations, we will allow the delamination model θ to be suitably complex through the use of a flexible model that will allow the dimension of θ to change.

Prior to attempting to characterize the damage, however, we must first ensure that the model used to generate the measured data accurately captures the mechanics of the physical inspection (i.e. we need to validate the model). This is crucial in model-based inversion methods, as the FE model results are to be directly compared to the experimentally acquired data. Therefore, prior to attempting to characterize the damage, we will first use the model-based inversion approach to characterize the homogenized laminate material properties. In this initial material characterization validation step, the model parameter vector θ_{mat} will take on values of the material properties as described in Section (2.3.1) below.

2.3.1. Laminate material model

Fundamental to model-based inversion methods is that the mechanics of the simulation model must accurately capture that of the physical inspection. If the FE simulation does not accurately represent the physics of the inspection, then there is little hope of using the high fidelity FE model to invert the inspection data to characterize damage. As discussed in section (2.1), one of the reasons that inspection data and simulated data (even in a case with no damage) may not agree is due to using inaccurate homogenized ply level material properties. In addition to fiber waviness, warping, porosity, resin rich regions and prepreg material property changes due to material aging or out-of-freezer time prior to manufacture; the cured laminate material properties may also change over time due to aging and exposure to the environment. For this reason, before attempting to characterize the complex delamination site, we first employ the model-based inversion approach to characterize the ply-level anisotropic material tensor.

For the CFRP material used in our studies, we assume that each ply layer can be modeled as a transversely isotropic material. For ply layers oriented at angles other than the reference angle (0° defined as fiber direction in the top ply, as shown in Figure 3), the material tensors are obtained by applying a rotation transformation to the reference material tensor. Therefore, the number of sought-after material properties are assumed to depend only on the unique material properties in a single layer, and are not proportional to the number of ply layers. The material properties for a single ply layer of IM7/8552 reported in literature are given in Table (1) [17, 18]. As can be seen in the table, there are seven unique material properties due to the material symmetry that can be used to populate the material tensor that describes the ply layer and comprise parameter vector θ_{mat} . An initial material property sensitivity study will be discussed in Section 3.2 to determine whether all seven of the unique material properties need to be included in θ_{mat} , as unnecessarily including extra parameters to solve for increases the inversion and characterization difficulty. Thus, special care should be given to determining which parameters are included.

Table 1: IM7/8552 material properties (for a single ply, with E_1 in the fiber direction)

Property	IM7/8552
ρ (kg/m^3)	1570
E_1 (GPa)	171.4
E_2 (GPa)	9.08
E_3 (GPa)	9.08
G_{12} (GPa)	5.29
G_{13} (GPa)	5.29
G_{23} (GPa)	2.80
ν_{12} (GPa)	0.32
ν_{13} (GPa)	0.32
ν_{23} (GPa)	0.5

2.3.2. Delamination model

As will be discussed further in (3.1), the goal is to characterize impact-induced delamination damage. Micro-focus computed tomography (microCT) scans of the damaged quasi-isotropic laminate were taken at a resolution of $23.4 \mu m$, and can be seen in Fig. (1), where red indicates delaminated regions. Because the microCT images indicate that the delaminated regions are relatively open, we model a delamination in the FE model through stress-free boundary conditions between ply layers in the regions that have been identified as delaminated. The stress-free boundary condition is effected through de-bonding the ply layers in the regions of delaminations. Additionally, in order to simplify the delamination implementation in the model, we assume that each individual delamination petal is a perfect circle. This choice is made in order to balance accuracy with as simple of a damage model as possible. An initial sensitivity study (for a single delamination) is described in Section 3.3 to illustrate how various delamination parameters effect the wavenumber spectrum that is to be used to evaluate the posterior distributions.

The parameters that describe a delaminated region are: the interface number between layers l ; the in-plane location of the center of the delamination circle c_x and c_y ; and the radius of the delamination petal r (see Figure 2 for an example). The stress-free boundary condition is indicated by the function Γ ,




Figure 1: Delaminations images created by segmenting microfocus x-ray CT data of the damaged composite laminate. Red regions are the open delaminations, while blue regions are pristine (non-delaminated) composite material. The through-thickness locations of the images from the bottom of the laminate are: (a) - 2.72 μm ; (b) - 1.94 μm ; (c) - 1.69 μm

$$\Gamma(\mathbf{x}; \theta_d) = \sum_{j=1}^{n_k} \begin{cases} 1, & \text{if } \|\{\mathbf{c}_j, l_j\} - \mathbf{x}\| \leq r_j \\ 0, & \text{else} \end{cases} \quad (7)$$

where n_k is the number of delamination petals in the laminate. If $\Gamma(\mathbf{x}; \theta_d) \geq 1$, then a stress-free (i.e. delamination) boundary condition exists at layer l and location \mathbf{x} , with radius r . As evidenced in Fig. (1), the complex damage pattern contains several delamination petals of different locations and sizes. It is not expected that the exact number of petals n_k is known *a priori*. Therefore, we allow n_k to be an additional unknown parameter. This enables the inversion method to automatically determine the number of delamination petals as well as their locations and size, that best agree with the measured inspection data.

Utilizing a damage model in which the number of delamination sites is not known *a priori* enables flexibility in the damage characterization, as the inversion method can estimate the number of delaminations. This flexibility, however, comes at the price of adding model complexity to the problem. Adding complexity to the damage model can allow for a more refined characterization of the delamination, but it also increases the difficulty in fitting the parameters (the “curse of dimensionality”). Additionally (as will be described in (2.5), the ability to increase or decrease the number of delamination petals n_k in the inversion scheme must be addressed in the algorithm, as we require a means of comparing the posterior probability distributions of two damage parameter vectors of different dimensions.

2.4. MCMC sampling

Evaluating the conditional posterior in Eq. (5) is difficult, in that it requires evaluating a multi-dimensional integral. Rather than analytically solving for the posterior, it has been shown that one may numerically draw samples from the posterior PDF [19]. More specifically, the Markov chain Monte Carlo (MCMC) algorithm builds a first order Markov chain such that the resulting sequence of values $\theta_l(i), i = 1 \dots N$ are, in fact, samples from the desired posterior distribution (i.e., each $\theta_l(i) \sim f_\theta(\theta_l)$). Variants of the MCMC algorithm abound in the statistics literature, and we describe one such implementation (the Metropolis-Hastings algorithm [19]) here.

Taking a sample from the prior PDF as the starting value, (e.g., $\theta_l(0) \sim f_\pi(\theta_l)$), the algorithm generates a candidate delamination petal θ_l^* via a probabilistic model $q(\theta_l^*|\theta_l(i))$ which specifies the probability of attaining petal parameters θ_l^* given the current value $\theta_l(i)$. Specific choice of $q(\cdot, \cdot)$ is discussed in [20]. Both current and candidate delamination petal values are then used to evaluate the ratio r :

$$r = \frac{f_\eta(\mathbf{y}|\theta_l^*)}{f_\eta(\mathbf{y}|\theta_l(i))} \times \frac{f_\pi(\theta_l^*)}{f_\pi(\theta_l(i))} \times \frac{q(\theta_l(i)|\theta_l^*)}{q(\theta_l^*|\theta_l(i))} \quad (8)$$

which is simply the ratio of posterior PDFs (evaluated using the finite element model) multiplied by the ratio of the candidate generating PDFs $q(\cdot|\cdot)$. The value θ_l^* is then retained as $\theta_l(i+1)$ with probability $P_r = \min(r, 1)$, otherwise $\theta_l(i+1) = \theta_l(i)$. While the above-described approach is theoretically sound, there are several choices that need to be made in a practical implementation. The reader is referred to [20] for details on the practical aspects of the implementation.

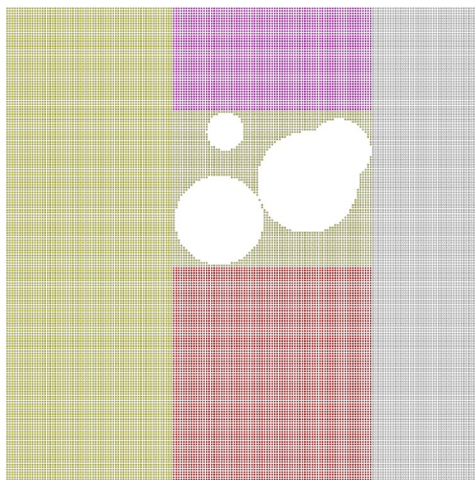


Figure 2: Potential delamination petals occurring in a single layer. A new delamination can emerge (birth); two delaminations could be merged or split apart; or a delamination can be removed (death). The flexible damage model enables complex delamination patterns to occur when delaminated regions overlap.

2.5. Reversible jump Markov chain Monte Carlo

The previous section describes a sampling method which is appropriate when the number of delamination petals is specified in the damage model *a priori*. However, we may also consider n_k to be a variable, in which case the number of parameters to estimate is equal to n_k . The sampling procedure for the case where the number of petals, n_k , varies may be treated using the reversible jump Markov chain Monte Carlo (RJMCMC) method described by Green [21]. RJMCMC differs from standard MCMC in that the former permits the dependent sampling of posterior distributions possessing different dimensions, in a way that preserves detailed balance within the method. Theoretical details can be found in the literature [21], [22].

The acceptance ratio given in Eq. (8) is augmented so that the appropriate acceptance probability is:

$$r = \frac{f_\eta(\mathbf{y}|\boldsymbol{\theta}^*)}{f_\eta(\mathbf{y}|\boldsymbol{\theta})} \frac{f_\pi(\boldsymbol{\theta}^*)}{f_\pi(\boldsymbol{\theta})} \frac{p(n_k \rightarrow n_k^*)}{p(n_k^* \rightarrow n_k)} \frac{1}{g(\mathbf{u})} \left| \frac{\partial \boldsymbol{\theta}^*}{\partial(\boldsymbol{\theta}, \mathbf{u})} \right|, \quad (9)$$

The acceptance ratio of the reverse, dimension-decreasing move is given by:

$$r = \frac{f_\eta(\mathbf{y}|\boldsymbol{\theta})}{f_\eta(\mathbf{y}|\boldsymbol{\theta}^*)} \frac{f_\pi(\boldsymbol{\theta})}{f_\pi(\boldsymbol{\theta}^*)} \frac{p(n_k^* \rightarrow n_k)}{p(n_k \rightarrow n_k^*)} g(\mathbf{u}) \left| \frac{\partial \boldsymbol{\theta}^*}{\partial(\boldsymbol{\theta}, \mathbf{u})} \right|^{-1} \quad (10)$$

In Eqs. (9) and (10), $f_\eta(\cdot)$ is the likelihood, defined as in Eq. (4).

In addition to the usual Metropolis-Hastings move described in Sec. (2.4), in which a candidate parameter value is uniformly sampled around the original value, we may also consider moves that change the number of delamination petals. To ensure reversibility, each move that changes the value of n_k must have a counter move, capable of restoring the proposed parameter vector ($\boldsymbol{\theta}^*$) back to the original parameter vector ($\boldsymbol{\theta}$). This is done through special operations resulting in the creation or deletion of a petal (birth/death); the splitting of a petal into two petals; or merging two petals into one petal (split/merge). The flexible defect model enables complex delamination patterns to emerge when delaminated regions overlap. The reader is referred to [20] and [23] for practical implementations on the RJMCMC algorithm.

3. Results

3.1. Experimental setup

A Polytec OFV-505 laser Doppler vibrometer (LDV) was used to collect relevant experimental data for the undamaged laminate plate comparisons. The LDV is attached to a two-axis scanning system to record out-of-plane velocity on a Cartesian grid. The out-of-plane velocity signal recorded by the LDV was digitized by a 14

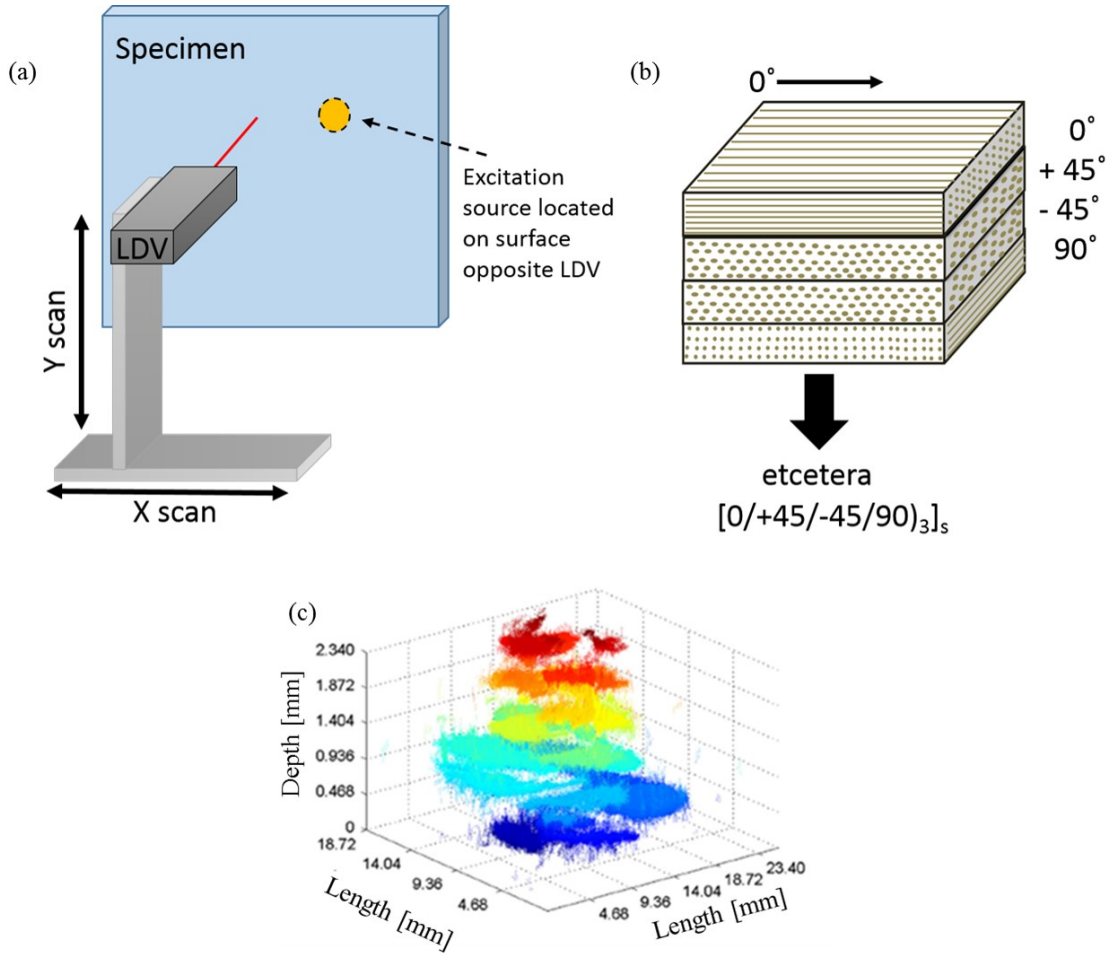


Figure 3: (a) - Diagram showing the LDV scan setup, (b) - Diagram showing the quasi-isotropic composite layup of the impact delaminated laminate, with the 0° reference direction defined, (c) - 3D image created from the microCT data of the impact induced delamination damage (color corresponds to depth through the thickness, z-axis, and is provided to more readily differentiate delamination depths).

bit high speed digitizer card at a rate of 10 MHz. Guided waves were excited in the experimental composite specimens using a 12.7 mm contact transducer. LDV data was recorded for two cases: 1) an 8 ply IM7/8552 unidirectional ($[0_8]$)pristine laminate excited using a 200 kHz 6.5 cycle Hann windowed sine wave, 2) a 26 ply IM7/8552 quasi-isotropic laminate of layup $[(0/+45/-45/90)_3]_s$ containing impact induced delamination damage and excited using a 400 kHz 5 cycle Hann windowed sine wave. The delamination was introduced via a quasi-static indentation technique [24]. For the pristine laminate case the LDV scan grid spacing was set to 1 mm and for the delamination case the grid spacing was set to 0.5 mm. A diagram of the experimental setup, the quasi-isotropic laminate layup with the 0° direction defined, and a volumetric microCT image of the delamination damage is shown in Figure 3.

3.2. Material characterization

Prior to characterizing the delamination damage using model-based inversion, a material validation study is conducted to ensure that the material properties used in the FE simulation of the experimental setup described in (3.1) are accurate. The pristine $[0_8]$ laminate used for the material characterization validation has an average thickness of 0.925 mm. While this laminate has a different thickness and layup than the 26-ply laminate containing the delaminations, we are confident that the material properties at a ply-level are similar to the 26-ply case. Using a thinner laminate in the material characterization validation has the added benefit of a




Figure 4: (a) - single snapshot in time of the experimental out-of-plane wavefield velocity data recorded by the LDV. (b) - simulated wavefield velocity using the theoretical material properties in Table 1.

smaller (and thus faster) computational model. Fast computational models are particularly important since MCMC analyses requires large number (i.e. thousands) of evaluations of the deterministic model. The inversion computational time is directly proportional to the time to evaluate a single model.

Fig. (4) shows the out-of-plane velocity $v(t, \mathbf{x})$ wavefield data. It should be noted that the A_0 mode has significantly higher out-of-plane motion compared to the S_0 mode for this laminate and excitation scenario; thus, the wavefield and wavenumber figures shown in this section primarily display the A_0 mode behavior. Figure Fig. (4a) shows the wavefield data recorded by the LDV. The velocity data appears off-centered because the LDV scan region was not centered directly over the excitation source. Fig. (4b) shows the simulated wavefield data. Note that in the simulation, a single element is used to excite the model as opposed to explicitly modeling the 12.7 mm transducer used in the experiment. This choice will not significantly alter the guided wave mode wavenumber spectrum values [11].

Fig. (5a) shows the 3D Fourier transformation of the experimental data acquired by the LDV. The red amplitude peak is due to the LDV scan grid being non-symmetric relative to the excitation source, resulting in more waves recorded in the positive k_x direction. The wavenumber spectrum for the FE model using the theoretical material properties given in Table (1) is shown in Fig. (5b). The computed wavenumbers in the 0° and 90° directions for the experiment and the simulation are shown in Table (2). For the material validation part of this study, these wavenumber values are the metrics we are using in evaluation of Eq. (4).

Table 2: Unidirectional layup model 1 wavenumber comparisons

Method	$0^\circ k (m^{-1})$	$90^\circ k (m^{-1})$
Experiment	129.5	263.4
FEM	140.0	250.5

As discussed in Sec. (2.3.1), there are seven unique material properties that describe the material tensor and density of the composite laminate. An initial material parameter study was undertaken to evaluate the sensitivity of the wavenumber values k_x and k_y in the 0° and 90° directions to changes in the material properties (e.g. E, G, ρ). The purpose of this initial study is to remove redundancies in the material parameterization (if possible) and to reduce the number of material properties from the θ_{mat} parameter vector to be characterized. Reducing the dimension of θ_{mat} exponentially decreases the amount of required samples (i.e. the computational time) to fully characterize the material, and thus care should be taken in determining which parameters truly need to be characterized.

Fig. (6) shows the likelihood evaluation (i.e. Eq (4)) as a function of various material property values, where red indicates a higher likelihood, and blue a lower likelihood. While there are seven unique material properties, visualizing more than three at a time is difficult. Additionally, showing all material property combinations




Figure 5: The wavenumber spectra ((a) - experimental; (b) - simulated) resulting from the spatio-temporal wavefield velocity data in Fig. (4). The yellow lines correspond to the experiment wavenumbers at 0° and 90° . The simulation was performed using the material property values reported in the literature.




Figure 6: The colored dots represent the likelihood function value (Eq. 4) as a function of varying material properties (where red indicates a higher likelihood, and blue a lower likelihood). It is clear that the likelihood is less sensitive to changes in the shear modulus G_{23} , given that the likelihood does not change for changing shear modulus values.

would result in several plots similar to Fig. (6). Thus, only two are shown for demonstration. As can be seen in Fig. (6a), the likelihood changes as each material property (i.e. E_{11} , E_{22} , ρ) changes, indicating that the likelihood is sensitive to those material properties. In Fig. (6b), however, the likelihood (indicated by the colors) does not change in the z-axis corresponding to material property G_{23} . This same observation occurs when G_{23} is plotted against other material properties (not shown in the figure). Because G_{23} does not appear to influence the likelihood, we will remove it from θ_{mat} and assume it is known (taken to be the value reported in the literature in Table 1). Additionally, we will assume the two unique Poisson's ratios are perfectly known in order to reduce the amount of redundant material properties to characterize. This reduces the number of unique material properties to be determined from seven to four.

The prior distribution of the material parameters are assumed to be uniform $f_\pi(\theta_{mat}) \sim U[a, b]$, where the bounds a and b for each material parameter are given in Table 3. The range in values in Table 3 was chosen such that the span of each parameter is $\pm 10\%$ to 15% of the material property value reported in scientific literature [17], [18]. It is generally not possible to know the required number of iterations required for convergence of the Markov chains *a priori*, and therefore the convergence of the chains is typically monitored using variance-based methods as in [25]. In this work, however, the chains appear to mix well and thus visually monitoring the chain

Table 3: IM7/8552 material parameter prior bounds

Property	a	b
ρ (kg/m^3)	1175	2000
E_1 (GPa)	130	210
E_2, E_3 (GPa)	7.5	10.5
G_{12}, G_{13} (GPa)	4.0	7.0

convergence was done. The Markov chains are initiated by randomly sampling from the prior distribution, and $N = 10,000$ iterations are sampled. The first 3,000 are discarded as burn-in, and the remaining 7,000 iterations are assumed to have been sampled from the posterior distribution of the material parameters. The authors in [25] recommend discarding the first half of the Markov chains, but this is conservative. In this work, we discarded the iterations that appeared to be in the transient phase of the Markov chain.

The results of the MCMC analysis for the material characterization are shown in Fig. (7) for the four material properties comprising θ_{mat} . The trace histories are shown in the top quadrants of Fig. (7). From the trace histories, it would appear that the material properties are highly correlated with respect to the posterior. As can be seen, the trace histories for all four material parameters appear to have similar trends over time (i.e. iteration). This indicates that several combinations of material parameters are equally as likely.

In addition to the trace histories, histograms of the Markov chains are also provided in Fig. (7). The histograms allow us to visualize the frequency with which the material parameter values appear in the Markov chain. As the Markov chains have been constructed such that the values contained in the data are sampled from the posterior, the histograms are thus the marginal posterior distributions of the material parameters. The median values of the material parameters are taken and used to evaluate the finite element model of the inspection. The resulting simulation wavenumber spectrum is shown in Fig. (8b), with the experimental wavenumber spectrum is shown in Fig. (8a). In both plots the yellow “experiment” line indicates the experimental wavenumber values k_x and k_y in the 0° and 90° for comparison. As can be seen in Fig. (8), the simulation wavenumber spectrum agrees very closely with the experimental wavenumber spectrum in Fig. (8a). Note that for the purposes of comparing the simulated wavenumber spectrum using the validated material properties against the experimental wavenumber spectrum, the source is explicitly modeled as a 12.7 mm diameter transducer (as opposed to the point source used in generating Fig. (5)). For the forward model evaluations for the material characterization and the delamination characterization, however, the point source is used.

3.3. Single Delamination Study

In this section we describe the analysis undertaken to understand the effects of a single delamination on the wavenumber spectrum. The goal is to understand how the delamination parameters (i.e. the in-plane location, the through-thickness location, and radius) affect the wavenumber spectrum, as the wavenumber spectrum is the metric by which we will compare the experimentally recorded data and the simulated data.

Fig. (9), shows the side view simulation of the laminate with a single delamination located half way through the thickness of the laminate, for three different time steps. The transducer source is located at the bottom center of the laminate, and the LDV measures the out-of-plane velocity on the top surface of the laminate (i.e. the opposite side of the source). As shown in Fig. (9), the delamination (the the right of the source) causes waves to be trapped above and below the delamination, resulting in higher wavenumbers. This observation matches with the expected behavior based on guided wave theory and published literature [26, 5].

The effect of the delaminations at different locations through the thickness of the laminate can be seen in Fig. (10). In the top two images in Fig. (10), a delamination is located within the black ring at approximately halfway through the thickness. This delamination depth results in higher wavenumbers above the delamination, as can be seen in the reduced wavelengths within the black circle. These decreased wavelengths result in increased wavenumbers, which manifest in the perturbation in the wavenumber spectrum. In Fig. (10c) and Fig. (10d), the delamination is located closer to the surface, resulting in even further reduction in wavelengths and increased wavenumbers, as can be seen in the corresponding wavenumber spectrum. These single delamination results indicate that the wavenumber spectrum is sensitive to the through-thickness delamination location.

In addition to seeing the effects on the wavenumber spectrum due to varying the ply layer below which the delamination occurs, we also show how other delamination locations in the x-y plane affect changes in

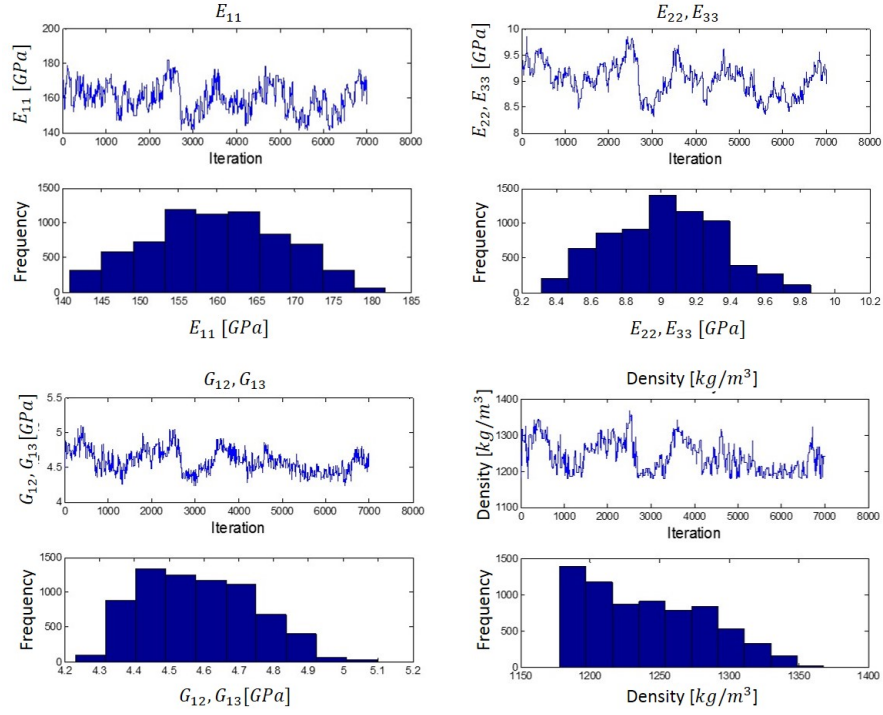


Figure 7: MCMC results for the four unique material properties. The trace histories are converted to histograms, which represent the marginal posterior probabilities. The correlations in the trace histories over iteration indicate a non-unique solution.




Figure 8: The updated wavenumber spectra using the median posterior material property values. It is clear that the simulated results shown in (b) are in significantly better agreement with the experimental measurements shown in (a) when compared to the simulation results in Figure 5 based on literature values.

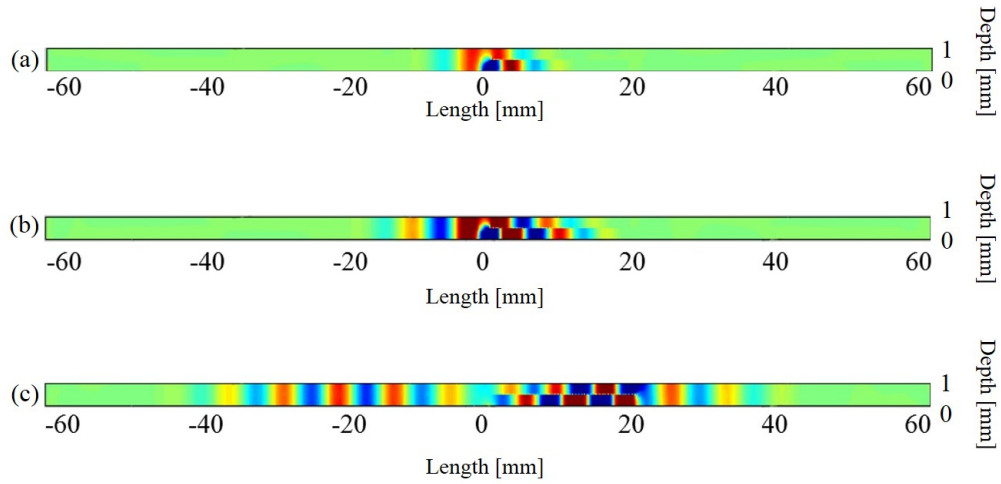


Figure 9: A side view of the laminate excited with a transducer on the bottom of the laminate at three different times. The delamination located halfway through the thickness affects the wavelengths above and below the delamination. (a) - $8.5\mu s$; (b) - $13.6\mu s$; (c) - $30.\mu s$. Supplementary video of this simulation can be found in the online version of this article.




Figure 10: Top view of the ultrasonic simulation with a delamination located at ply 14 (a), delamination located at ply 22 (closer to the surface of the laminate) (c). The resulting wavenumber spectra are shown in (b) and (d). The closer the delamination is to the surface of the laminate, the higher the wavenumbers are.




Figure 11: Examples of the wavenumber spectrum differences between a reference spectrum K_{ref} associated with a reference delamination (θ_{ref} : delamination centered at $x = 10$ mm; $y = 10$ mm; $r = 10$ mm; ply number = 14) and the wavenumber spectra associated with different delaminations (K_θ).

- (a): delamination centered at $x = 0$ mm; $y = 10$ mm; $r = 10$ mm; ply number = 14.
- (b): delamination centered at $x = -10$ mm; $y = -10$ mm; $r = 10$ mm; ply number = 14.
- (c): delamination centered at $x = 10$ mm; $y = 10$ mm; $r = 5$ mm; ply number = 14).
- (d): delamination centered at $x = 10$ mm; $y = 10$ mm; $r = 10$ mm; ply number = 22).

the wavenumber spectrum in Fig. (11). We assume a reference delamination θ_{ref} has generated a reference wavenumber spectrum K_{ref} . We perturb the delamination and see how much the wavenumber spectrum changes by plotting the difference between the reference wavenumber spectrum and the new wavenumber spectrum associated with new delamination θ . When $K_{ref} - K_\theta$ is large over the wavenumber domain, it indicates that the wavenumber spectra is sensitive to changes in the delamination feature. In general, when the wavenumber spectrum associated with the experimentally recorded data minus a candidate delamination’s wavenumber spectrum ($K_{expt} - K_\theta$) is zero over the domain, then we have found the delamination θ that exists in the experiment. From Fig. (11), we see that the wavenumber spectrum is most sensitive to changes in the delamination depth.

Table 4: Single delamination “ground truth” parameters θ_{gt} for use in MCMC verification

θ_{gt}	Value
In-plane center c_x	10 mm
In-plane center c_y	10 mm
Radius r	10 mm
Ply layer N	14

The Markov chain Monte Carlo inversion described in Sec. 2.4 is then used to characterize the delamination parameters of a single delamination. This single-delamination analysis is numerical only (inversion of simulated data) and conducted to verify that the MCMC inversion can correctly characterize a single delamination. The “ground truth” delamination parameters are given in Table 4. A signal-to-noise ratio (SNR) of 250 with a normal distribution is added to the spatio-temporal wavefield data generated from the ground truth delamination parameters, to mimic variance that could be expected in an experimental setting.




Figure 12: The MCMC results for the single delamination numerical study. The red dashed line indicates the “ground truth” delamination values. The delamination in-plane location and radius appear normally distributed while the ply number where the delamination occurs appears uniformly distributed.

The MCMC results for the single delamination case are shown in Fig. (12), where each quadrant represents the marginal posterior likelihood for each delamination parameters. The Trace plots above the histograms show the Markov chains for each parameter, where the red dashed line indicates the true value. The histograms represent the marginal posterior distributions (i.e. Eq. 6). For this single delamination, numerical problem, the MCMC is able to characterize the delamination parameters very well. The delamination in-plane location and radius appears to be approximately normally distributed. The ply number (representing the through-thickness location) appears to be more uniformly distributed. This is because the the ply number parameterization in the PZFlex model is described by an integer number. In the finite element model, the ply number is rounded to the nearest integer and the delamination is inserted at the integer ply level. Therefore, any numbers between 13.5 and 14.5 will round to integer 14 and result in the same simulation. Note that in the MCMC, the delamination petal is allowed to occur between any two ply layers.

3.4. Complex Delamination

The overall goal of this work is to determine the damage parameters that describe the complex, impact-induced delamination defect cluster shown in Fig. (13). The analyses described in the previous sections serve to illustrate that the finite element model is a validated model for use in the model-based inversion approach. Additionally, the single delamination numerical study was conducted to determine whether the wavenumber spectrum and MCMC inversion algorithm are suitably able to characterize a single delamination. In this section, we use the experimental UT data in Fig. (14) in the RJMCMC algorithm to characterize the delamination cluster. For this study of the algorithm for inversion in a case with a complex defect, we chose to avoid the complication of reflections due to the specimen boundaries and only analyzed the experimental data contained within the white box shown in Fig. (14a) and for times below 30 μs . Because the total time of the experiment




Figure 13: MicroCT scan of the impact-induced damage. (b) - approximate dimensions of the delamination cluster.

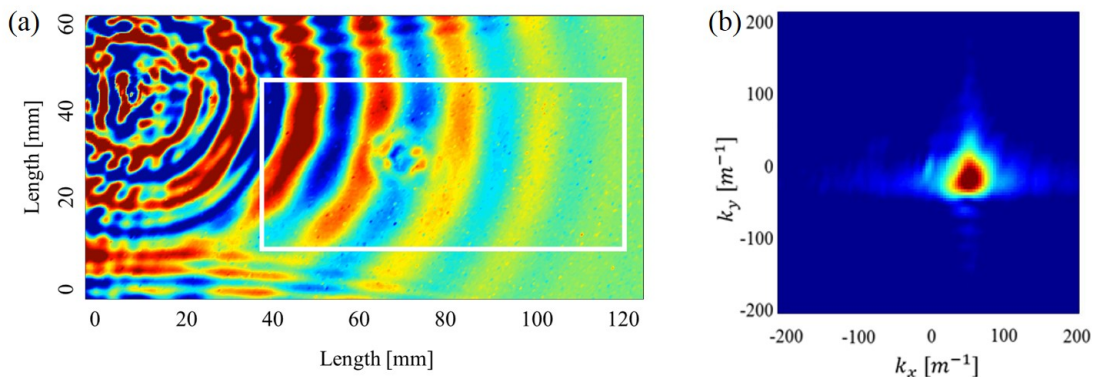


Figure 14: Experimental LDV data. The source is located in the upper left region of the laminate. The data within the white box is used in the inverse analysis (a) The associated wavenumber spectrum (b). Supplementary video of these results can be found in the online version of this article.

(and simulation) is truncated at $t = 30 \mu s$, only the faster S_0 mode is captured in the data, as can be seen in Fig. (14b).

The experimental delamination case is much more complex than the single delamination case. Using the microCT data, delamination petals were placed into the model to demonstrate and visualize the complexity of the wave interactions occurring within the laminate. The UT inspection simulation is then evaluated with the manually placed delamination petals. A side view of the UT inspection simulation (i.e. where a slice of the laminate cuts through the delamination cluster) at three time steps can be seen in Fig. (15). The several delamination petals trap waves within the plies. On the side of the delamination cluster opposite the excitation source, the waves result in the wave pattern as in the non-damaged laminate area. Again, the simulation for Fig. (15) is provided to the reader for demonstration of wave complexity created by the real delamination shape.

In the RJMCMC algorithm, we assume the number of delamination petals is not known *a priori*; rather, we let the algorithm estimate the number of petals n_k and the parameters (i.e. location including depth) and radius) of each petal. We initialize the damage parameter vector θ with one petal ($n_k = 1$), located in the lower left region of the white box in Fig. (14). We require information about the petals for use in the prior distribution $f_\pi(\theta)$. We assume that there can be zero, one, or two petals in between each layer and any two adjacent layers can have petals between them. The centers of the petals can be located anywhere in the white box in Fig. (14), and the petal radii are uniformly distributed between 1 mm and 15 mm. The petals are allowed to occur between any two ply layers. In addition to priors for the damage parameter values, a prior for the number of petals, n_k , is required, since the damage model size is unknown. Consistent with Occam's razor heuristic [27], [28], a prior for n_k is chosen such that fewer petals are favored over more petals, following the prior work in [23].

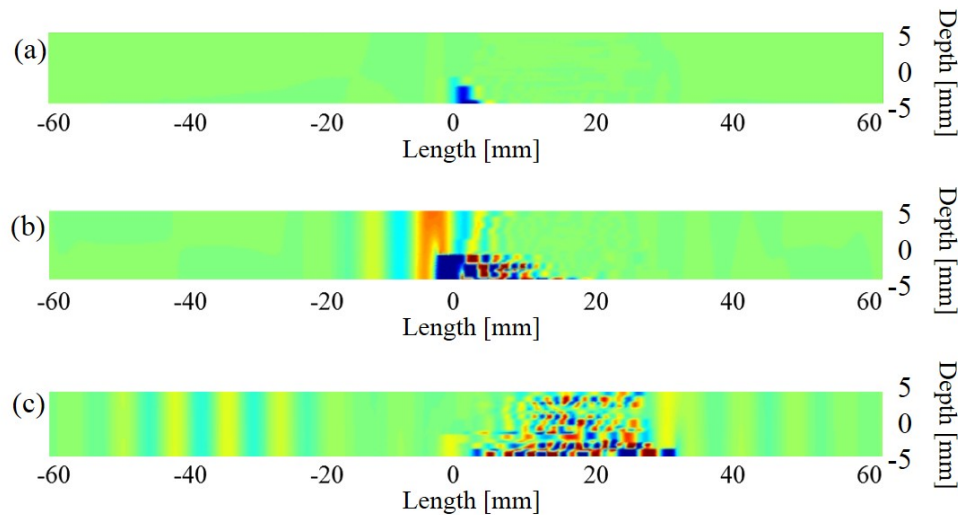


Figure 15: Side view of the 26 ply laminate where damage petals have been inserted into the model based on the CT data. (a) - $1.6\mu s$; (b) - $3.0\mu s$; (c) - $10\mu s$. Supplementary video of this simulation can be found in the online version of this article.

The Markov chain results are shown in Figs. 16-18 for each petal parameter (the ply number where the delamination occurs; horizontal and vertical centers of the delamination; and the radius of each ply), where the first $N = 1,000$ iterations are discarded as burn-in. In the case of the complex damage, the objective is to locate the center of the delamination petal cluster and envelope the size of the damage site. It can be seen from the five colored lines in the top halves of Figs. 16-18 (corresponding to the trace histories of each delamination petal) that the individual chains have converged to values after $N = 1,000$. Additionally seen in the trace history plots, the trace history lines begin at varying iteration numbers. This is because the RJMCMC algorithm is able to introduce new delamination petals during the algorithm. If a new delamination petal is introduced and the resulting posterior probability is greater than the posterior probability of the current damage state θ , then the new petal is accepted and n_k increases by one. The histograms in Figs. 16-18 correspond to the pooled Markov chain results at the final iteration ($N = 9,000$), which is why the histograms appear “spikey”. In Fig. (17), the histogram does encompass the “true” monolithic defect center, indicating the delamination petals are aggregating such that their center is located near the true defect center.

As can be seen in Figs. 16-18, the solution θ is initialized far from where the delaminations actually are. After approximately 1,000 iterations, the blue trace history line jumps to a new delamination, closer to where the delamination damage site is. New delamination petals subsequently appear after $N = 1,000$ iterations. As can be seen from the petal centers results in Fig. (17), the new delaminations emerge with centers near each other. The ply number results in Fig. (16), indicate that some of the delamination petals occur near the top of the laminate where the data is being captured (corresponding to the high ply numbers) and some of the delaminations are closer to the bottom of the laminate. According to Fig. (18), the delamination petal with the smallest radius occurs near the top of the laminate, which is consistent with the microCT data.

While the individual petal results are interesting, we are more interested in how the individual petals aggregate into a monolithic damage site. The individual petals’ convergence is shown in Fig. (19) at six different iterations. At the time $N = 0$ box in Fig. (19), the delamination petal corresponds to the value in which θ was initialized. As iterations increase, more delamination petals arise, and the monolithic defect image begins to appear. While the box corresponding to $n = 9,000$ (i.e. the final iteration), appears to only have four delaminations, there are actually five petals, where one petal is hidden beneath the four visible petals. This is a common inspection issue (especially in scenarios with single-sided access), where shallow delaminations shadow the regions beneath them, making it difficult or not possible to detect details of deeper delamination damage. The inversion approach applied in this study does not require that an inspector be able to image the delaminations, but rather the algorithm places the petals where the resulting inspection data best “matches” the experimental data.

Finally, one of the goals with this analysis is to be able to determine the extent of the delamination damage. We




Figure 16: The results of the ply numbers where the individual petals occur. The pooled histogram data appears “spiky” because the petals occur throughout the laminate thickness.




Figure 17: The results of the individual petal centers. The petal centers appear to converge to near the true solution of $x = 76$ mm and $y = 32$ mm. The pooled histograms surround the true values (red dashed lines).




Figure 18: Results of the individual petal radii. Top - according to the blue chain, the initial delamination with radius of 5 mm jumps to < 2 mm, and soon after a new larger delamination petal appears.




Figure 19: Snapshots of the individual petals during the algorithm. At $N = 0$, the delamination is initialized far from the truth so as not to bias the solution. The colors correspond to location of the petals through the thickness, where red indicates closer to the top of the laminate. Supplementary video of these results can be found in the online version of this article.




Figure 20: Aggregated Markov chain results for the monolithic defect size. (a) - trace histories of the defect size, where jumps correspond to the addition of new petals. (b), (c) - the pooled histogram data. The defect size is underestimated according to the red, dashed lines representing the true size from the microCT data.




Figure 21: Simulated experiment (a) with predicted damage parameters. Wavenumber spectrum associated with simulated experiment (b). Supplementary video of this simulation can be found in the online version of this article.

can accomplish this by considering the individual petals as a whole. The aggregated defect size overall dimensions are estimated by considering the spread of the delaminations from each of the individual petals. The aggregated data are shown in Fig. (20). The histograms indicate that the aggregated defect size is underestimated with this approach (the vertical red dashed lines correspond to the approximated true defect size from the microCT data). It is possible that if the RJMCMC analysis was allowed to continue, then more delamination petals could have been added to the results which may have increased the total defect size. However, Monte Carlo algorithms such as this are notoriously computationally expensive, as they involve evaluating expensive finite element forward solvers. The final Markov chain solution and the resulting simulation and wavenumber spectra are shown in Fig. (21). It can be seen that the simulation compares well with the experimental data. In the simulation, however, absorbing boundary conditions are placed along the laminate in-plane boundaries to mitigate the effects of boundary reflections, which are present in the experimental data but were time-gated out. This was done to reduce the finite element model size while still having the ability to run the model for the same amount of time as the experiment. As in the experiment, the data recorded for use in the inverse problem is contained within the white box in Fig. (21a). The associated wavenumber spectrum in Fig. (21b) has a peak at the wavenumbers corresponding to the wavenumbers of the slower S_0 mode, where any deviations in the peak are due to the effects of the delaminations.

4. Conclusions

In this work we present a model-based inversion methodology for characterizing the extent of impact-induced damage in composite laminates. In the model-based inversion, experimental inspection data in the form of LDV data on the laminate surface is compared to a defect-parameterized simulation model. The model used as the forward, deterministic model consists of a finite element model created using the finite element program PZFlex. The experimental and simulated data are analyzed by considering their wavenumber spectra, where deviations from the wavenumbers associated with the pristine laminate indicate delaminations. A sensitivity and validation study was undertaken to ensure that the model parameters and the delamination defects accurately represent the mechanics of the inspection.

The inversion is effected through Reversible jump Markov chain Monte Carlo (RJMCMC), in which the number of delamination petals is allowed to increase or decrease. This flexible damage model allows the inversion algorithm to automatically estimate the number of delamination petals that can suitably describe the extent of the damage. Additionally, the use of RJMCMC enables visualization of the variance on the delamination values. While the RJMCMC algorithm does slightly underestimate the extent of the delamination cluster, it is able to accurately locate where the delaminations occur within the laminate. Additional Markov chain iterations or improved knowledge about the correlations of the delamination petals (e.g. exploiting the spiral pattern typical of impact-induced delaminations) could result in a more accurate solution, and these are areas for future work.

Acknowledgments

The authors acknowledge support from the nonreimbursement Space Act Umbrella Agreement SAA1-21887 between Thornton Tomasetti Weidlinger Applied Science and the National Aeronautics and Space Administration (NASA) Langley Research Center.

- [1] A. Williams. Aerospace growth climbs with carbon fiber. <http://www.compositesworld.com/columns/aerospace-growth-climbs-with-carbon-fiber>, 2016.
- [2] I. Daniel and O. Ishai. *Engineering mechanics of composite materials*. Oxford Press, New York, 2006.
- [3] P. Joyce and T. Moon. Compression strength reduction in composites with in-plane fiber waviness. *ASTM International*, 7:76–96, 1998.
- [4] K. Mizukami, Y. Mizutani, A. Todoroki, and Y. Suzuki. Detection of in-plane and out-of-plane fiber waviness in unidirectional carbon fiber reinforced composites using eddy current testing. *Composites Part B: Engineering*, 86:84–94, 2016.
- [5] M. Rogge and C. Leckey. Characterization of impact damage in composite laminates using guided wavefield imaging and local wavenumber domain analysis. *Ultrasonics*, 53:1217–1226, 2013.
- [6] X. Yu, M. Ratassepp, and Z. Fan. Damage detection in quasi-isotropic composites bends using ultrasonic feature guided waves. *Composites Science and Technology*, 141:120–129, 2017.
- [7] Z. Su, L. Ye, and Y. Lu. Guided lamb waves for identification of damage in composite structures: A review. *Journal of sound and vibration*, 295:753–780, 2006.
- [8] J. Michaels, A. Dawson, T. Michaels, and M. Ruzzene. Approaches to hybrid shm and nde of composite aerospace structures. *Proceedings SPIE*, 9064:906427, 2014.
- [9] C.T. Ng, M.Veidt, L.R.F. Rose, and C.H. Wang. Analytical and finite element prediction of Lamb wave scattering at delaminations in quasi-isotropic composite laminates. *Journal of Sound and Vibration*, 331:4870–4883, 2012.
- [10] D. Singh, R. Guerjouma, and M.Bentahar. Interaction of fundamental lamb modes with a point impact damaged zone in composite. *Proceedings of the Societe Francaise d’Acoustique*, 2012:2423–2428, 2012.
- [11] C. A. C. Leckey, M. D. Rogge, and F. R. Parker. Guided waves in anisotropic and quasi-isotropic aerospace composites: Three-dimensional simulation and experiment. *Ultrasonics*, 54:385–394, 2014.
- [12] B. Murat, P. Khalili, and P. Fromme. Scattering of guided waves at delaminations in composite plates. *Journal of the Acoustical Society of America*, 139:3044–3052, 2016.
- [13] R. Roberts and S. Holland. Model based defect characterization in composites. *AIP Proceedings*, 1806:090015, 2017.
- [14] A.A. Fahim, R. Gallego, N. Bochud, and G. Rus. Model-based damage reconstruction in composites from ultrasound transmission. *Composites: Part B*, 45:50–62, 2013.
- [15] N. Bochud, A.A. Fahim, A.M. Gomez, and G. Rus. Impact damage characterization in composites using signal processing techniques. *Procedia Engineering*, 14:169–176, 2011.
- [16] PZFlex. <http://pzflex.com>.
- [17] P.P. Camanho, P.Maimi, and C.G. Davila. Prediction of size effects in notched laminates using continuum damage mechanics. *Composite Science and Technology*, 67:2715–2727, 2007.
- [18] A.S. Kaddour, M.J. Hinton, S. Li, and P.Smith. Instructions to the contributors of the third world-wide failure exercise. *WWFE-III: part A*, 2008.
- [19] W. K. Hastings. Monte Carlo sampling methods using Markov chain and their applications. *Biometrika*, 57(1):97–109, 1970.
- [20] H.M. Reed and C.J. Earls. Stochastic identification of the structural damage condition of a ship bow section under model uncertainty. *Ocean Engineering*, 103:123–143, 2015.
- [21] P.J. Green. Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82:711–732, 1995.
- [22] S. Richardson and P.J. Green. On bayesian analysis of mixtures with an unknown number of components (with discussion). *Royal statistical society: Series B*, 59:731–792, 1997.

- [23] P.S. Koutsourelakis. A multi-resolution, non-parametric, bayesian framework for identification of spatially-varying model parameters. *Journal of Computational Physics*, 228:6184–6211, 2009.
- [24] P. Juarez and C. Leckey. Multi-frequency local wavenumber analysis and ply correlation of delamination damage. *Ultrasonics*, 62:56–65, 2015.
- [25] S. P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, 1988.
- [26] Joseph L. Rose. *Ultrasonic Waves in Solid Media*. Cambridge University Press, 1999.
- [27] W. Jeffreys and J. Berger. Ockham’s razor and bayesian analysis. *American Scientist*, 80(1):64–72, 1992.
- [28] C.E. Rasmussen and Z.Ghahramani. Occam’s razor. *Neural Information Processing Systems*, 13:294–300, 2001.