



Stochastic identification of the structural damage condition of a ship bow section under model uncertainty



H.M. Reed*, C.J. Earls

School of Civil and Environmental Engineering Cornell University Ithaca, NY 14850, United States

ARTICLE INFO

Article history:

Received 3 March 2014

Accepted 27 April 2015

Available online 27 May 2015

Keywords:

Ship bow structure

Fluid structure interaction

Stochastic inversion

Bayesian estimation

Damage identification

Reversible jump Markov chain Monte Carlo

ABSTRACT

A accurate, quantifiable means of assessing structural damage condition are paramount for maintaining the structural integrity of ship hull forms. Toward this end, precise knowledge of the location and magnitude of any imperfections (i.e. geometric imperfections in the form of denting and corrosion patches) must be determined, along with concomitant uncertainties accompanying such predictions. The current paper describes a non-contact approach to identifying and characterizing such imperfections within the submerged bow section of a representative ship hull. By monitoring the pressure field local to the acoustically excited hull section, it is shown how the resulting data can be used to identify the parameters describing the structural damage field. In order to perform the identification, a fluid-structure model that predicts the spatio-temporal pressure field is required. A Bayesian, reversible jump Markov chain Monte Carlo approach is then used to generate the imperfection parameter estimates and quantify the uncertainty in those estimates. This approach is particularly appealing as it does not allow for the damage model to be explicitly known *a priori*. Convergence of the Markov chains is assessed, and estimates of the Monte Carlo standard error (MCSE) are provided.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Advances in the state-of-the-art in naval architecture have led to the ability to field new hull forms that frequently are made from new material grades. These advancements precipitate increased uncertainties surrounding long term performance, due to a lack of experience within the design community. Essentially, new hull forms, made from new materials, are being fielded in response to performance demands and economic constraints of modern service contexts. Because of the uncertainties surrounding the performance of the new designs, it may be advisable to monitor the resulting hull structures throughout their life in such a manner that an accurate assessment of the condition of the state of the structure, at any point in time, is possible. Additionally, as decades of experience and data are not available for these new designs, a means for making a prognosis for the future state of the structure is desirable.

Since the publication of Koiter's seminal dissertation in 1945, it has become well known that initial geometric imperfections (e.g. denting) in shell type structures, for example ship hulls, may lead to dramatic erosions in ultimate buckling strength (Bazant and Cedolin, 1991; Featherston, 2003; Singer et al., 2002). Such imperfections may arise due to manufacturing, fabrication,

construction, or service conditions. A recent and comprehensive survey of research developments during the period 1996–2006 (Edlund, 2007) highlights the fact that understanding the nature and effects of imperfections in shell structures continues to be a fertile line of inquiry to this day. In addition to dent-type imperfections as a source of reduced load carrying capacities, another geometric imperfection relevant to ship hull structures is corrosion. Corrosion erodes the cross-sectional area of structural components within the hull; thus reducing the ability of the hull to resist external loads (Akpan et al., 2001). However, the jump from the notional understanding that geometric imperfections adversely affect hull strength, to the realization of a practically useful means for predicting the actual strength of in-service ship hulls, is formidable. It is virtually impossible to rationally guess the precise damaged condition that may be manifest in a given ship structure. One might attempt to infer the damage field from observations, however practical methods for identifying the exact location and magnitude of the damage, from observed hull response data, are currently lacking. The problem is even more challenging if this identification is to be made *in situ*.

The traditional structural health monitoring (SHM) problem is fundamentally one that can be described as statistical pattern recognition (SPR), composed of the following four step process (Farrar and Worden, 2007):

1. Operational evaluation.
2. Data acquisition, normalization and cleansing.

* Correspondence to: Weidlinger Associates, Inc., New York, NY 10005, United States. Tel.: +1 774 218 5137.

E-mail address: heather.reed@wai.com (H.M. Reed).

3. Feature selection and information condensation.
4. Statistical model development for feature discrimination.

In an SPR SHM approach, the goal is to determine the presence of (and in advanced cases quantify the values of) structural damage by statistically correlating features extracted from sensor response data with data features of known damage. However, this typically requires some *a priori* knowledge of how the structure will behave when it is damaged. The approach described herein is a *model-based* approach, in which a high-fidelity computational analog of the monitored structure (i.e. a finite element model) is parameterized with features describing the damage to be determined. The solution is then to quantify these damage parameters in the finite element model, so that the model generated response data “matches” the observed sensor response data.

Identifying localized damage in structures is a difficult problem, made so by the limited influence the damage typically has on observed response data from which the damage related parameters are to be estimated. The result is a challenging system identification problem for which an obvious solution is unclear. The two requirements for any system identification approach are the following: (1) a means of acquiring observations from the object of interest and (2) a model describing the influence of the parameters of interest (in this case, parameters describing the denting imperfections and corrosion) on the observed data. The present work focuses on the *a posteriori* characterization of the actual damaged condition (denting imperfections or corrosion) within the bow section of a representative ship (in this case a prototype of the Joint High Speed Sealift vessel) (Cusanelli and Chesnakas, 2007). The current research employs a sparse array of acoustic fluid pressure sensors (i.e. hydrophones) to help identify structural condition, simply by “listening” to the pressure waves induced by ambient vibrations occurring during the ship's service condition. A model describing the fluid–structure interaction (FSI) and the influence of the damage on the model response is also developed, and is described in detail in section 2.1.

In an earlier work (Reed et al., 2014), this basic approach was used to generate marginal posteriors of geometric imperfection damage parameters, enabling uncertainty quantification of the parameter estimates. This approach was also used in Reed et al. (2013) to generate *maximum likelihood estimates* (MLEs) of the imperfection location and magnitude. While MLE is a powerful approach to draw statistical inference about unknown parameters, the associated methods for quantifying the uncertainty in the estimate require either replication (many trials) or asymptotic approximations. The present research furnishes a Bayesian approach for estimating the parameters describing the imperfections, estimating the number of imperfections, and quantifying the uncertainty in those estimates. It will be shown that this can be accomplished through the solution of a stochastic inverse problem effected using dependent sampling techniques.

1.1. Background and motivation

Discrepancies between theoretical predictions and experimentally observed ultimate strengths in shell structures (e.g. ship hull structures, as considered in this work) may arise from imperfections related to boundary conditions, material properties, and shell thickness (Papadopoulos and Papadrakakis, 2004, 2005; Papadopoulos and Pavlos, 2007; Papadopoulos et al., 2009). The principal cause of this variation, however, is the presence of geometric imperfections, (e.g. denting and corrosion), resulting in deviations from the assumed perfect shell geometry (Singer et al., 1998, 2002). Structural damage (i.e. any condition that changes the material and/or geometric properties of structural systems) may arise during manufacturing, fabrication, construction, service conditions, or natural aging. The resulting differences

between observed and theoretical ultimate strength make the development of principled design specifications difficult.

A recent study (Kristanic and Korelc, 2008) explored an efficient optimization-based approach for obtaining a worst case initial geometric imperfection field within a shell structure, as constrained by bounds on feasibility, as related to hypothetical deformations being explored by the optimization algorithm. The current research is different in that its approach aims to uncover the actual imperfection field present in the structure, as opposed to hypothetical, worst-case deformations. Such knowledge could allow for improved inspection protocols (e.g. physically inspecting the structure when there is evidence that damage exists, versus inspecting the structure at regular, timed intervals).

The goal of the present work is to model, and subsequently identify and characterize, dent-like and corrosion patch structural damage in a partially submerged ship hull bow section. Furthermore, we wish to quantify uncertainty in the damage estimates. In a previous study, Reed et al. (2014), a similar problem was solved using a Bayesian estimation approach to identify and characterize the parameters of a dent-like imperfection present in a simple, idealized submerged shell structure, where the damage model was assumed to be perfectly known. The Markov chain Monte Carlo approach employed in the previous work was able to successfully characterize the damage, as well as the uncertainty surrounding the parameter estimates. A reversible jump Markov chain Monte Carlo example was also provided in that work to demonstrate the efficacy of estimating the damage field when the damage model is not perfectly known *a priori*.

In another work, Reed et al. (2013), the same simple, idealized submerged shell structure damage identification problem was solved using a modified differential evolution (DE) algorithm. The algorithm was designed to quickly find the MLEs of the damage parameters, thereby providing point estimates of the imperfections. Such MLE-based approaches are attractive from a practical point of view (e.g. they may be employed as part of fast inverse solution algorithms), yet they can sometimes present challenges in terms of quantifying uncertainty. Typically one either repeats the MLE many times and uses the resulting spread of estimates to produce a confidence interval. Or, one could use an asymptotic approximation and base confidence intervals on the curvature of the likelihood function. However for many structural systems the likelihood contains numerous maxima with a very high degree of curvature (“spikey” likelihoods are not uncommon) (Reed et al., 2013, 2014; Stull et al., 2011), a situation that violates the underlying assumptions of this approach. The Bayesian formulation described in Section 2.2 neither requires a single-peaked likelihood nor relies on asymptotic approximation. Additionally, the reversible jump Markov chain Monte Carlo approach described in 2.5 does not require that the damage model be known exactly *a priori*, allowing for a robust stochastic inversion solution.

Much of the research involving damage detection within the context of FSI and shell-type structures has focused on structures that are filled (either completely or partially) with fluid, such as spherical storage tanks (Curadelli and Ambrosini, 2011), composite fuel tanks (Zhou et al., 2010), and cylindrical laminated composite shells (Yu et al., 2007, 2007). The works by Curadelli and Ambrosini (2011) and Zhou et al. (2010) explore the extent to which the presence of structural damage can be detected by monitoring changes in modal parameters between the undamaged (perfect) and a damaged, fluid-filled structure. If damage identification can be classified into four levels (Rytter, 1993):

1. Level 1: Determination that damage is present in the structure,
2. Level 2: Determination of the geometric location of the damage,

3. Level 3: Quantification of the severity of the damage,
4. Level 4: Prediction of the remaining service life of the structure,

then Curadelli and Ambrosini (2011) and Zhou et al. (2010) focus on level 1.

A more sophisticated approach to damage identification is provided by Yu et al. (2007, 2007), which addresses the second level outlined above. In that work, a damage index is computed as the change (between the undamaged and damaged structure) in the energy spectrum of the decomposed wavelet signals of structural dynamic responses, and then used within the context of an *artificial neural network* (ANN) to identify the location and length of a crack in a laminated composite shell partially filled with fluid. While successfully able to locate the damage, this approach furnished a single solution to the inverse problem. Moreover, such “data-driven” approaches are only capable of identifying cases on which they have been trained (i.e., previously unseen damage cases may or may not be classified). These types of approaches can be similarly confused by fluctuations in covariates (temperature, sensor noise, etc.) that were not present in the training data.

The works described above require some knowledge of the response of the undamaged fluid–structure model *a priori*, for the purpose of calculating damage feature indices. The current work, however, requires no knowledge of the actual undamaged hull’s response (which is typically not known, and is impossible to observe once the structure has been damaged), but rather directly models the physics of the unknown damage condition in the hull structure. The present work not only estimates the structural damage condition, but also quantifies the uncertainty surrounding those damage parameter estimates, addressing both levels 2 and 3, and enabling level 4 to be approached in a principled manner.

The approach taken in this research involves solving an inverse problem (i.e. recovering the damage parameters) within the context of a Bayesian analysis, effected using a reversible jump Markov Chain Monte Carlo sampling approach. The current research employs a sparse array of acoustic fluid pressure sensors (i.e. hydrophones) to help identify the damage field simply by “listening” to the pressure waves induced by vibrations during the hull’s service condition response. A similar approach was adopted in Zein et al. (2008), which compared the accuracy and computation time of a Markov Chain Monte Carlo simulation and a *simultaneous perturbation stochastic approximation* (SPSA) method (essentially an MLE method) of estimating the density and Lamé coefficients of a layered solid in contact with a fluid medium, by measuring the variation of the pressure in the fluid as a result of the propagation of a seismic wave. The authors of that study concluded that the SPSA method yields the same accuracy in the parameter estimations as the MCMC approach, but in considerably less computation time. While this may be true, the SPSA provides little information about the uncertainty in the final solution.

As computational experiments increasingly become the standard for modeling complex, physical behavior, the development of solution techniques that affect accurate and meaningful quantification of errors in model-based predictions becomes paramount (Ghanem and Wojtkiewicz, 2004). The present research furnishes a method of quantifying the uncertainty in the parameters characterizing the structural damage condition within the context of a Bayesian analysis, under damage model uncertainty, through the solution of a stochastic inverse problem effected using dependent sampling techniques. The techniques are subsequently demonstrated on a realistic ship hull bow section problem.

2. Problem description

2.1. Forward problem

The fluid structure domain we consider is closed and bounded: $\bar{\Omega} = \bar{\Omega}_s \cup \bar{\Omega}_f \subset \mathbb{R}^3$, where the boundary $\partial\Omega = \partial\Omega_s \cup \partial\Omega_f$; with the FSI boundaries of the solid and fluid domains denoted as $\partial\Omega_s^{fsi}$ and $\partial\Omega_f^{fsi}$, respectively. In this way, the fluid–structure interface $\partial\Omega_{fsi}$ may be thought of as a disjoint union of the sets of the solid and fluid boundaries that are in physical contact with one another: as approached from within $\partial\Omega_s$, in the case of $\partial\Omega_s^{fsi}$, and as approached from within $\partial\Omega_f$, in the case of $\partial\Omega_f^{fsi}$.

On this problem domain, we require a coupled dynamic equilibrium between the solid and the fluid. The solid response is governed by the equation of motion (neglecting a body force term) within the context of a linear elastic material response:

$$\nabla \cdot \bar{\sigma} = \rho \ddot{u}(\vec{x}, t), \quad \bar{\sigma} = C^{IV} : \bar{\epsilon}(u(\vec{x}, t)), \quad \forall \vec{x} \in \Omega_s \subset \mathbb{R}^3 \quad (1)$$

where $u(\vec{x}, t)$ is the displacement field satisfying all prescribed boundary conditions on the solid boundary; C^{IV} is the fourth order tensor relating the infinitesimal strain tensor, $\bar{\epsilon}$, and the Cauchy stress tensor, $\bar{\sigma}$; and ρ is the solid material density.

The requirement for the dynamic equilibrium of the fluid response is governed by the homogeneous acoustic wave equation:

$$\nabla^2 p(\vec{x}, t) = \frac{1}{c^2} \ddot{p}(\vec{x}, t), \quad \forall \vec{x} \in \Omega_f \subset \mathbb{R}^3 \quad (2)$$

where p denotes the scalar pressure in the fluid satisfying all prescribed boundary conditions on the fluid boundary; and c is the speed of sound in the fluid.

To arrive at the finite element weak formulation, we multiply the partial differential equations (PDE) by admissible test functions $w_s(\vec{x}, t)$ and $w_f(\vec{x}, t)$, and then apply Green’s lemma (Langtangen, 1999) to the subsequent weak forms. The fluid–structure system is then coupled via continuity of linear momentum at the fluid–structure interface:

$$\rho \ddot{u}_n = \frac{\partial p}{\partial n}, \quad (3)$$

i.e. the forces due to the nodal accelerations at the interface are offset by the pressure gradient in the fluid, normal to the interface. The normal derivative is then considered in the context of the analogy presented in Everstine (1981), in which the pressure gradient is proportional to the force of the fluid on the structure divided by the discrete nodal tributary area, resulting in the coupled system of equations Eq. (4) to be solved via integration in time. In this case, the Newmark integration scheme is used with optimal parameters, $\delta = \frac{1}{2}$ and $\alpha = \frac{1}{4}$, in accordance with the recommendations in Bathe (1996). The monolithically coupled FSI system may be expressed as

$$\begin{bmatrix} \bar{M} & \bar{0} \\ \rho(\bar{G}\bar{A})^T & \bar{Q} \end{bmatrix} \begin{Bmatrix} \ddot{u} \\ \ddot{p}_s \end{Bmatrix} + \begin{bmatrix} \bar{K} & \bar{G}\bar{A} \\ \bar{0} & \bar{H} \end{bmatrix} \begin{Bmatrix} u \\ p_s \end{Bmatrix} = \begin{Bmatrix} \bar{F} - \bar{G}\bar{A}p_i \\ \rho \bar{A} \ddot{u}_{ni} \end{Bmatrix} \quad (4)$$

where \bar{M} and \bar{Q} are the solid and fluid mass matrices, respectively; \bar{K} and \bar{H} are the solid and fluid stiffness matrices, respectively; \bar{A} is the matrix of tributary areas along the fluid–structure interface; and \bar{G} is a matrix of direction cosines that converts a vector of outward normals at the interface to its components with respect to the global coordinate system. On the right hand side of Eq. (4), \bar{F} is the vector of applied tractions on the structure; p_i are the incident pressures in the fluid; and \ddot{u}_{ni} are the incident nodal accelerations on the interface (Everstine, 1997).

2.2. Stochastic inverse problem

We assume that our submerged structure has M pressure sensors that are sampled at discrete times $t_i = i\Delta t$, $i = 1, \dots, N$ with constant sampling frequency $1/\Delta t$. We further assume that only sensor noise will contribute to discrepancies between FSI model and observations, such that our model for the observed pressure field may be written as

$$\mathbf{y} \equiv y_{ij} = p_{ij}(\boldsymbol{\theta}) + \eta_{ij}, \quad i = 1, \dots, N; j = 1, \dots, M. \quad (5)$$

Here y_{ij} is the observed sensor response data at spatial location j and discrete sampling time i ; $p_{ij}(\boldsymbol{\theta})$ is the associated deterministic FSI model output, given the damage parameter vector $\boldsymbol{\theta} = \{\theta_1, \dots, \theta_P\}$; and η_{ij} are 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 the value y_{ij} , given the model response $p_{ij}(\boldsymbol{\theta})$ as

$$f_\eta(y_{ij} | p_{ij}(\boldsymbol{\theta})) = \frac{1}{\sqrt{2\pi\sigma_\eta^2}} e^{-(1/2\sigma_\eta^2)(y_{ij} - p_{ij}(\boldsymbol{\theta}))^2} \quad (6)$$

Eq. (6) is the probability density function (PDF) for each observation; and clearly relates the damage parameters to the observed data. Because the noise is assumed i.i.d. and stationary, we may express the PDF for the entire collection of data as the product of the individual (marginal) distributions of each spatio-temporal point:

$$\begin{aligned} f_\eta(y_{ij} | p_{ij}(\boldsymbol{\theta})) &= \prod_i \prod_j \frac{1}{\sqrt{2\pi\sigma_\eta^2}} e^{-(1/2\sigma_\eta^2)(y_{ij} - p_{ij}(\boldsymbol{\theta}))^2} \\ &= \frac{1}{(2\pi\sigma_\eta^2)^{N \times M/2}} \exp \left[-\frac{1}{2\sigma_\eta^2} \sum_{i=1}^N \sum_{j=1}^M (y_{ij} - p_{ij}(\boldsymbol{\theta}))^2 \right] \\ &\equiv f_\eta(y_{ij} | \boldsymbol{\theta}). \end{aligned} \quad (7)$$

Eq. (7) is a probabilistic model, predicting the sequence of observed pressures, given a damage field $\boldsymbol{\theta}$, and is referred to as the *likelihood* of the observations. An MLE therefore provides the values $\boldsymbol{\theta}$ that maximize Eq. (7). However, rather than focusing on a single “best” value (in the sense of maximizing the likelihood), we treat each damage parameter 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, one of the main goals of this paper.

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 $\boldsymbol{\theta}$ in terms of the

likelihood (7) as

$$f_\theta(\boldsymbol{\theta} | \mathbf{y}) = C^{-1} f_\eta(\mathbf{y} | \boldsymbol{\theta}) f_\pi(\boldsymbol{\theta}) \quad (8)$$

where $f_\pi(\boldsymbol{\theta})$ is the *prior* PDF, which encapsulates any prior information about the parameters $\boldsymbol{\theta}$; and C^{-1} is defined as the multidimensional integral $\int_{\mathbb{R}^P} f_\eta(\mathbf{y} | \boldsymbol{\theta}) f_\pi(\boldsymbol{\theta}) d\boldsymbol{\theta}$, 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 parameter $\theta_l \in \boldsymbol{\theta}$, $l = 1, \dots, P$, we are required to integrate

Eq. (8) over each of the parameters in $\boldsymbol{\theta}$, excluding θ_l . If we define the parameter vector that excludes parameter θ_l as $\boldsymbol{\theta}_{-l}$, then the expression for the marginal posterior distribution for parameter θ_l is given as

$$f_\theta(\theta_l) = \int_{\mathbb{R}^{P-1}} f_\theta(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}_{-l} \quad (9)$$

Eq. (9) is the end goal of this Bayesian analysis; providing a complete probabilistic description of our unknown damage parameters. The problem is that Eq. (9) 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 (Fig. 1).

2.3. Acoustic FSI modeling with damage

The deterministic solver considered in this work consists of a representative bow section whose exterior is in contact with an acoustic fluid. Within the fluid domain, there exists an array of 36 sensors that record the acoustic scattering pressure emanating from the vibrating ship structure consistent with ambient service conditions. The temporal output from these sensors comprise the model response, \mathbf{p} , defined in the previous section. The loading is chosen so as to be consistent with the vibration of the ship's engine. It is applied sinusoidally as if coming from the stern toward the bow, as illustrated in Fig. 2. Using Wiggins et al. (2012) as the basis for our loading, the amplitude and frequency are chosen so as to be realistic with respect to actual ship loads, and are depicted in Fig. 2.

Since only half of the bow is being considered, the nodes are fixed in the y direction along the edge of the hull (see Fig. 3). For numerical stability, two nodes are fixed in the vertical direction (z), and a single node is fixed in the direction of the applied load (x), as given in Fig. 3.

A fluid domain is imposed on the hull, extruding into the y direction. Two boundary conditions for the fluid are considered: a pressure release (i.e. $p=0$) condition imposed along the free surface of the fluid; and an absorbing condition imposed along all of the other edges of the fluid domain (except at the fluid–structure interface), (see Fig. 4). The absorbing boundary condition is employed to simulate an endless fluid domain, the details of which are described in Everstine (1997).




Fig. 1. Schematic of an acoustic FSI problem.




Fig. 2. Bow section of the JHSS model. Load is applied sinusoidally as depicted.




Fig. 3. Ship structure shown with Dirichlet boundary conditions.




Fig. 4. Surrounding fluid shown with pressure release boundary condition along the free surface of the fluid and absorbing boundary condition along the other faces of the fluid domain.

In this work, $N=31$ time samples are used in the evaluation of the likelihood. The analysis starts at time $t=0.0$ s, and ends at time $t=0.012$ s. The sampling time step Δt is constant at $3.9e-4$ s. The end time parameter is chosen to furnish enough time for the acoustic wave, emitted from the hull, to make it to the hydrophones, and the time step size is chosen to adequately resolve the wave itself.

Two types of damage are considered in this work: one in the form of an initial imperfection, consistent in form to a dent; and the other as a patch of corrosion. We require a way to parametrize forms of the hull damage. A natural approach for modeling such damage is via Gaussian Radial Basis Functions (RBFs) (Stull et al., 2008). The denting damage field is thus modeled as

$$u(\mathbf{x}; \boldsymbol{\theta}) = u_0(\mathbf{x}) + \sum_{j=1}^{n_k} \omega_j \Psi_j(\mathbf{x}), \quad (10)$$

where

$$\Psi_j(\mathbf{x}) = e^{-\left(\|\mathbf{c}_j - \mathbf{x}\|_{\mathcal{R}^2} / \sqrt{2}\sigma_j\right)^2} \quad (11)$$

where u and u_0 describe the finite element nodal coordinates of the damaged model and the original, undamaged coordinates, respectively; n_k is the number of Gaussian RBFs used in the approximation of the damaged configuration; \mathbf{c}_j are the radial

basis centers of each RBF; σ_j are the standard deviations; and ω_j are the Fourier coefficients of the initial imperfection approximation. The corrosion damage field is modeled as

$$\Gamma(\mathbf{x}; \boldsymbol{\theta}) = \Gamma_0(\mathbf{x}) - \sum_{j=1}^{n_k} \omega_j \Psi_j(\mathbf{x}), \quad (12)$$

where Γ is the hull side shell thickness, ω_j corresponds to the maximum reduction in shell thickness; and $\Psi_j(\mathbf{x})$ is given in Eq. (11). While the variables σ_j and ω_j of the Gaussian RBFs are the same for the two damage models, they are interpreted differently, as the damage of each type is included in the finite element model differently. It is these parameters that comprise $\boldsymbol{\theta}$, and that we are attempting to indirectly identify and characterize by “listening” to the structure’s response within the fluid.

2.3.1. Additional parameter uncertainty

In this work, the unknown and sought for parameters $\boldsymbol{\theta}$ are the damage parameters described in Eqs. (10)–(12) (i.e. the RBF kernel centers, standard deviation, and amplitudes for each kernel). While outside the scope of this work, in addition to estimating the structural damage parameters, estimations of the structure and/or fluid material properties are possible, along with the variance of the noise distribution (σ_η) could also be estimated. To facilitate these additional estimates, the parameter vector $\boldsymbol{\theta}$ would need to be augmented to include the additional model parameters. Thus the computationally expensive forward model would potentially need to be evaluated for each of these additional parameters; depending on the variable and the prior imposed on the variable. For example, if one wished to model the uncertainty of the noise distribution, a prior that is conjugate to the likelihood could be used, in order that evaluation of the forward model would not be needed. In any case, increasing the number of parameters to be inferred during the stochastic inversion not only extends the solution time (due to the extra forward model evaluations needed to properly sample a larger model parameter space), but also increases the dimension and complexity of the likelihood function, making representative sampling (or even finding the global maximum) even more difficult, and potentially requiring additional iterations. Therefore, great care should be given in terms of deciding which parameters are of most importance during the inferencing. Estimating material parameter uncertainty, in addition to the currently sought for structural damage parameters (i.e. denting and corrosion), is beyond the scope of this work. Material properties are assumed perfectly known, herein.

2.4. MCMC sampling

Evaluating the conditional posterior in Eq. (8) is difficult, in that it requires evaluating a multi-dimensional integral. Rather than analytically solving for the posterior, it has been shown (Hastings, 1970) that one may numerically draw samples from the posterior PDF. 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 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 value θ_l^* via a probabilistic model $q(\theta_l^* | \theta_l(i))$ which specifies the probability of attaining value θ_l^* given the current value $\theta_l(i)$. Specific choice of $q(\cdot, \cdot)$ is discussed shortly. Both current and candidate values are

then used to evaluate the ratio:

$$r = \frac{f_{\eta}(\mathbf{y}|\theta_l^*)}{f_{\eta}(\mathbf{y}|\theta_l(i))} \times \frac{f_{\pi}(\theta_l^*)}{f_{\pi}(\theta_l(i))} \times \frac{q(\theta_l(i)|\theta_l^*)}{q(\theta_l^*|\theta_l(i))} \quad (13)$$

which is simply the ratio of posterior PDFs (evaluated using the deterministic solver) 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 $\theta_l(i+1) = \theta_l(i)$. The procedure just described draws samples from the posterior distribution under a change in a single parameter, θ_l , within the parameter vector, $\boldsymbol{\theta}$. However we would like to sample from the posterior PDF for each of the $l = 1 \dots P$ parameters. It turns out that we can simply perform the above process (generate candidate, compute r , discard or retain θ_l^*) for each parameter in turn while holding the other $P-1$ parameters fixed (i.e. by using a so-called “Metropolis within Gibbs” approach). Thus, we are really sampling from the conditional posterior distribution $f_{\theta_l}(\theta_l|\theta_1, \dots, \theta_{l-1}, \theta_{l+1}, \dots, \theta_P)$. The ratio in Eq. (13) can be further simplified by using symmetric priors on the parameters.

While the above-described approach is theoretically sound, there are several choices that need to be made in a practical implementation. The first of these is selection of the candidate generating function. In this work, a uniform candidate generating function is used; having the form:

$$q(\theta_l^*|\theta_l(i)) = \frac{1}{2A_l}, \quad |\theta_l^* - \theta_l(i)| < A_l, \quad l = 1, \dots, P, \quad (14)$$

By design this function is symmetric in its arguments, hence it cancels out the final ratio in Eq. (13). Other function choices are certainly possible, however we have found the uniform distribution to work well in our problems. Despite the symmetry, this function still requires the selection of the parameter A_l which can be heuristically thought of as the size of the perturbations used in generating the candidate values θ_l^* . To facilitate fast convergence, a tuning algorithm is employed in the current work, where the distance A_l is modified during the early iterations of the algorithm. These first B iterations are referred to as the “burn-in” phase of the algorithm, and are used to enable the Markov chain to sufficiently settle into a stationary distribution.

Practical methods for tuning A_l involve adjusting the values at each iteration such that the probability of accepting the candidate parameter value, P_{accept} , is optimal (equal to 0.25 in this work, based on the recommendations of Gelman et al., 2004). As described in Link and Barker (2010), A_l is modified as follows:

$$A_l = A_l \times \left(\alpha - (\alpha - 1) \frac{i}{B} \right) \quad (15)$$

if the candidate parameter value is accepted; and A_l is modified as

$$A_l = \frac{A_l}{\left(\frac{1 - P_{\text{accept}}}{1 - P_{\text{accept}} \left(\alpha - (\alpha - 1) \frac{i}{B} \right)} \right)} \quad (16)$$

if the candidate parameter value is rejected. In this work, $\alpha = 1.01$.

2.5. Reversible jump Markov chain Monte Carlo

The previous section describes a sampling method for when the number of Gaussian RBFs 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 $4n_k$. The sampling procedure for the case where the number of kernels (i.e. number of RBFs), n_k , varies may be treated using the reversible jump Markov chain Monte Carlo (RJMCMC) method due to Green (1995). 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. Details can be found in the literature (Green, 1995; Richardson and Green, 1997).

Consider a situation where $\boldsymbol{\theta} \in \boldsymbol{\Theta}_{n_k}$ and $\boldsymbol{\theta}^* \in \boldsymbol{\Theta}_{n_{k+1}}$ so that $m = \dim(\boldsymbol{\Theta}_{n_{k+1}}) - \dim(\boldsymbol{\Theta}_{n_k})$ is the difference in parameter vector cardinality between the original and candidate parameter vectors (assuming $\dim(\boldsymbol{\Theta}_{n_{k+1}}) > \dim(\boldsymbol{\Theta}_{n_k})$). The ratio r in Eq. (13) hardly makes sense in this situation since the posteriors of $\boldsymbol{\theta}^*$ and $\boldsymbol{\theta}$ are not defined on the same space. The original parameter vector must be augmented with a vector $\mathbf{u} \in \mathbb{R}^m$, where $\mathbf{u} \sim g(\mathbf{u})$. Now that the original and candidate parameter vectors are of the same dimension, consider a differential and one-to one mapping $h: \boldsymbol{\Theta}_{n_{k+1}} \rightarrow \boldsymbol{\Theta}_{n_k+1}$ that maps $\boldsymbol{\theta}^* = h(\boldsymbol{\theta}, \mathbf{u})$.

The acceptance ratio given in Eq. (13) is augmented so that the appropriate acceptance probability is

$$r = \frac{f_{\eta}(\mathbf{y}|\boldsymbol{\theta}^*) f_{\pi}(\boldsymbol{\theta}^*) p(n_k \rightarrow n_{k+1})}{f_{\eta}(\mathbf{y}|\boldsymbol{\theta}) f_{\pi}(\boldsymbol{\theta}) p(n_{k+1} \rightarrow n_k) g(\mathbf{u})} \left| \frac{\partial \boldsymbol{\theta}^*}{\partial (\boldsymbol{\theta}, \mathbf{u})} \right|, \quad (17)$$

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

$$r = \frac{f_{\eta}(\mathbf{y}|\boldsymbol{\theta}) f_{\pi}(\boldsymbol{\theta}) p(n_{k+1} \rightarrow n_k) g(\mathbf{u})}{f_{\eta}(\mathbf{y}|\boldsymbol{\theta}^*) f_{\pi}(\boldsymbol{\theta}^*) p(n_k \rightarrow n_{k+1})} \left| \frac{\partial \boldsymbol{\theta}^*}{\partial (\boldsymbol{\theta}, \mathbf{u})} \right|^{-1} \quad (18)$$

In Eqs. (17) and (18), $f_{\eta}(\cdot)$ is the likelihood, defined as in Eq. (7). The second factor, $f_{\pi}(\cdot)$, is the joint prior, which now includes the prior on the model size, n_k , and is of the form $f_{\pi}(n_k, \boldsymbol{\theta})$. The distribution $p(n_k^* \rightarrow n_k)$ is the probability that proposal from n_k^* to n_k is made. The distribution $g(\mathbf{u})$ is used to generate the augmenting vector \mathbf{u} . The last term, $|\partial \boldsymbol{\theta}^* / \partial (\boldsymbol{\theta}, \mathbf{u})|$, is the Jacobian of the mapping h . It is pointed out that the two Eqs. (17) and (18) preserve detailed balance in RJMCMC in an analogous way that Eq. (13) does in standard MCMC.

In addition to the usual Metropolis–Hastings move described in Section 2.4, in which a candidate parameter value is uniformly sampled around the original value $\pm A_l/2$, we may also consider moves that change the number of kernel functions describing the structural damage. 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 an RBF (birth/death); the splitting of an RBF into two kernels; or merging two RBFs into one kernel (split/merge). In an attempt to achieve good acceptance ratios, we only change the value of n_k by at most 1 in each iteration. A description of the aforementioned suite of dimension altering operators is now described; closely following the moves described in Koutsourelakis (2009):

• Birth and death:

For obvious reasons, p_b , the probability of a birth move, is zero when $n_k = n_{k_{\max}}$; and p_d , the probability of a death move, is zero when $n_k = n_{k_{\min}}$.

- Death – A kernel j ($1 \leq j \leq n_k$) is uniformly selected from the n_k RBFs and is removed from the parameter vector.
- Birth – A new kernel $j = n_{k+1}$ is added to the parameter vector, and $\mathbf{c}_{n_{k+1}}$, $\omega_{n_{k+1}}$ and $\sigma_{n_{k+1}}$ are sampled from the prior.

The vector \mathbf{u} (for use in evaluating Eqs. (17) and (18)) is thus $\mathbf{u} = \{\mathbf{c}_{n_{k+1}}, \omega_{n_{k+1}}, \sigma_{n_{k+1}}\}$, and $g(\mathbf{u})$ is just the product of the parameter priors. The Jacobian of the birth/death pair is 1.

• Split and merge:

Similar to the birth/death moves, the probability of a split move, p_s , is zero when $n_k = n_{k_{\max}}$; and the probability of a merge move, p_m , is zero when $n_k = n_{k_{\min}}$.

- Merge: In order to achieve good acceptance ratios, merge moves between any two kernels (j_1 and j_2 , $j_1 \neq j_2$), in the parameter vector θ of dimension n_k , are only allowed if the following two conditions are met:

$$\frac{\|\mathbf{c}_{j_1} - \mathbf{c}_{j_2}\|}{\sqrt{2\sigma_{j_1}^2 + 2\sigma_{j_2}^2}} \leq \delta_c \quad (19)$$

$$|\omega_{j_1} - \omega_{j_2}| \leq \delta_\omega, \quad (20)$$

where values of $\delta_c = 2$ and $\delta_\omega = 0.3$ were used in this work. During a merge move, Eqs. (19) and (20) are evaluated for every possible combination of kernel pairs in the parameter vector. If both conditions are met for a particular pair, then that pair is eligible to be merged. This ensures that the average of (10) (when integrated in \mathbb{R}^2) of the split expansion (i.e. with kernels j_1 and j_2) is the same as that of the merged kernel (i.e. with kernel j in place of kernels j_1 and j_2). A new kernel j is added to the parameter vector, replacing kernels j_1 and j_2 , with the following parameters:

$$\sigma_j = \sqrt{\sigma_{j_1}^2 + \sigma_{j_2}^2} \quad (21)$$

$$\omega_j = \frac{\omega_{j_1}\sigma_{j_1} + \omega_{j_2}\sigma_{j_2}}{\sigma_j^2} \quad (22)$$

$$\mathbf{c}_j = \frac{\mathbf{c}_{j_1} + \mathbf{c}_{j_2}}{2} \quad (23)$$

- Split – The split move must be reversible with respect to the merge move (i.e. an application of the split move must return a merged kernel (j) back to the original two kernels (j_1 and j_2), and vice versa), and thus the parameters of the two new kernels and the parameters of the previous kernel must satisfy Eqs. (19)–(23). In this work, a kernel j is selected randomly from the n_k kernels and is substituted by two new kernels j_1 and j_2 as follows:

- * A scalar u_σ is drawn from $U[0, 1]$, and

$$\sigma_{j_1}^2 = u_\sigma \sigma_j^2, \quad \sigma_{j_2}^2 = (1 - u_\sigma) \sigma_j^2, \quad (24)$$

which satisfies Eq. (21).

- * A vector \mathbf{u}_c is drawn uniformly from a circle of radius R , where $R = \sqrt{2}\delta_c\sigma_j$, and

$$\mathbf{c}_{j_1} = \mathbf{c}_j - \mathbf{u}_c, \quad \mathbf{c}_{j_2} = \mathbf{c}_j + \mathbf{u}_c, \quad (25)$$

satisfying Eqs. (19) and (23).

- * A scalar u_ω is drawn uniformly from $U[-\delta_\omega/2, \delta_\omega/2]$, and

$$\omega_{j_1} = \hat{\omega} - u_\omega, \quad \omega_{j_2} = \hat{\omega} + u_\omega, \quad (26)$$

where $\hat{\omega} = (\omega_j + u_\omega(\sqrt{u_\sigma} - \sqrt{1 - u_\sigma})) / (\sqrt{u_\sigma} + \sqrt{1 - u_\sigma})$. This satisfies Eqs. (20) and (22).

The vector \mathbf{u} is thus $\mathbf{u} = \{u_\sigma, \mathbf{u}_c, u_\omega\}$, and $g(\mathbf{u})$ is the product of the uniform distributions above so that $g(\mathbf{u}) = (1/\pi R^2)1/\delta_\omega$. The Jacobian of the split move is

$$\frac{4\sigma_j}{\sqrt{u_\sigma}(\sqrt{1 - u_\sigma} + \sqrt{u_\sigma})\sqrt{1 - u_\sigma}} \quad (27)$$

For the moves that do not involve changes to n_k , three moves are considered: updates to ω , σ , and \mathbf{c} , and are proposed with probability p_ω , p_σ , and p_c , respectively. If one of these three moves is proposed, a kernel j is selected randomly, and its appropriate parameter is updated. The acceptance ratio in Eqs. (17) and (18) is reduced to the same acceptance ratio as in Eq. (13).

2.6. Convergence criteria

To monitor the convergence of the RJMCMC simulations, a scalar summary of the data, ψ , will be monitored, and we will follow a method outlined in Brooks (1998). As simulation convergence cannot be established with a single chain, $m=10$ independent chains are compared. In a previous work (Reed et al., 2014), the expectations of each parameter value ($E[\theta_i]$) were taken as the scalar summaries of the data, and convergence was monitored for each parameter. Convergence was said to occur when the ratio of the variance of the pooled data (i.e. from all m chains) to the average of the variance of each of the m chains reached a value of 1.1. In that work, however, the number of kernels describing the damage model (n_k) was fixed. Because n_k is to be determined in this work, simply monitoring the expectations of the parameter values will not work, as it is possible that the m different chains may not converge to the same value for n_k (Figs. 5 and 6).

Therefore, a scalar summary of data must be chosen such that it is not sensitive to the number of kernels estimated by the analysis. One choice is to monitor the actual value of the structural damage field described by the parameter vector θ , for a specific set of nodal coordinates chosen *a priori*, see Fig. 7. In this case, the scalar summary would be the expectation of the value of the damage field (i.e. the out of plane displacement or the % reduction in thickness) that results from the parameter vector, and a ratio of the variance of the pooled data to the average variances of the chains could be monitored. However, from past experiences (Reed et al., 2013, 2014), it is known that a wide range of amplitude parameter (ω) values yields a relatively similar likelihood density. Therefore, it is likely that while the m different chains are able to converge to the same damage locations (\mathbf{c}) and standard deviations (σ), the amplitude value of the damage field, governed by ω , may not converge to the same value, and so convergence may not be diagnosed.

Instead of monitoring the expected damage field value at a finite set of points, convergence of the Markov chains will be monitored by considering credible intervals of the damage field, marginalized in the horizontal and vertical directions, see Fig. 8. In




Fig. 5. Schematic of sensor placement within the fluid domain.




Fig. 6. Illustration of birth/death operations (left) and split/merge operations (right).




Fig. 7. Convergence by monitoring the damage field values at a finite set of nodal coordinates.




Fig. 8. Convergence by monitoring the horizontal and vertical credible interval lengths of the damage field.

keeping with the interval length method described in Brooks (1998), convergence is assessed as follows:

1. For each of the m individual chains, calculate the interval lengths of the $100(1-\alpha)\%$ credible intervals of the damage field (in the horizontal and vertical directions).
2. Calculate the interval lengths of the $100(1-\alpha)\%$ credible intervals of the damage field (in the horizontal and vertical directions) for the pooled data (i.e. for all m chains combined).
3. Plot the \hat{R}_h and \hat{R}_v equal to the ratios of the interval lengths of the credible intervals of the pooled data to the average length of the credible intervals of the individual chains in the horizontal and vertical directions, respectively.

A good rule of thumb is to say that convergence is reached when \hat{R}_h and $\hat{R}_v \approx 1.1$ (Brooks, 1998).

2.7. MCMC standard error

In addition to monitoring whether or not the Markov chains are converging to the same distributions, it is important to report the precision of the results. This is accomplished by estimating the Monte Carlo standard error by the batch means method described in Flegal et al. (2008). Suppose the goal of the MCMC analysis is to estimate $E_{f_\theta}[g] \equiv \int_{\mathbb{R}^p} g(\theta) f_\theta(\theta) d\theta$, where g is a real-valued, f_θ -integrable function on θ (Jones et al., 2006). According to the ergodic theorem:

$$\bar{g}_N \equiv \frac{1}{N} \sum_{i=0}^{N-1} g(\theta_i) \rightarrow E_{f_\theta}[g] \quad \text{as } n \rightarrow \infty, \quad (28)$$

with probability 1.

We define the Monte Carlo error as $\bar{g}_N - E_{f_\theta}[g]$, and although we do not know $E_{f_\theta}[g]$, and thus cannot compute the Monte Carlo error exactly, we can assess the Monte Carlo error by estimating the variance on the asymptotic distribution of \bar{g}_N . Under regularity conditions, the error will admit a central limit:

$$\sqrt{N}(\bar{g}_N - E_{f_\theta}[g]) \xrightarrow{d} N(0, \sigma_g^2), \quad (29)$$

where \xrightarrow{d} is defined as convergence in distribution. Therefore, we can assess the precision of the point estimate \bar{g}_N by calculating and reporting an estimate of σ_g^2 (i.e. $\hat{\sigma}_g^2$) (Jones et al., 2006).

The Monte Carlo standard error (MCSE) estimate $\hat{\sigma}_g^2$ will be calculated via the non-overlapping batch means method, in which the output from the Markov chains is divided into a batches of size b such that $N=ab$. The batch means estimate of σ_g^2 is calculated as

$$\hat{\sigma}_g^2 = \frac{b}{a-1} \sum_{j=1}^a (\bar{Y}_j - \bar{g}_N)^2, \quad (30)$$

where

$$\bar{Y}_j \equiv \frac{1}{b} \sum_{i=(j-1)b+1}^{jb} g(\theta_i) \quad \text{for } j=1, \dots, a. \quad (31)$$

We can then form an asymptotically valid confidence interval for $E_{f_\theta}[g]$, where the half-width of the interval is given by

$$h_a = t_{a_{N-1}} \frac{\hat{\sigma}_g}{\sqrt{N}} \quad (32)$$

where $t_{a_{N-1}}$ is an appropriate quantile from the Student's t distribution with a_{N-1} degrees of freedom (Flegal et al., 2008).




Fig. 9. Iterative k -means clustering algorithm. (a) Unassigned damage centers \mathbf{c} . (b) Cluster centers ζ initialized randomly. (c) Damage centers assigned to their nearest cluster center. (d) Cluster centers recomputed to be the mean of the damage centers previously assigned to the cluster. (e) Damage centers assigned to their nearest cluster center. (f) Cluster centers recomputed to be the mean of the damage centers previously assigned to the cluster. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

Table 1
Parameters of prior distributions of the damage parameters.

Parameter	Geometric imperfection	Corrosion patch
σ_{min}	19.0 in	19.0 in
σ_{max}	66.5 in	66.5 in
ω_{min}	0.24 in	0.09 in
ω_{max}	2.16 in	0.81 in




Fig. 10. Support of dent center parameters (\mathbf{c}). (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)




Fig. 11. Natural log of the prior probability of model size, n_k .

The choice of $g(\theta)$ to be used in this work is not trivial, as the various chains are not guaranteed to converge to the same estimation of the number of Gaussian kernels in the parameter vector θ . Therefore, the choice of g should be something that is not very sensitive to the number of kernels, and it should provide a meaningful summary of the results, or else there is very little point in calculating the error associated with it. As the goal of this work is to identify the parameters that describe the structural damage




Fig. 12. Ground truth damage field for the denting imperfection on bow structure model.

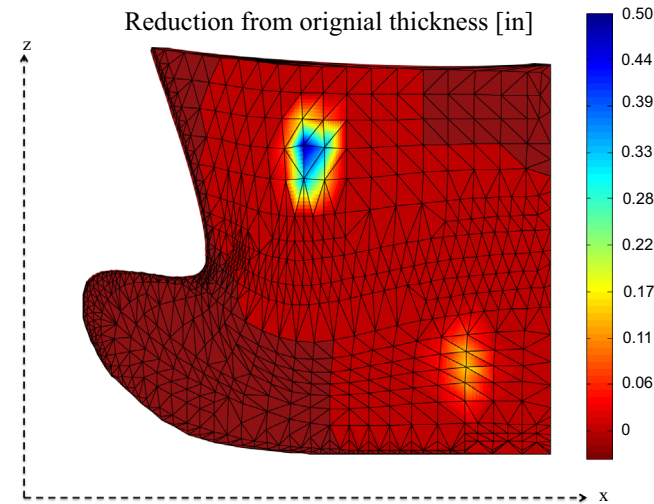


Fig. 13. Ground truth damage field for the corrosion patch on bow structure model.

Table 2
Ground truth parameters: denting imperfection.

Kernel	c_1	c_2	ω	σ
$j=1$	855.0	190.0	1.20	23.8
$j=2$	0.000	−190.0	1.92	47.5

Table 3
Ground truth parameters: corrosion patch.

Kernel	c_1	c_2	ω	σ
$j=1$	380.0	237.5	0.77	37.5
$j=2$	902.5	−166.4	0.14	47.5

condition within a ship bow section, $g(\theta)$ should be chosen to describe the damage, invariant of the size of the parameter vector. In this work, g will be the center points of the “clusters” of damage (i.e. damage neighborhoods within the hull side shell). The calculation of the cluster centers is described in Section 2.7.1.

2.7.1. K-means clustering

It is very likely that the parameter vectors that result from the Markov chains will have damage centers, \mathbf{c} , that “cluster” together around certain points. When the damage field is computed using the parameter vector estimated during the inverse solution, the resulting field may appear to have only a few damage sites, despite an estimated relatively large number of Gaussian kernels, n_k . For example, if n_k is estimated to be 10, but 5 of the kernels are more or less “on top of each other” forming a “cluster” that manifests as one large damage site, then we seek an algorithm that will quantify how many clusters (i.e. damage sites) there are. Using the K-means clustering algorithm described in Bishop (2006), we will identify the center of the damage clusters as $g(\theta)$ and report its Monte Carlo standard error.

The goal of the k -means clustering algorithm is to find the n_c cluster centers ζ (where n_c is defined as the number of clusters) that minimize the following objective function:

$$J = \sum_{j=1}^{n_k} \sum_{k=1}^{n_c} r_{jk} \|\mathbf{c}_j - \zeta_k\|^2 \quad (33)$$

where r_{jk} is a binary indicator variable such that a value of 1 indicates that kernel center \mathbf{c}_j is assigned to cluster k . Minimizing Eq. (33) is done via an iterative two-step process, first minimizing Eq. (33) with respect to r_{jk} and then with respect to ζ_k . For example, Fig. 9a represents $n_k=7$ estimated kernels. It is clear that there are $n_c=2$ clusters, which the k -means clustering algorithm will determine, along with the cluster centers. The k -means clustering algorithm works as follows:

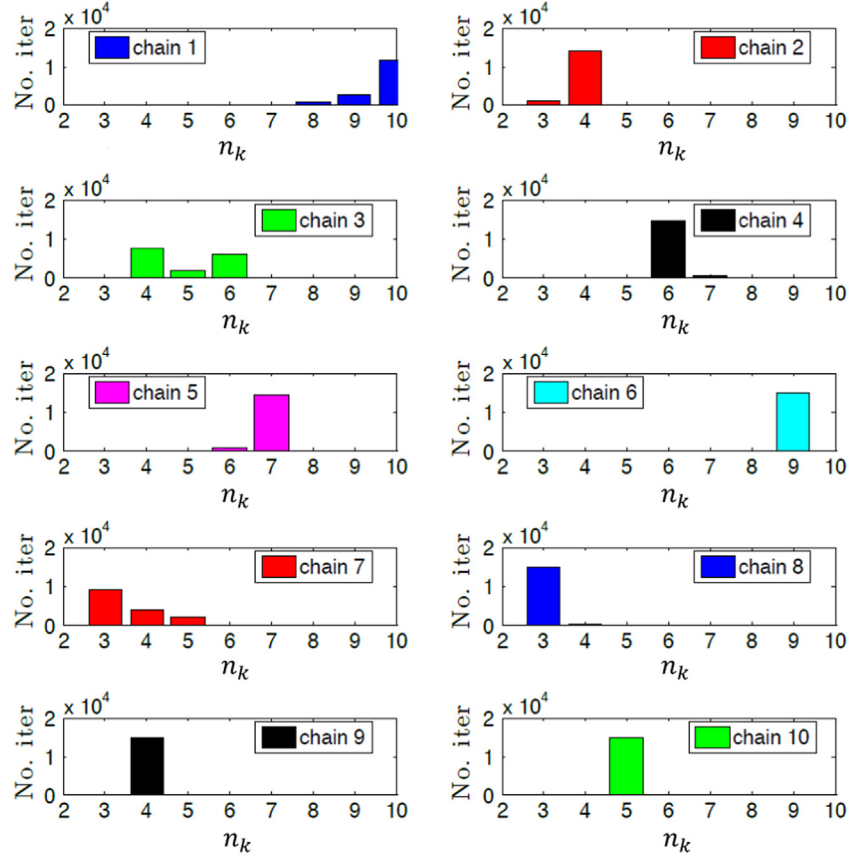
First, some initial starting values for ζ_k are chosen, see Fig. 9b, where the two colors (red and blue) represent different clusters. Then a line is drawn that is halfway between the cluster centers and orthogonal to a line connecting the cluster centers. Each damage center is then assigned to its nearest cluster center (Fig. 9c) by determining which side of the purple line it falls on, minimizing Eq. (34) with respect to r_{jk} . Next, the ζ_k are recomputed to be the mean of the damage centers assigned to that cluster (see Fig. 9d), minimizing Eq. (33) with respect to ζ_k . This process is repeated until Eq. (33) is minimized (Bishop, 2006).

Because the number of damage clusters, n_c , is not known *a priori*, we need a means of determining the appropriate number of damage clusters from the Markov chain output. In keeping with the principle of “Occam’s Razor”, this is accomplished by selecting the minimum number of clusters such that the weighted average distance, \bar{d}_k , between the cluster center and each of the \mathbf{c}_j assigned to cluster k is less than $1.5\sigma_{max}$, where σ_{max} is the upper bound on the prior for the damage standard deviation parameter, see Table 1. The weighted average distance, \bar{d}_k , is the distance weighted by both the amplitude and standard deviation associated with each of the damage centers. This is done to ensure that the contribution to \bar{d}_k from a damage center relatively far from a computed cluster center does not significantly increase the \bar{d}_k , if the amplitude and standard deviation associated with the far away damage center are relatively small. The distances weighted by the standard deviation alone are

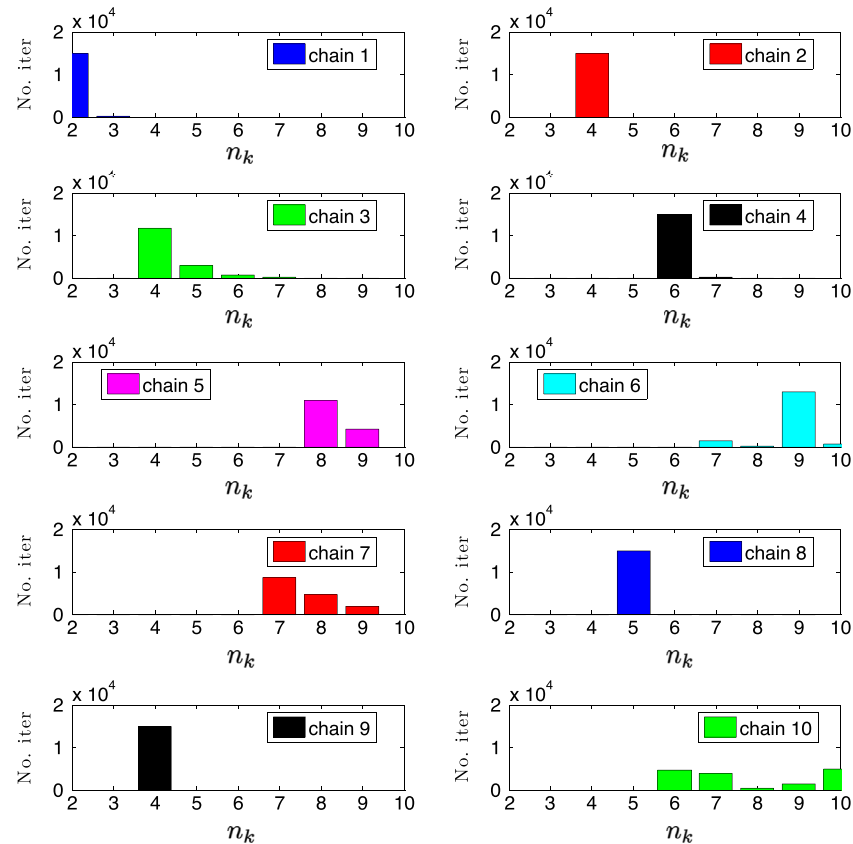
$$d_j^\sigma = \frac{\sigma_i d_{lk} r_{jk} \delta_{il}}{\sum_{n=1}^{n_c} r_{mn} \sigma_m}, \quad (34)$$

where σ_i are the damage standard deviations; d_{lk} are the distances from damage center l to cluster k ; r_{jk} is the binary indicator variable defined before; δ_{il} is the Kronecker delta, equal to 1 when $i=l$; and the denominator is the sum of the standard deviations associated with each cluster. Eq. (34) is a vector of σ -weighted distances from the damage centers to the assigned cluster centers.

Histogram of model size – denting imperfection

Fig. 14. Histogram results for model size n_k for each of the 10 chains: denting imperfection.

Histogram of model size - corrosion

Fig. 15. Histogram results for model size n_k for each of the 10 chains: corrosion patch.




Fig. 16. Resulting damage fields for each of the $m=10$ chains for the denting imperfection (left) and corrosion patch (right). The first row contains chains 1 and 2 (from left to right); the second row contains 3 and 4 (from left to right); etc.

The weighted average distance, \bar{d}_k , is now computed by weighting Eq. (34) with the appropriate damage amplitudes:

$$\bar{d}_k = \frac{(\omega \odot \mathbf{d}^\sigma)_j r_{jk}}{\sum_{n=1}^{n_c} r_{mn} \omega_m}, \quad (35)$$

where $(\omega \odot \mathbf{d}^\sigma)_j$ denotes the j th component of the Hadamard product (i.e. the component wise product) of the amplitude and the vector in Eq. (34); and the denominator is the sum of the amplitudes associated with each cluster.

3. Results

The aim of this paper is to demonstrate the ability to identify and characterize relatively small dent-like imperfections and corrosion patches in a multi-physics environment, through the solution of a stochastic inverse problem. The centers of the structural damage (\mathbf{c}) are assumed to be within the red lines in Fig. 10, where the top red line coincides with the free surface of the fluid domain, and \mathcal{D} = area enclosed by the red lines. The bounds on the standard deviation (σ) of the damage (both the

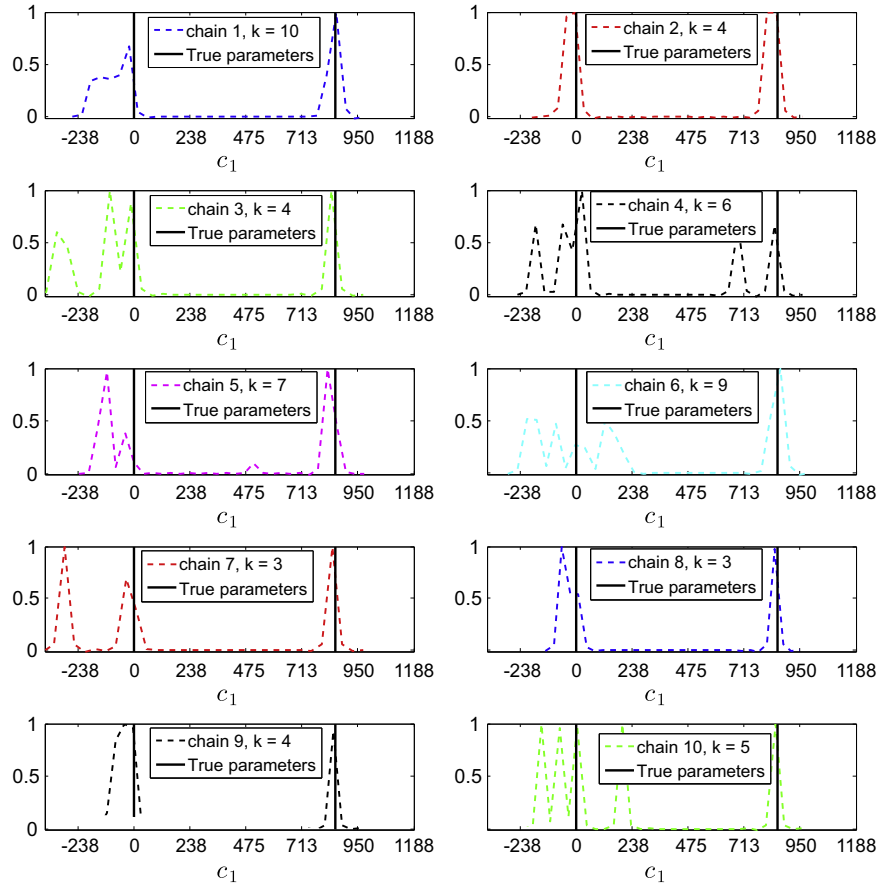


Fig. 17. Kernel density estimates of the c_1 parameter for the most probable model: denting imperfection.

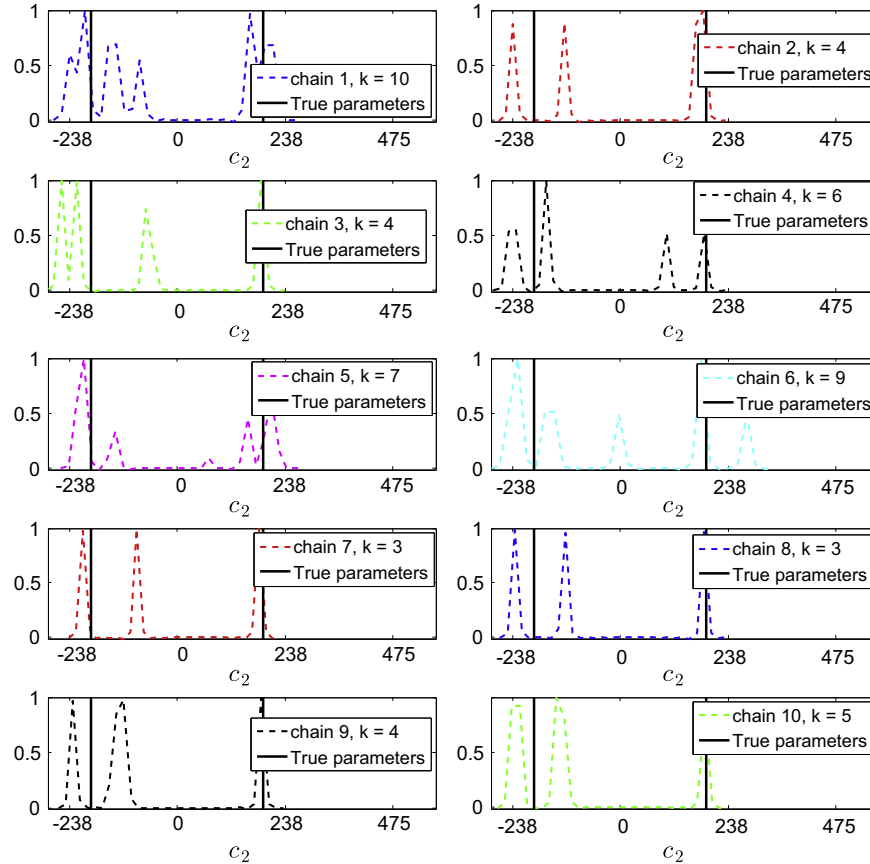


Fig. 18. Kernel density estimates of the c_2 parameter for the most probable model: denting imperfection.

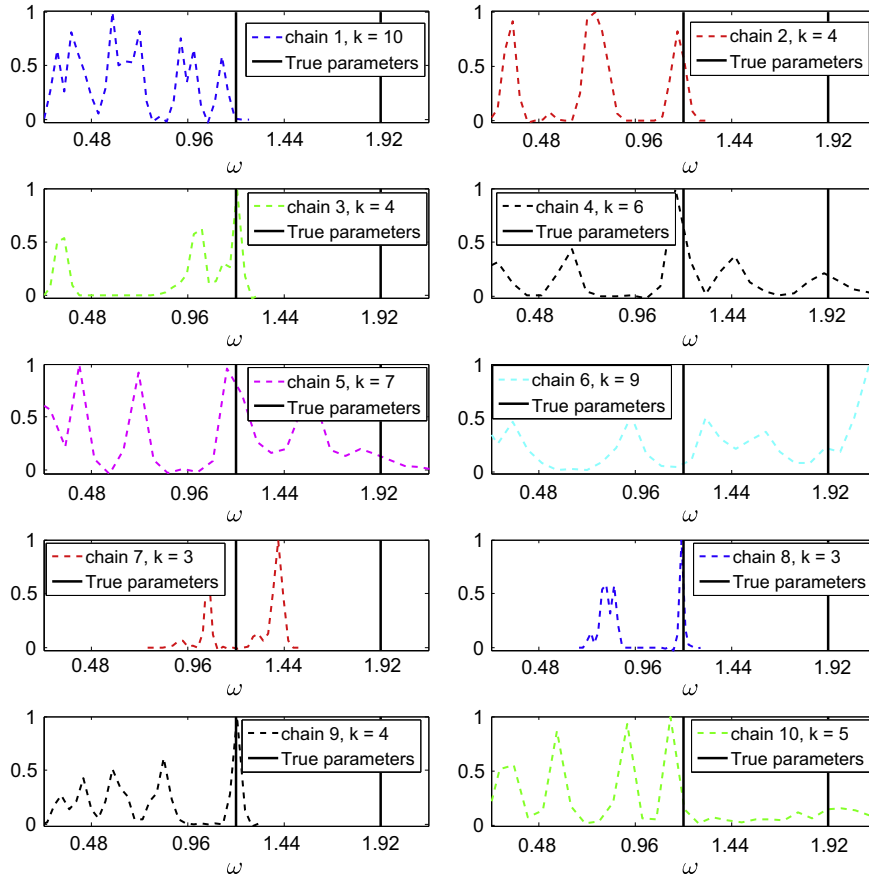


Fig. 19. Kernel density estimates of the ω parameter for the most probable model: denting imperfection.

denting imperfection and the corrosion patch) are chosen such that the lower bound yields damage that can be resolved within the mesh, and the upper bound is small enough that the damage still remains “small”, relative to the overall structure characteristic length. The amplitude parameter (ω) for the geometric imperfection is chosen such that the minimum and maximum amplitudes of a dent (corresponding to the out of plane displacement) are approximately 25% and 250% of the thickness of the structure, respectively. The minimum and maximum amplitudes of the corrosion damage (i.e. the % reduction in thickness of the structure) correspond to 10% and 90% reduction in side shell thickness, respectively. The priors on the parameters are uniforms:

$$\mathbf{c} \sim \frac{1}{|\mathcal{D}|}, \quad \mathcal{D} = \text{area of damage center support} \quad (36)$$

$$\sigma \sim U[\sigma_{\min}, \sigma_{\max}] \quad (37)$$

$$\omega \sim U[\omega_{\min}, \omega_{\max}], \quad (38)$$

where the prior parameter values are given in Table 1.

The priors for the parameter values are independent of the model size. The possible number of kernels is allowed to vary between $n_{k_{\min}} = 1$ and $n_{k_{\max}} = 20$. In addition to priors for the parameter values, a prior for the model size, n_k , is required, since the model size is unknown. Consistent with the Occam's razor heuristic (Jefferys and Berger, 1992; Rasmussen and Ghahramani, 2001), a prior for n_k is chosen such that fewer kernels are favored over more kernels, following the prior in Koutsourelakis (2009). In this work,

$$p(n_k | \lambda) \propto e^{-\lambda \frac{n_k}{n_k}} \quad \forall n_k : n_{k_{\min}} \leq n_k \leq n_{k_{\max}}, \quad (39)$$

where an exponential hyper-prior is used for hyper-parameter λ : $p(\lambda | s) = s e^{-\lambda s}$. Integrating out λ , the prior for n_k becomes

$$p(n_k | s) \propto \frac{1}{(s+1)^{n_k+1}}, \quad \forall n_k : n_{k_{\min}} \leq n_k \leq n_{k_{\max}}, \quad (40)$$

where $s = 10$ in this work (see Fig. 11).

The “ground truth” damage fields to be identified are presented in Figs. 12 and 13. The parameters that describe the damage fields are given in Tables 2 and 3, respectively. In both cases, $n_k = 2$.

In an attempt to ensure that the set of 10 Markov chains begin at “over dispersed starting points” (Gelman and Rubin, 1992; Brooks, 1998) relative to the posterior density, a mode-finding algorithm is employed to locate regions of high likelihood. In this work, a type of evolutionary algorithm (EA) known as the modified differential evolution algorithm (DE), described in Reed et al. (2013), is employed to find starting points for 4 of the 10 Markov chains. The other 6 chains' starting points are selected uniformly from the parameter support. In keeping consistent with Occam's razor, the chains initially assume $n_k = 1$, and are allowed to increase according to the RJMCMC rules. Based on literature describing the noise levels within various ships (Wiggins et al., 2012) and ambient sea noise (Frisk, 2012), a signal to noise ratio (SNR) of 1000 is used to generate the zero mean, Gaussian additive noise assumed by the signal model in this work. The chains are run for $N = 20,000$ iterations, and the first 5000 are discarded as burn-in.

The results of the estimated model size (n_k) are presented for each chain in Figs. 14 and 15, for both damage cases. It is clear that the chains do not all converge to the same number of kernels, and only one chain actually predicts the correct number of kernels ($n_k = 2$). This result is discussed in detail in Section 3.1. In a Markov chain Monte Carlo analysis, in which n_k is assumed to be known, the results can be typically summarized by assembling the damage

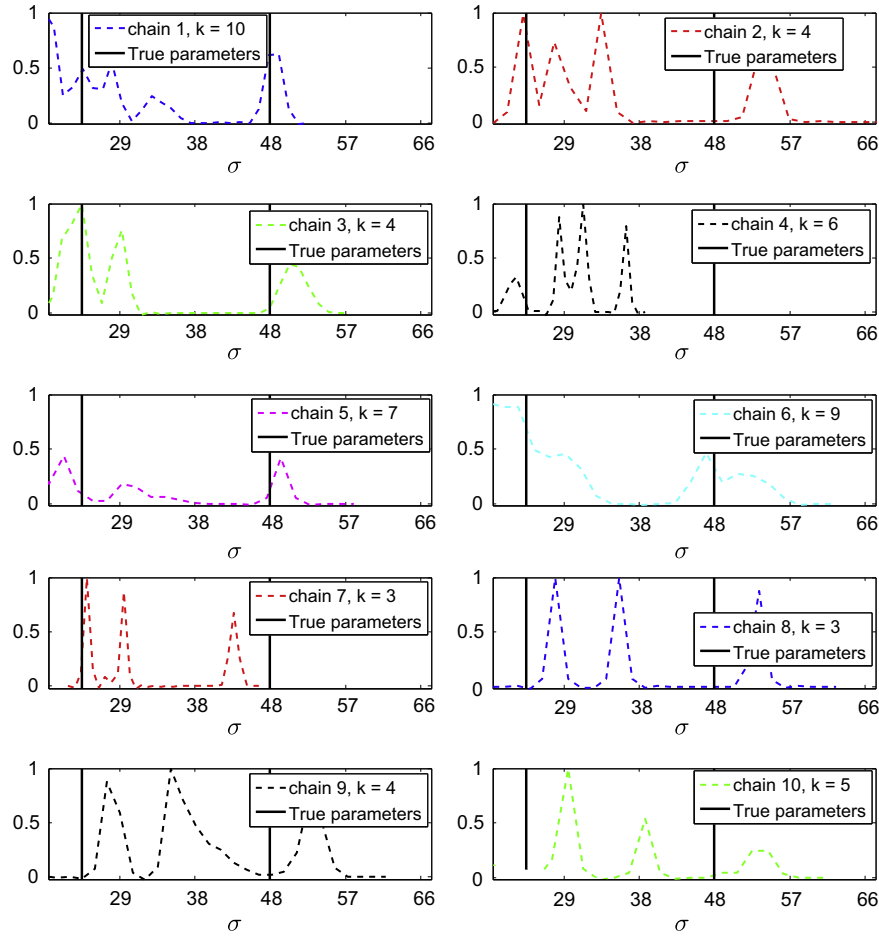


Fig. 20. Kernel density estimates of the σ parameter for the most probable model: denting imperfection.

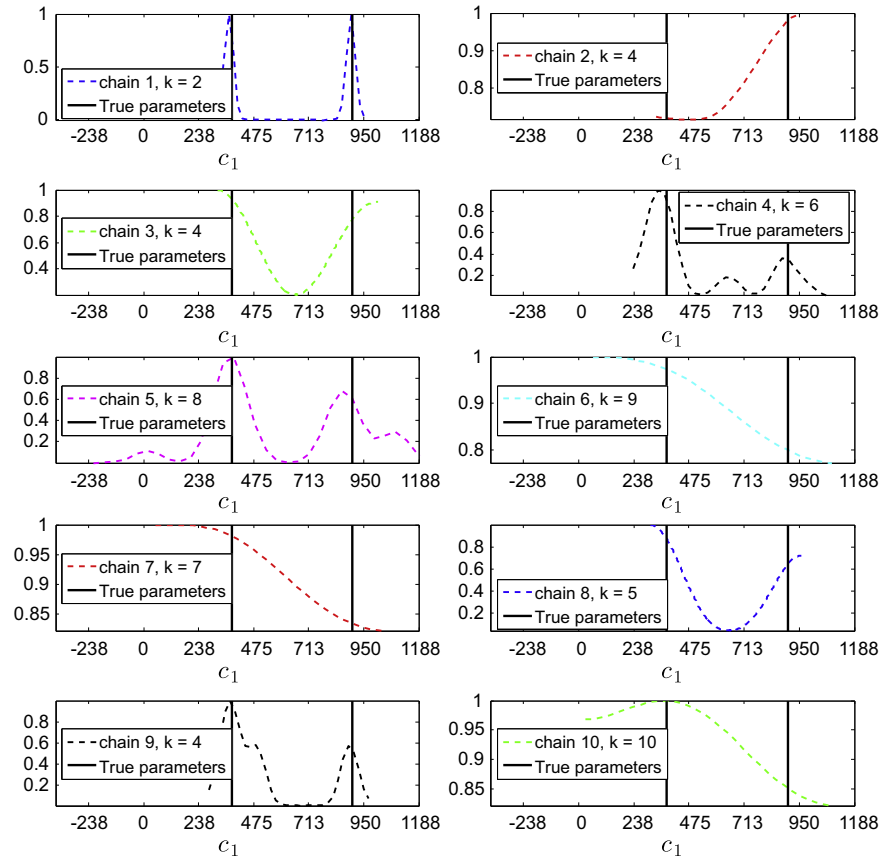


Fig. 21. Kernel density estimates of the c_1 parameter for the most probable model: corrosion patch.

field that results from considering the mean or median parameter values from the output. However in the RJMCMC analysis, simply taking averages of the parameters does not summarize the damage field (i.e. Eq. (10)) since the damage field is calculated by summing the contribution from each of the Gaussian kernels. Therefore, to visualize the results, the damage field landscapes are presented for each of the 10 chains in Fig. 16. The parameters used to generate these damage landscapes are taken from the very last iteration of each chain. The chains are assumed to have converged well before then, and so the last iteration can be assumed to be a random sample from the highest probability region within the posterior, and is thus an appropriate representation of the resulting damage field.

Kernel density estimates (KDEs) are calculated so that the uncertainty surrounding each parameter may be visualized. The KDEs are calculated for the most probable model sizes for each chain, as shown in the histograms in Figs. 14 and 15. The kernel density estimates of the marginal posteriors conditioned on the model size (i.e. $f_{\theta}(\theta|\mathbf{y}, n_k)$) are presented in Figs. 17–24.

As discussed in Section 2.6, the chain convergence results (the iterative values of \hat{R}_h and \hat{R}_v) are shown in Figs. 25 and 26. The \hat{R} statistic is called the “potential scale reduction factor” and describes the amount one could expect the pooled posterior density to decrease in the event that $N \rightarrow \infty$. For example, for an $\hat{R} = 1.2$, the width of the pooled posterior could be expected to decrease by approximately 20%. It is also suggested to monitor the behavior of the numerator and denominator of \hat{R} to ensure that those values are stabilizing. If they continue to change, it could indicate that at least one of the chains is still exploring the parameter space, and suggest that convergence has not been reached. For the geometric imperfection case (i.e. denting), while \hat{R} in the horizontal direction as well as the associated numerator and denominator (the top two figures in Fig. 25) appear to stabilize, the value of \hat{R} is still relatively high. The horizontal \hat{R} appears to have reached an acceptable value, although it would appear that it is still stabilizing. For the corrosion patch damage case (i.e. Fig. 26), both the horizontal and vertical \hat{R} have reached acceptable values. However, the numerators and denominators have yet to stabilize. This may be a sign that more iterations are needed in order to state that the individual chains have converged to the posterior distribution.

In addition to monitoring whether or not the Markov chains are converging to the same distributions, it is important to report the precision of the results. This is accomplished by estimating the Monte Carlo standard error associated with the n_c cluster centers ζ , by the batch means method described in Section 2.7. The number of clusters is determined for each chain by calculating the minimum n_c such that the weighted average distances for each cluster defined in Eq. (35) are less than 1.5 times the maximum allowable standard deviation parameter. As can be seen in Fig. 27, the majority of the chains estimate two damage clusters, and therefore in order to compare the MCSE for each of the chains' cluster centers, n_c will be assumed to be two for both damage scenarios.

The cluster centers, ζ , are displayed in Fig. 28 for each chain. While it is visually clear that the cluster centers appear to be precise with respect to the other chains' cluster centers, it is important to also provide the MCSE associated with each cluster center (i.e. h_a defined in Eq. (32)). This information is provided in Table 4 for the geometric imperfection and Table 5 for the corrosion patch, where the subscript on ζ and h indicates the horizontal (1) and vertical (2) parameters, and the superscript indicates which cluster.

While most of the cluster centers appear to be relatively close to the ground truth parameters, there appears to be a very large range in the MCSE (as much as 4 orders of magnitude) between

some of the chains. Chain 1 of the corrosion patch damage case is the only chain that correctly estimates two Gaussian kernels in the damage model. Therefore it is not a surprise that the MCSE associated with the cluster centers of that chain is the lowest. This may be because there are only two kernels (eight total parameters) to estimate, and so all of the iterations have been focused on those parameters, as compared to a chain that is trying to estimate the parameters associated with nine kernels. Additionally, the high MCSE associated with chains that predict a relatively large number of kernels could be due to “over-fitting” (i.e. the parameters fit to the noise present in the observed pressure field). It is likely that chains with several kernels will need far more iterations to be able to accurately identify the parameters. Exceptionally high MCSEs could also be explained by a lack of convergence of model size. For example, for chain 5 of the corrosion damage case, the histogram for n_k (Fig. 15) appears to have most of its mass for $n_k=8$, but also a significant portion of values in $n_k=9$. It is possible that this switch in model size significantly changes the value of the cluster center so that the error becomes greater. For the chains that have not converged to a single values of n_k , additional iterations would likely decrease the MCSE associated with the cluster centers.

3.1. Discussion of results

Almost every chain for both damage cases overestimates the number of Gaussian kernels n_k that describe the damage, despite the prior on n_k favoring simpler models. Additionally, there is a large variation in the predicted number of kernels among the different chains. Despite a high value of n_k , the kernel density estimates (KDEs) of the (damage center) parameter values seem to qualitatively suggest relatively few modes. This is because the kernels tend to be centered around a damage “cluster”, and several small kernels are used to represent a single damage site. This also explains the estimation of the amplitude parameters, which appear to be significantly under estimated. Several kernels with relatively small amplitudes appear to look the same (with respect to the damage landscape) as a single kernel with a large amplitude. For chain 1 of the corrosion damage (the only case when n_k is correctly estimated to be 2), all of the parameters' KDEs include the ground truth values in the density estimates, and the variance associated with each mode is relatively small. This is likely due to two reasons: first, the chain has correctly identified the number of kernels that was used to generate the noisy data; and secondly, since there are only two kernels, more iterations are spent sampling from fewer parameters. Also, the relative heights of the different modes in each of the KDEs are a reflection of the number of kernels centered about each cluster and not the probability of that parameter value. For example, a chain that estimates $n_k=3$, where two kernels are centered around cluster 1 and the other kernel is centered near cluster 2, would show a KDE where the damage center density appears to be higher for cluster 1 because 2/3 of the damage center parameters appear to be at cluster 1 versus cluster 2. While the KDEs are useful for estimating the uncertainty surrounding the parameter estimates (i.e. by observing the variance surrounding each mode), the damage landscapes in Fig. 16 are the best way to visualize the actual damaged condition of the structure.

Convergence of the Markov chains is determined by considering the interval lengths of the $100(1-\alpha)\%$ credible intervals associated with the damage landscapes, marginalized in the horizontal and vertical directions. The convergence monitoring results (Figs. 25 and 26) would indicate that more iterations are necessary to determine that all of the chains have converged. While most of the \hat{R} s appear to be of an acceptable value, the numerators and denominators do not appear to have stabilized.

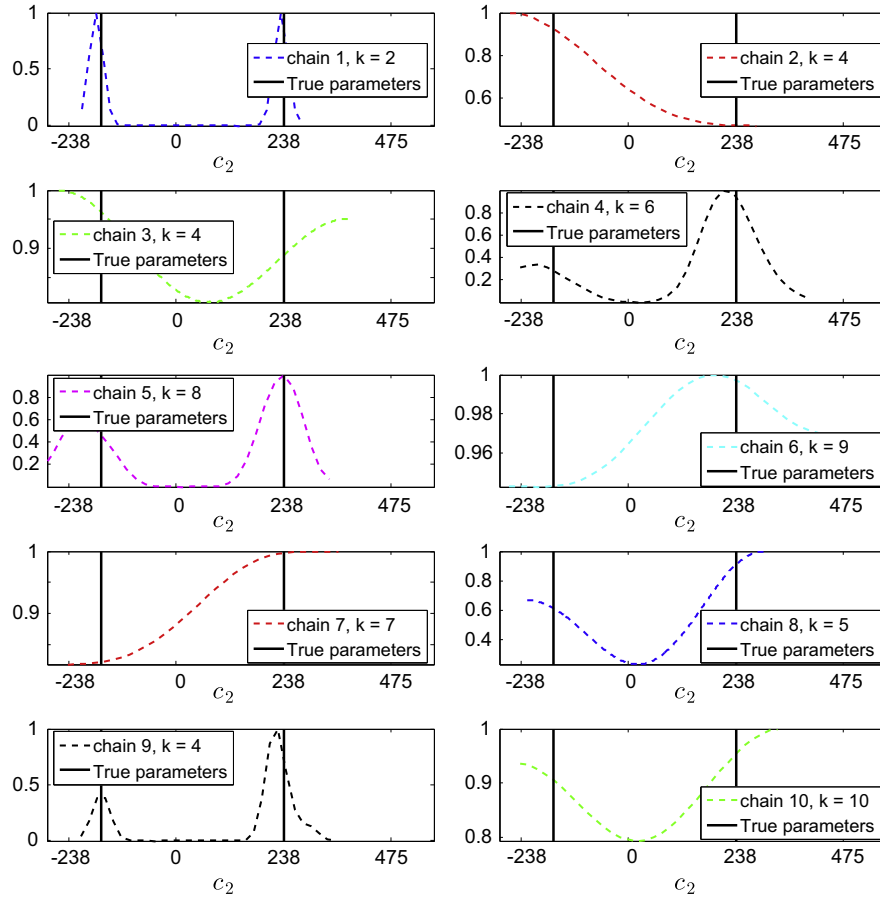


Fig. 22. Kernel density estimates of the c_2 parameter for the most probable model: corrosion patch.

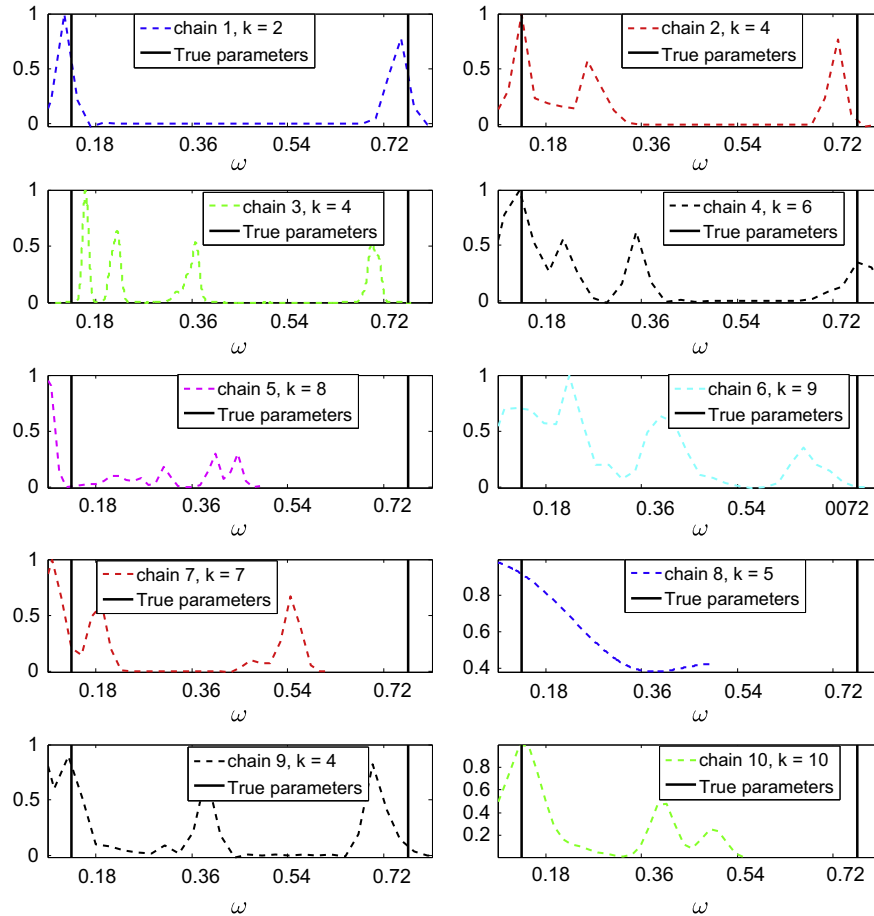


Fig. 23. Kernel density estimates of the ω parameter for the most probable model: corrosion patch.

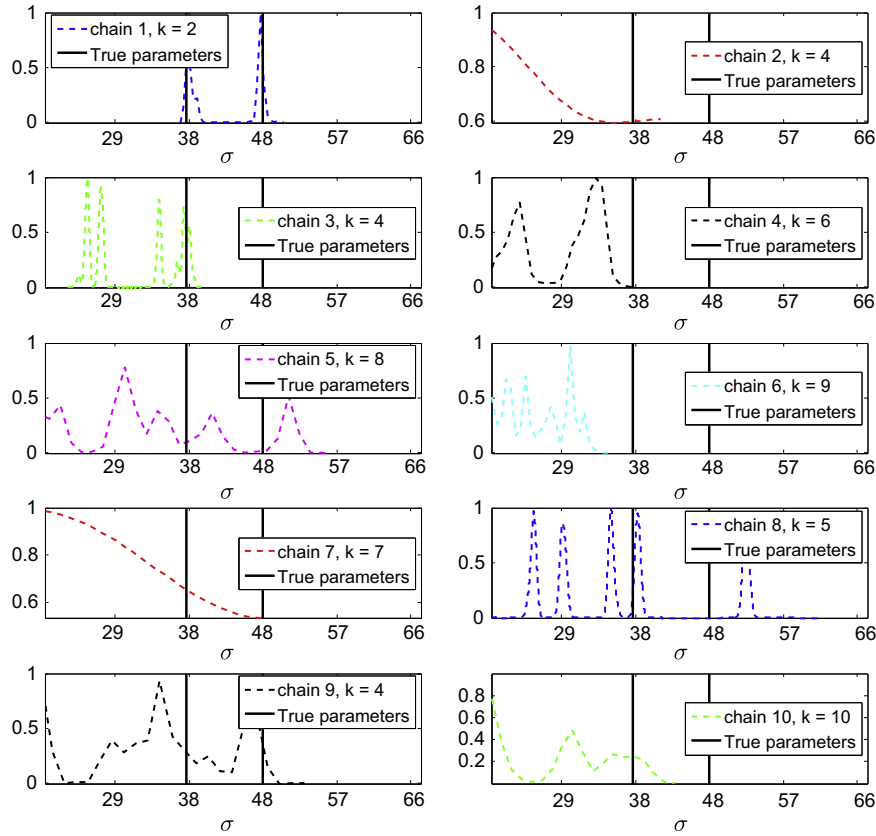


Fig. 24. Kernel density estimates of the σ parameter for the most probable model: corrosion patch.




Fig. 25. Iterative \hat{R} (left). Pooled and average (over all chains) credible interval lengths (right): denting imperfection

This is an indication that at least one of the chains is still exploring the parameter space, and the analysis could benefit from more iterations. Another likely scenario that would result in parameters

continuing to explore the parameter space is that there are several parameter vector combinations that result in qualitatively the same damage vector landscape field, and thus the same posterior




Fig. 26. Iterative \hat{R} (left). Pooled and average (over all chains) credible interval lengths (right): corrosion patch.

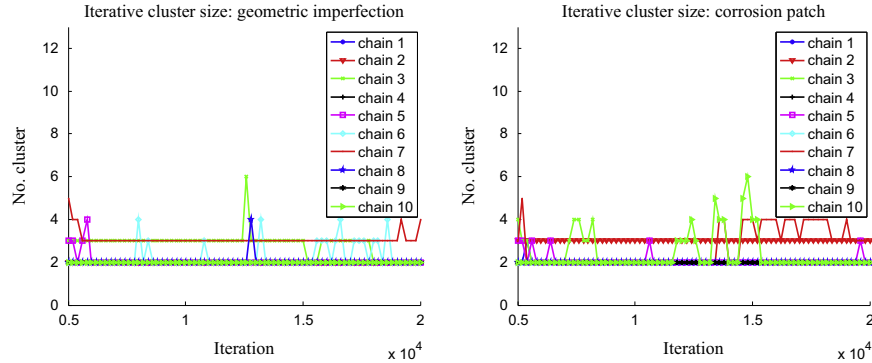


Fig. 27. Estimated number of clusters per iteration: geometric imperfection (left); corrosion patch (right).




Fig. 28. Cluster center estimation for each chain: denting imperfection (left); corrosion patch (right).

probability. For this reason, parameters could easily continue to “move” in the parameter space and still result in a stable, converged landscape. The total analysis time of the inverse solution is proportional to the number of times the deterministic solver (i.e. the finite element model) must be evaluated. In this work, the deterministic solver takes approximately two minutes to evaluate. Thus, the total inverse solution time for the $N=20,000$

iterations in this work is approximately 28 days using a Mac Pro workstation with dual 2.66 GHz, 6-core Intel Xeon processor, with 64 GB, 1333 MHz DDR3 memory.

In addition to monitoring how the chains are converging, it is also important to monitor how precise the results are. By computing the Monte Carlo standard error associated with the $n_c=2$ damage clusters, the precision of the results can be assessed, independent

Table 4
Monte Carlo standard error (MCSE) associated with cluster centers: denting imperfection.

Chain Number	ϵ_1^1	$h_{a_1}^1$	ϵ_2^1	$h_{a_2}^1$	ϵ_1^2	$h_{a_1}^2$	ϵ_2^2	$h_{a_2}^2$
Chain 1	862.3169 \pm	0.0100	194.0465 \pm	0.0769	–64.0215 \pm	1.4222	–162.1631 \pm	1.1552
Chain 2	847.0737 \pm	0.0157	185.0035 \pm	0.0204	–14.0382 \pm	1.2521	–172.8098 \pm	7.1602
Chain 3	856.5072 \pm	1.3091	189.0277 \pm	5.4031	–116.3731 \pm	3.1996	–158.3360 \pm	5.0792
Chain 4	782.6832 \pm	0.3695	149.6564 \pm	0.2822	–33.1070 \pm	0.3410	–188.9018 \pm	0.0427
Chain 5	833.4483 \pm	35.0374	191.1543 \pm	4.5049	–80.7234 \pm	1.4526	–189.4566 \pm	1.8957
Chain 6	861.3940 \pm	0.0057	219.3179 \pm	0.0119	–1.9907 \pm	0.0128	–159.8826 \pm	0.0670
Chain 7	861.3679 \pm	7.7976	190.4057 \pm	3.4319	–117.6627 \pm	7.4708	–136.8133 \pm	0.0784
Chain 8	855.3524 \pm	0.2247	189.3079 \pm	1.0963	–15.4560 \pm	0.0114	–168.4165 \pm	0.0256
Chain 9	855.1995 \pm	0.0128	189.7796 \pm	0.0033	–28.1276 \pm	0.2095	–151.6575 \pm	0.0964
Chain 10	855.1867 \pm	0.0157	190.8526 \pm	0.0076	10.0700 \pm	0.0485	–171.9647 \pm	0.0613
Ground truth	855.0		190.0		0.0		–190.0	

Table 5
Monte Carlo standard error (MCSE) associated with cluster centers: corrosion patch.

Chain Number	ϵ_1^1	$h_{a_1}^1$	ϵ_2^1	$h_{a_2}^1$	ϵ_1^2	$h_{a_1}^2$	ϵ_2^2	$h_{a_2}^2$
Chain 1	902.3732 \pm	0.0014	–165.4962 \pm	0.0038	381.1226 \pm	7.9966	237.3565 \pm	0.9776
Chain 2	798.2893 \pm	0.0494	–147.9264 \pm	0.0641	384.3847 \pm	0.0000	236.9728 \pm	0.0000
Chain 3	907.2942 \pm	7.3886	–147.1493 \pm	73.2863	409.0779 \pm	38.3311	277.6423 \pm	0.9709
Chain 4	904.2309 \pm	22.7159	–171.3221 \pm	11.4656	401.5084 \pm	18.1958	233.6658 \pm	6.4325
Chain 5	929.1694 \pm	22.4418	–124.7958 \pm	89.2820	347.4936 \pm	214.5603	163.4043 \pm	171.8645
Chain 6	901.3320 \pm	2.2567	–166.4690 \pm	1.9504	379.4287 \pm	14.3692	204.7578 \pm	2.7559
Chain 7	921.5432 \pm	6.8637	–102.2599 \pm	197.0143	378.5490 \pm	5.6658	189.3188 \pm	13.4140
Chain 8	902.0716 \pm	0.0993	–152.7505 \pm	0.0884	376.7994 \pm	0.0537	230.3622 \pm	0.0337
Chain 9	897.4403 \pm	4.0936	–157.4407 \pm	2.0886	424.4014 \pm	–1.0122	239.9548 \pm	2.5674
Chain 10	795.4241 \pm	24.2369	–138.1120 \pm	10.7412	353.2642 \pm	16.3281	210.2065 \pm	17.8861
Ground truth	902.5		–166.4		380.0		237.5	

of the estimated number of kernels. While the estimated cluster centers appear to be accurate relative to the ground truth values (see Tables 4 and 5), the precision varies significantly from chain to chain. The MCSE also appears to be higher for the chains whose histograms of n_k estimate more than one value. For example, for the denting imperfection, chains 3 and 7 (Fig. 14) estimate three different values for n_k , and the MCSE associated with those chains is relatively high (on the order of 10^{-1}). Similarly for the corrosion patch, the histogram of the model size for chain 10 (Fig. 15) estimates 5 different values of n_k , and the MCSE associated with chain 10 is the highest.

Because convergence of the chains, as well as the Monte Carlo standard error, appears to be a function of n_k , a method of minimizing the number of kernels should be employed. In this work, a prior is placed on n_k that favors fewer kernels. Another option could be to run a preliminary RJMCMC analysis and use the k -means clustering algorithm to estimate the number of damage clusters. The k_{max} could be reduced based on the number of estimated clusters. Having to estimate the values of fewer parameters allows for faster convergence and smaller errors.

4. Conclusions

In this work we have described both the modeling and estimation tools needed to successfully identify and characterize dent-like imperfections as well as section loss damage typical in corrosion within a submerged ship bow structure, in a non-contact manner, by using hydrophones. A Bayesian approach to damage parameter estimation is adopted, whereby the unknown parameters describing the damage are treated as random variables; each described by a probability density function. It is shown how numerical samples from these functions can be obtained and subsequently used to estimate the error associated with damage parameters, independent of n_k .

This work demonstrated the feasibility of a reversible jump Markov chain Monte Carlo (RJMCMC) approach, for when a practitioner is not confident in specifying a damage model (in this case the number of RBFs in the damage model), allowing for a robust damage model. In the RJMCMC approach, much of the implementation remains similar to that of an MCMC analysis. However, moves that change the cardinality of the parameter vector must be made reversible, with respect to the irreducibility, aperiodicity, and non-transience of the Markov chains. Therefore a mapping between parameter vectors of different dimensions is introduced, and this mapping (i.e. diffeomorphism) is included in the acceptance ratio by way of a Jacobian matrix. By analyzing the posteriors of the parameter values, conditioned on the model size, we may gain an understanding of the probability of the various parameter values with respect to each model.

In the RJMCMC analysis in this work, 10 independent Markov chains are generated for two structural damage scenarios for sampling from the posterior distribution. Using more than one chain allows for monitoring whether or not the chains have converged to the stationary, posterior distribution. Calculating the Monte Carlo standard error of the damage parameters allows for estimating how precise the solutions are, a feature that is not inherent in a maximum likelihood estimation approach.

The Bayesian analysis foundation in the current approaches provides both parameter estimates and quantify the uncertainty in the predictions.

Acknowledgements

The Cornell researchers gratefully acknowledge the financial support of the Office of Naval Research under Grants N00014-12-1-0024 and N00014-12-1-0379. The helpful technical input from Dr. Paul Hess was also appreciated.

References

- Akpan, U.O., Koko, T.S., Ayyub, B., Dunbar, T.E., 2001. Risk assessment of aging ship hull structures in the presence of corrosion and fatigue. *Mar. Struct.* 15, 211–231.
- Bathe, K.J., 1996. *Finite Element Procedures*. Prentice Hall, New Jersey.
- Bazant, Z.P., Cedolin, L., 1991. *Stability of Structures*. Oxford University Press, Oxford.
- Bishop, C.M., 2006. *Pattern Recognition and Machine Learning*. Springer.
- Brooks, Stephen P., Gelman, Andrew, 1998. General methods for monitoring convergence of iterative simulations. *J. Comput. Graph. Stat.* 7 (4), 434–455.
- Curadelli, O., Ambrosini, D., 2011. Damage detection in elevated spherical containers partially filled with liquid. *Eng. Struct.* 33 (4), 2708–2715.
- Cusanelli D.S., Chesnakas C.J., 2007. Joint High Speed Sealift (JHSS) Baseline Shaft and Strut (BSS) Model 5653-3: Series 2, Propellor Disk LDV Wake Survey; and Series 3 Stock Propellor Powering and Stern Flap Evaluation Experiments. Technical Report Hydrodynamics Section Report NSWCCD-50-TR-2007/084, Naval Surface Warfare Center, Carderock Division, West Bethesda, Maryland, USA.
- Edlund, N.L.O., 2007. Buckling of metallic shells: buckling and postbuckling behaviour of isotropic shells, especially cylinders. *Struct. Control a Health Monit.* 14, 693–713.
- Everstine, G.C., 1981. Structural analogies for scalar field problems. *Int. J. Numer. Methods Eng.* 17, 471–476.
- Everstine, G.C., 1997. Finite element formulations of structural acoustics problems. *Comput. Struct.* 65, 307–321.
- Farrar, C.R., Worden, K.W., 2007. An introduction to structural health monitoring. *Philos. Trans. R. Soc. A* 365, 303–315.
- Featherston, C.A., 2003. Imperfection sensitivity of curved panels under combined compression and shear. *Int. J. Non-Linear Mech.* 38, 225–238.
- Flegal, J.M., Jiang, D., Haran, M., Jones, G.L., 2008. Markov chain Monte Carlo: can we trust the third significant figure?. *Stat. Sci.* 23 (2), 250–260.
- Frisk, G.V., 2012. Noiseconomics: the relationship between ambient noise levels in the sea and global economic trends. *Sci. Rep.* 2 (437).
- Gelman, Andrew, Carlin, Jhon B., Stern, Hal S., Rubin, Donald B., 2004. *Bayesian Data Analysis*. Chapman & Hall/ CRC, 2nd ed.
- Gelman, A., Rubin, Donald B., 1992. Inference from iterative simulation using multiple sequences. *Stat. Sci.* 7 (4), 457–472.
- Ghanem, R.G., Wojtkiewicz, S.F., 2004. Special issue on uncertainty quantification. *SIAM J. Sci. Comput.* 26 (2) vii–vii.
- Green, P.J., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82 (4), 711–732.
- Hastings, W.K., 1970. Monte Carlo sampling methods using Markov chain and their applications. *Biometrika* 57 (1), 97–109.
- Jefferys, W., Berger, J., 1992. Ockham's razor and Bayesian analysis. *Am. Sci.* 80 (1), 64–72.
- Jones, G.L., Harran, M., Caffo, B.S., Neath, R., 2006. Fixed-width output analysis for Markov chain Monte Carlo. *J. Am. Stat. Assoc.* 101 (476), 1537–1548.
- Koutsourelakis, P.S., 2009. A multi-resolution, non-parametric, Bayesian framework for identification of spatially-varying model parameters. *J. Comput. Phys.* 228, 6184–6211.
- Kristanic, Korelc, J., 2008. Optimization method for the determination of the most unfavorable imperfection of structures. *Comput. Mech.* 42, 859872.
- Langtangen, H.P., 1999. *Computational Partial Differential Equations, Numerical Methods and Diffpack Programmin*. Springer, Berlin.
- Link, W.A., Barker, R.J., 2010. *Bayesian inference with ecological examples*. Academic Press, San Diego.
- Wiggins, S.M., McKenna, M.F., Ross, D., Hildebrand, J.A., 2012. Underwater radiated noise from modern commercial ships. *J. Acoust. Soc. Am.* 131 (1), 92–103.
- Papadopoulos, V., Papadrakakis, M., 2004. Finite-element analysis of cylindrical panels with random initial imperfections. *J. Eng. Mech.* 8, 867–876.
- Papadopoulos, V., Papadrakakis, M., 2005. The effect of material and thickness variability on the buckling load of shells with random initial imperfections. *Comput. Methods Appl. Mech. Eng.* 194, 1405–1426.
- Papadopoulos, V., Pavlos, I., 2007. The effect of non-uniformity of axial loading on the buckling behavior of shells with random imperfections. *Int. J. Solids Struct.* 44, 62996317.
- Papadopoulos, V., Stefanou, G., Papadrakakis, M., 2009. Buckling analysis of imperfect shells with stochastic non-Gaussian material and thickness properties. *Int. J. Solids Struct.* 46, 28002808.
- Rasmussen, C.E., Ghahramani, Z., 2001. Occams razor. In: *Neural Information Processing Systems*, vol. 13, pp. 294–300.
- Reed, H.M., Earls, C.J., Nichols, J.M., 2014. Stochastic identification of imperfections in a submerged shell structure. *Comput. Methods Appl. Mech. Eng.* 272, 58–82.
- Reed, H.M., Nichols, J.M., Earls, C.J., 2013. A modified differential evolution algorithm for damage identification in submerged shell structures. *Mech. Syst. Signal Process.* 39, 396–408.
- Richardson, S., Green, P.J., 1997. On Bayesian analysis of mixture with an unknown number of components. *J. R. Stat. Soc. Ser. B* 59 (4), 731–792.
- Rytter A., 1993. *Vibration Based Inspection of Civil Engineering Structures* (Ph.D. dissertation). Department of Building Technology and Structural Engineering, Aalborg University, Denmark.
- Singer J., Arbocz J., Weller T., 1998. *Buckling Experiments*, vol. 1. John Wiley and Sons, Inc.
- Singer J., Arbocz J., Weller T., 2002. *Buckling experiments*, vol. 2. John Wiley and Sons, Inc.
- Stull, C.J., Earls, C.J., Aquino, W., 2008. A posteriori initial imperfection identification in shell buckling problems. *Comput. Methods Appl. Mech. Eng.* 198, 260–268.
- Stull, C.J., Nichols, J.M., Earls, C.J., 2011. Stochastic inverse identification of geometric imperfections in shell structures. *Comput. Methods Appl. Mech. Eng.* 200, 2256–2267.
- Yu, L., Cheng, L., Yam, L.H., Yan, Y.J., Jiang, J.S., 2007. Experimental validation of vibration-based damage detection for static laminated composite shells partially filled with fluid. *Compos. Struct.* 79, 288–299.
- Yu, L., Cheng, L., Yam, L.H., Yan, Y.J., Jiang, J.S., 2007. Online damage detection for laminated composite shells partially filled with fluid. *Compos. Struct.* 80, 334–342.
- Zein, S., Canot, E., Erhel, J., Nassif, N., 2008. Determination of the mechanical properties of a solid elastic medium from a seismic wave propagation using two statistical estimators. *Math. Mech. Solids* 13, 388–407.
- Zhou, W., Wu, Z., Mevel, L., 2010. Vibration-based damage detection to the composite tank filled with fluid. *Struct. Health Monit.* 9, 433–445.