

A Computationally-Efficient Probabilistic Approach to Model-Based Damage Diagnosis

James E. Warner¹, Geoffrey F. Bomarito¹, Jacob D. Hochhalter¹, William P. Leser¹, Patrick E. Leser¹, and John A. Newman¹

¹ NASA Langley Research Center, Hampton, VA, 23666, USA

james.e.warner@nasa.gov

geoffrey.f.bomarito@nasa.gov

jacob.d.hochhalter@nasa.gov

william.p.leser@nasa.gov

patrick.e.leser@nasa.gov

john.a.newman@nasa.gov

ABSTRACT

This work presents a computationally-efficient, probabilistic approach to model-based damage diagnosis. Given measurement data, probability distributions of unknown damage parameters are estimated using Bayesian inference and Markov chain Monte Carlo (MCMC) sampling. Substantial computational speedup is obtained by replacing a three-dimensional finite element (FE) model with an efficient surrogate model. While the formulation is general for arbitrary component geometry, damage type, and sensor data, it is applied to the problem of strain-based crack characterization and experimentally validated using full-field strain data from digital image correlation (DIC). Access to full-field DIC data facilitates the study of the effectiveness of strain-based diagnosis as the distance between the location of damage and strain measurements is varied. The ability of the framework to accurately estimate the crack parameters and effectively capture the uncertainty due to measurement proximity and experimental error is demonstrated. Furthermore, surrogate modeling is shown to enable diagnoses on the order of seconds and minutes rather than several days required with the FE model.

1. INTRODUCTION

Structural health monitoring (SHM) is the driving technology behind the transition from time-based to condition-based maintenance. Motivated by both safety and economic drivers, this paradigm shift from *offline* inspection to *online* (i.e., while operating) monitoring is critically important to industries including manufacturing, aerospace, and defense that seek to

detect damage in structural and mechanical systems at the earliest possible time. For example, an online monitoring system onboard an aircraft would be capable of producing a damage diagnosis well within the time period between flights and, ideally, during the course of a single flight. While SHM is in the process of making the transition into the application domain, the evolution of the technology to enable damage prognosis to forecast residual life has very few deployed applications (Farrar & Worden, 2013). As damage prognosis is inherently probabilistic in nature and presumes a properly characterized initial damage state, its practical use is predicated on not just the detection and localization of damage from SHM, but on a thorough assessment of the extent of the damage along with rigorous uncertainty quantification (UQ).

In order to deliver a more comprehensive online health management system for practical use, a SHM system should possess several key characteristics to enable integration with damage prognosis. Since an explicit quantification of damage is required for prognosis, model-based (inverse problem) SHM is preferred to a data-based approach since the latter is generally limited to detection and localization in the absence of training data from damage states (Barthorpe, 2010). In this case, high-fidelity modeling (e.g., finite element (FE) analysis) is needed to allow for arbitrary geometries and damage types to be considered (limited only by sensitivity of sensors to the damage indices). The damage diagnosis approach must also effectively incorporate UQ to facilitate probabilistic prognostics rather than providing only deterministic assessments. Finally, in order to make online application of the framework feasible, the algorithms deployed must also be computationally efficient. Unfortunately, model-based SHM with high-fidelity modeling implies time-consuming simulations and UQ often requires tens of thousands of such analyses, so taking such an approach is generally impractical if not impossible.

James Warner et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Because of this, existing damage detection techniques have largely been deterministic in nature and have identified structural anomalies based on changes in measured mechanical response (e.g., vibrations (Kim & Stubbs, 2002; Mal, Ricci, Banerjee, & Shih, 2005), ultrasonic wave characteristics (L. Wang & Yuan, 2007; Kehlenbach & Hanselka, 2003), and strains (Krishnamurthy & Gallegos, 2011; Hochhalter, Krishnamurthy, Aguilo, & Gallegos, 2016)). While deterministic approaches have been successfully used to accurately locate and sometimes quantify damage in a computationally-efficient manner, these methods neglect the impact of uncertainty that is ubiquitous in SHM systems deployed in the field due to effects such as sensor noise and modeling assumptions.

More recently, there has been increased focus on uncertainty quantification for damage diagnosis using Bayesian inference in order to explicitly account for measurement and model uncertainties in practice. Several studies (Moore, Murphy, & Nichols, 2011; Nichols, Link, Murphy, & Olson, 2010; Huhtala & Bossuyt, 2011) have used noisy vibrations data to detect structural damage, while in one such study (Nichols, Moore, & Murphy, 2011), the emphasis was on the development of an efficient numerical sampling algorithm for exploring the resulting probability distribution. A Bayesian imaging method was developed to probabilistically estimate delamination location and size in composite laminates using Lamb wave measurements (Peng, Saxena, Goebel, Xiang, & Liu, 2014). Additionally, Bayesian inference and the extended FE method has been used to inversely estimate the probability distribution of crack location and size using strain data (Yan, 2012). Most recently, displacement data were used to estimate the parameters of a continuum mechanics model within a Bayesian framework while Kalman filters were subsequently used to update and evolve the system state in time (Prudencio, Bauman, Faghihi, Ravi-Chandar, & Oden, 2015).

Compared to deterministic methods, Bayesian approaches have the advantage of quantifying uncertainty in the estimates provided, but also incur a substantial computational penalty. Here, the computational expense results from the numerical sampling algorithms, e.g., Markov Chain Monte Carlo (MCMC) (Gamerman & Lopes, 2006), which can exhibit slow convergence and involve the evaluation of a potentially time-consuming computational model for each sample drawn. To alleviate this computational burden, advanced MCMC methods have been developed to reduce sampling time by improving sampling convergence (Haario, Laine, & Mira, 2006; Nichols et al., 2011) or through parallelization of the algorithms themselves (Vrugt et al., 2009; Neiswanger, Wang, & Xing, 2013; Prudencio & Cheung, 2012; Warner, Zubair, & Ranjan, 2017). Another common approach is to replace the original physics-based model with a computationally-efficient surrogate model using probabilistic spectral methods (Marzouk, Najm, & Rahn, 2006) or machine learning algorithms (Meeds & Welling, 2014).

The development of surrogate model-accelerated Bayesian approaches for model-based SHM applications remains relatively limited in comparison to data-based approaches. However, machine learning has been used in conjunction with physics-based models to learn the inverse-map directly from measurement to damage directly, rather than the forward-map in traditional surrogate modeling (Katsikeros & Labeas, 2009; Sbarufatti, Manes, & Giglio, 2013). Most notably, this inverse-mapping approach was demonstrated using artificial neural networks (ANNs) to perform damage localization and quantification using experimentally measured strains in a simplified structure resembling a helicopter fuselage (Sbarufatti et al., 2013). The uncertainty quantification effort, however, was limited to confidence intervals based off the scatter in predictions from various ANNs.

Motivated by online, integrated SHM and damage prognosis, this study demonstrates an efficient and general approach to probabilistic model-based damage diagnosis. Given measurement data, probability distributions of unknown damage parameters are estimated using Bayesian inference and MCMC sampling. The framework is applicable to arbitrary component geometries and damage types as well as different sensor data and is formulated as such. This generality is enabled under the assumption that a properly-calibrated, high-fidelity model (e.g., via FE modeling) is available with adequate predictive capability of the quantity being measured by sensors. For computational speedup, it is proposed that this potentially expensive model is only used *offline* to generate training data for the development of surrogate models that can be used for rapid *online* diagnosis. Furthermore, the surrogate training data can also be used to generate an informed initial guess for the unknown damage, which can significantly improve convergence of MCMC sampling for additional gains in efficiency.

The general model-based diagnosis approach is applied to the problem of strain-based crack characterization and experimentally validated using digital image correlation (DIC) (Peters & Ranson, 1982) strain data from two cracked lab specimens. Access to full-field DIC data facilitates the study of the effectiveness of strain-based diagnosis as the distance between the locations of damage and strain measurements is varied. It also demonstrates the ability of the Bayesian framework to capture the growing uncertainty in diagnosis as the signal-to-noise ratio is decreased in this way. It is shown that surrogate modeling provides orders of magnitude computational speed up with respect to using a FE model while retaining satisfactory accuracy.

With respect to the preliminary work done on the damage diagnosis framework (Warner, Hochhalter, Leser, Leser, & Newman, 2016), this study emphasizes the practicality of the approach and provides explicit quantification of the accuracy versus efficiency tradeoff associated with surrogate modeling.

The formulation provided here serves as a general, application-independent procedure for probabilistic damage diagnosis in terms of a generic damage model and sensor data type. It also provides more details on the required model calibration step, a commonly-stated hurdle to implementing model-based SHM. The strain-based crack characterization example is tailored in this paper to mimic a practical scenario where fiber optic strain sensors (Meltz & Snitzer, 1981; Li, Li, & Song, 2004) are used to provide monitoring data. It is shown that accurate damage diagnoses can be obtained with just one component of strain along two linear arrays (as opposed to two strain components used in the preceding study). Finally, a rigorous comparison of the surrogate-accelerated diagnosis framework versus a reference implementation with a FE model is provided here to quantify the gains in performance obtained. The comparison includes a demonstration of the improvements in MCMC convergence obtained by the simple scheme for generating a favorable initial guess versus a randomly chosen guess.

The remainder of the paper is organized as follows. First, a complete formulation of the proposed damage diagnosis approach is provided in the following section, with individual subsections devoted to model-based diagnosis, Bayesian inference, MCMC sampling, and surrogate modeling. Next, the specifics of applying the diagnosis framework to the problem of strain-based crack characterization are illustrated. Experimental validation of the approach is then provided, including the experimental strain data obtained with digital image correlation, the development and performance of the surrogate models used, and results of the damage diagnosis method applied to both damage localization and characterization in two separate lab specimens containing cracks are presented. The gains in computational efficiency enabled through surrogate modeling are explicitly quantified here. Finally, the findings of the study are summarized in the conclusion section.

2. FORMULATION

In this section, the proposed probabilistic approach for model-based damage diagnosis is presented. Starting from a deterministic model-based approach, Bayesian inference is then used to deduce the probability distribution of damage parameters conditional on available sensor data. MCMC sampling is employed to explore the resulting distribution. Finally, surrogate modeling is introduced to provide computational speedup in the approach. These points are elaborated on in detail in the subsequent subsections.

2.1. Model-Based Diagnosis

Damage diagnosis methods operate under the assumption that the mechanical response of a structural component is altered in the presence of damage. To this end, the goal of diagnosis is to use measured response data $\mathbf{d}^{\text{meas}} \in \mathbb{R}^m$ to detect

if damage is present and then ideally estimate some parameters $\mathbf{c} \in \mathbb{R}^d$ that characterize the damage (location, size, etc.). Model-based approaches to diagnosis require a model of the structural component, \mathcal{M} , capable of predicting the mechanical response $\mathbf{y} \in \mathbb{R}^m$ for a given set of damage parameters

$$\mathcal{M}(\mathbf{c}; \mathbf{f}) = \mathbf{y} \in \mathbb{R}^m, \quad (1)$$

where \mathbf{f} are any additional free parameters (material properties, boundary conditions, etc.) affecting the output.

Prior to diagnosis, the parameters \mathbf{f} must be properly prescribed for the specific component being monitored through a model calibration procedure. Generally speaking, this is done by taking an initial measurement of the component, \mathbf{d}^{cal} , for a known, often undamaged state, \mathbf{c}^{cal} . Then, system-specific model parameters, $\hat{\mathbf{f}}$, can be found through an optimization problem of the following form

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} g(\mathcal{M}(\mathbf{c}^{\text{cal}}; \mathbf{f}) - \mathbf{d}^{\text{cal}}), \quad (2)$$

where $g(\cdot)$ is a scalar-valued norm function. After model calibration, it is assumed that $\mathbf{y} \approx \mathbf{d}^{\text{meas}}$ for a damage estimate \mathbf{c} that accurately characterizes the true damage. For the remainder of the formulation, the explicit dependence on \mathbf{f} will be suppressed to denote a calibrated model (e.g., $\mathcal{M}(\mathbf{c}) \equiv \mathcal{M}(\mathbf{c}; \hat{\mathbf{f}})$).

In the context of model-based diagnosis, \mathcal{M} is referred to as the forward model while the diagnosis problem of using \mathbf{d}^{meas} to infer \mathbf{c} is the associated inverse problem. A typical deterministic approach to solving this inverse problem is to first pose an error metric between the measured response data and corresponding model response

$$Q(\mathbf{c}, \mathbf{d}^{\text{meas}}) = \sum_{i=1}^m \|d_i^{\text{meas}} - \mathcal{M}_i(\mathbf{c})\|^2, \quad (3)$$

where $\mathcal{M}_i(\mathbf{c}) \equiv y_i$, the i^{th} component of the model output. Then, gradient-based or global optimization algorithms are employed to find the damage parameters that minimize Equation 3 to produce the so-called least squares estimator

$$\mathbf{c}^{LS} = \arg \min_{\mathbf{c}} Q(\mathbf{c}, \mathbf{d}^{\text{meas}}). \quad (4)$$

The primary drawback of such deterministic approaches for model-based diagnosis is that only a point estimate of the damage is produced with no regard to uncertainty inherent in the measurement data (noise, sparsity, etc.). A suitable regularization strategy must also be chosen and tuned in an effort to rectify the well known ill-posedness of the inverse problem (Isakov, 1998) (e.g., infinitely many solutions may exist).

2.2. Bayesian Inference

The Bayesian inference approach to model-based diagnosis reformulates the inverse problem (Equation 4) as one of deducing a probability distribution of the unknown damage parameters, \mathbf{c} , conditional on the observed measurement data \mathbf{d}^{meas} . This distribution, $p(\mathbf{c}|\mathbf{d}^{\text{meas}})$, known as the *posterior distribution*, is given according to Bayes' Theorem (Kaipio & Somersalo, 2004):

$$p(\mathbf{c}|\mathbf{d}^{\text{meas}}) = \frac{p(\mathbf{d}^{\text{meas}}|\mathbf{c})p(\mathbf{c})}{p(\mathbf{d}^{\text{meas}})} \propto p(\mathbf{d}^{\text{meas}}|\mathbf{c})p(\mathbf{c}), \quad (5)$$

integrating any knowledge about the damage prior to the measurement in the *prior distribution*, $p(\mathbf{c})$, with the information from the data, \mathbf{d}^{meas} , through the *likelihood function*, $p(\mathbf{d}^{\text{meas}}|\mathbf{c})$. Note that the normalizing constant, $p(\mathbf{d}^{\text{meas}})$, need not be computed in order to explore the posterior probability distribution with MCMC, as explained later.

The prior density function, $p(\mathbf{c})$, provides an effective way of incorporating an analyst's insight about likely damage characteristics into the Bayesian inference approach. From a mathematical point of view, prescribing such an informative prior density function is an approach to regularize the inverse problem (J. Wang & Zabarar, 2014). While prior distributions have been shown to increase the effectiveness of Bayesian damage diagnosis (Leser & Warner, 2017; Warner, Hochhalter, et al., 2016), a non-informative prior density function (e.g., $p(\mathbf{c}) \propto 1$) is chosen in this work to represent a common practical case where no reliable *a priori* knowledge about the damage is available.

The likelihood function, $p(\mathbf{d}^{\text{meas}}|\mathbf{c})$, models the discrepancy between the measurement data and the predicted values of response from the model. To this end, the following common assumption is made about this relationship (Kaipio & Somersalo, 2004)

$$d_i^{\text{meas}} = \mathcal{M}_i(\mathbf{c}) + \delta_i, \quad \delta_i \sim \text{Normal}(0, \sigma). \quad (6)$$

That is, the measurement data are polluted with errors, δ_i , that are treated as a sequence of independent, identically distributed (i.i.d.) samples drawn from a zero-mean Gaussian (Normal) distribution with variance σ (interpreted as the noise level). The i.i.d. assumption yields the following expression for likelihood function

$$p(\mathbf{d}^{\text{meas}}|\mathbf{c}) = \frac{1}{(2\pi\sigma^2)^{m/2}} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^m \|d_i^{\text{meas}} - \mathcal{M}_i(\mathbf{c})\|^2\right) \propto \exp\left(-\frac{1}{2\sigma^2} Q(\mathbf{c}, \mathbf{d}^{\text{meas}})\right). \quad (7)$$

Here, it is clear that as the error between computed and measured strains (Equation 3) increases, the value of the likelihood function (and hence posterior probability in Equation 5) decreases and vice versa. Note that in cases where signifi-

cant inadequacy or bias of the model remain after calibration, an additional input-dependent model discrepancy term can be added to the righthand side of Equation 6 in an attempt to correct this (Kennedy & O'Hagan, 2001).

The noise level parameter, σ , can be inferred on the fly during damage diagnosis (Warner, Hochhalter, et al., 2016) or estimated beforehand through a comparison of the calibrated model in the reference damage state, $\mathcal{M}(\mathbf{c}^{\text{cal}})$, with the measurement data used for calibration, \mathbf{d}^{cal} . To this end, the empirical error is first calculated for each measurement point

$$\hat{\delta}_l = d_l^{\text{cal}} - \mathcal{M}_l(\mathbf{c}^{\text{cal}}). \quad (8)$$

Then, the noise level, σ , can be estimated using the sample standard deviation

$$\sigma \approx \hat{\sigma} = \left(\frac{1}{1-n} \sum_{l=1}^n (\hat{\delta}_l - \hat{\mu})^2\right)^{1/2}, \quad (9)$$

where n is the number of data points used for calibration and $\hat{\mu}$ is the sample mean of $\hat{\delta}$. This approach will be demonstrated in the application section of this paper.

2.3. Markov Chain Monte Carlo

The solution of the model-based diagnosis problem as the posterior probability distribution in Equation 5 is not practically helpful since it can rarely be evaluated analytically. MCMC (Gamerman & Lopes, 2006) is a powerful tool for numerically forming probabilistic damage estimates based on $p(\mathbf{c}|\mathbf{d}^{\text{meas}})$. The goal of MCMC is to generate a collection of N damage parameter samples from the posterior probability distribution

$$\{\mathbf{c}^{(j)}\}_{j=1}^N \text{ where } \mathbf{c}^{(j)} \sim p(\mathbf{c}|\mathbf{d}^{\text{meas}}), \quad (10)$$

which can then be used to construct empirical probability distributions, credibility intervals, and moment estimates for \mathbf{c} . Algorithm 1 summarizes a very basic instantiation of MCMC, the Metropolis algorithm, used in this work:

Algorithm 1 Metropolis MCMC

```

Initialize  $\mathbf{c}^{(0)}$ 
for  $j = 1 : N$  do
  Sample  $u \sim \text{Uniform}(0, 1)$ 
  Sample  $\mathbf{c}^* \sim q(\mathbf{c}^*|\mathbf{c}^{(j-1)})$ 
  if  $u < A(\mathbf{c}^*, \mathbf{c}^{(j-1)}) = \min\{1, \frac{p(\mathbf{c}^*|\mathbf{d}^{\text{meas}})}{p(\mathbf{c}^{(j-1)}|\mathbf{d}^{\text{meas}})}\}$  then
     $\mathbf{c}^{(j)} = \mathbf{c}^*$ 
  else
     $\mathbf{c}^{(j)} = \mathbf{c}^{(j-1)}$ 
  end if
end for

```

Here, the method simply draws a trial sample, \mathbf{c}^* , at each iteration from a *proposal distribution*, $q(\mathbf{c}^*|\mathbf{c}^{(j-1)})$, and then

decides whether to accept or reject this sample based on the *acceptance probability*, $A(\mathbf{c}^*, \mathbf{c}^{(j-1)})$. The Metropolis algorithm assumes that the proposal distribution is symmetric, where a common choice is a Gaussian distribution centered at the previous sample

$$q(\mathbf{c}^* | \mathbf{c}^{(j-1)}) = \text{Normal}(\mathbf{c}^{(j-1)}, \Sigma_q), \quad (11)$$

and Σ_q is the user-specified covariance matrix. Algorithm 1 with Equation 11 constructs a Markov chain that, by design, is guaranteed to have a stationary distribution that reflects the true posterior distribution in Equation 5 (Gamerman & Lopes, 2006). Note that since the posterior probability distribution only appears as a ratio through A , the normalizing constant $p(\mathbf{d}^{\text{meas}})$ in Equation 5 is not required.

Although the Metropolis MCMC algorithm above is conceptually simple and straightforward to implement, the number of samples (N) required for convergence can be very large, making its application challenging and often infeasible. The convergence rate is directly related to the selection of Σ_q , an appropriate value for which is often unknown a priori and can be difficult to tune on the fly. Convergence of MCMC sampling is also greatly affected by the choice of initial guess, $\mathbf{c}^{(0)}$, where a lengthy *burn-in* period (samples discarded from the beginning of the chain) can be necessary for a value of $\mathbf{c}^{(0)}$ far from the true damage.

While advanced algorithms that rely on adaptively selecting Σ_q (Warner, Hochhalter, et al., 2016) or parallel computing (Warner et al., 2017) can be used to accelerate MCMC-based diagnosis, this work employs a simple scheme to generate a highly probable initial guess that yields efficient performance from the Metropolis algorithm in its most basic form. This approach is elaborated on at the end of the following section.

2.4. Surrogate Modeling

Combining Bayesian inference (Equation 5) and MCMC sampling (Algorithm 1), as described in the previous sections, yields a powerful method for generating probabilistic damage estimates. However, the primary challenge associated with the approach is that the model, \mathcal{M} , must be evaluated for each sample drawn with MCMC. Since large values of N are typically required for convergence of the sampling process, Bayesian diagnosis can be infeasible even for modestly expensive models, especially in its application to online SHM.

Surrogate modeling is a technique that can alleviate the computational burden associated with probabilistic model-based diagnosis when intensive, high fidelity simulations are required for the components being monitored (Meeds & Welling, 2014; Warner & Hochhalter, 2016; Warner et al., 2017). The approach relies on the (*offline*) pre-computation and storage of input-output pair datasets from an original computational model in an effort to replace it during (*online*) analysis by a

more efficient data-driven model. Furthermore, with a sufficient amount of pre-computed data and an effective regression/interpolation algorithm, a high degree of accuracy with respect to the original model can be maintained.

To utilize surrogate modeling for model-based damage diagnosis, a set of T damage parameter arrays, $\{\mathbf{c}^{(k)}\}_{k=1}^T$, is first selected. Then, the model responses corresponding to all m measurements are computed and stored for each damage state,

$$\mathcal{M}_i^{(k)} \equiv \mathcal{M}_i(\mathbf{c}^{(k)}), \quad k = 1, \dots, T, \quad (12)$$

for $i = 1, \dots, m$. The result is the following $T \times (d + m)$ input-output dataset

$$\mathcal{S} = \{\mathbf{c}^{(k)}; \mathcal{M}_1^{(k)}, \dots, \mathcal{M}_m^{(k)}\}_{k=1}^T. \quad (13)$$

From a machine learning perspective, \mathcal{S} is the training data and a variety of off-the-shelf regression and interpolation algorithms can be utilized to directly infer the input-output mappings. Specifically, a surrogate model that maps a new damage state, $\mathbf{c}^{(*)}$, to the predicted sensor response is generated offline for each individual measurement

$$\widetilde{\mathcal{M}}_i : \mathbf{c}^{(*)} \rightarrow \mathcal{M}_i^{(*)} \quad \text{for } i = 1, \dots, m. \quad (14)$$

Now, the original model, \mathcal{M} , is replaced by the set of surrogate models, $\{\widetilde{\mathcal{M}}_i\}_{i=1}^m$, in the posterior probability distribution (Equation 5) so that sampling can be conducted rapidly for damage diagnosis.

A couple of remarks about the surrogate modeling process are worth noting. First, the size, T , of the training dataset has a lower limit based on accuracy requirements and a practical upper limit based on the computational expense of the original model, \mathcal{M} , the computational resources available, and the training complexity and memory requirements of the regression/interpolation algorithm used. In this work, a testing dataset, $\hat{\mathcal{S}}$, of P randomly generated damage states and corresponding responses

$$\hat{\mathcal{S}} = \{\hat{\mathbf{c}}^{(n)}; \hat{\mathcal{M}}_1^{(n)}, \dots, \hat{\mathcal{M}}_m^{(n)}\}_{n=1}^P, \quad (15)$$

is used to evaluate surrogate model accuracy. The assessment is based on the relative error between the surrogate and original models when predicting the test data

$$\Delta_j = \frac{1}{P} \sum_{n=1}^P \frac{|\widetilde{\mathcal{M}}_j(\hat{\mathbf{c}}^{(n)}) - \hat{\mathcal{M}}_j^{(n)}|}{|\hat{\mathcal{M}}_j^{(n)}|}. \quad (16)$$

Equations 15 and 16 can be used to compare different learning algorithms and select the most accurate one (i.e., with lowest error, Δ_j) to use as a surrogate model for damage diagnosis. The equations can also be used to guide the selection of an appropriate training dataset size, T . That is, starting from an initial number of data points, new training data can be incrementally added to \mathcal{S} until a prescribed accuracy threshold

based on Δ_j is met.

It is also worth pointing out that the T model evaluations and the training of the m surrogate models is an offline cost associated with the diagnosis framework. That is, the computational burden of generating the surrogate models is a single upfront cost that then permits an arbitrary number of efficient damage diagnoses to be conducted by rapidly evaluating $\widetilde{\mathcal{M}}_j$ during the online analyses. Furthermore, the T executions of the FE simulation are completely independent of each other, and can therefore be run in parallel on as many computer processors as are available.

As shown in the preceding study (Warner, Hochhalter, et al., 2016), an additional benefit of surrogate modeling is that the training dataset, \mathcal{S} , can be used to generate a favorable initial guess for MCMC sampling, $\mathbf{c}^{(0)}$, that can reduce the required burn-in period. This is done by computing the least squares estimator (Equation 3) over the input-output dataset (Equation 13)

$$\mathbf{c}^{(0)} = \arg \min_{\mathbf{c} \in \mathcal{S}} Q(\mathbf{c}, \mathbf{d}^{\text{meas}}). \quad (17)$$

By only considering the precomputed training grid values, this computation can be done rapidly as it does not require any additional model evaluations. While multimodal distributions may still pose a challenge, Equation 17 provides a simple and systematic way to generate an initial guess that will reside in a high probability region of the posterior distribution as a good starting point (Smith, 2013). It will be shown in this work that even the basic Metropolis MCMC method (Algorithm 1) can be effective and robust when starting the algorithm in this fashion.

2.5. Summary

The formulation presented thus far has prescribed a general framework for using noisy sensor data to produce probabilistic damage estimates using a model-based diagnosis approach. The formulation is broadly applicable irrespective of the damage description, \mathbf{c} , type of measurement data, \mathbf{d}^{meas} , and computational model, \mathcal{M} , under two primary assumptions. First, the measurement data must be sufficiently sensitive to changes in the damage parameters chosen. This is dependent on the sensor density, the quality of the measurements (signal-to-noise ratio), and the complexity of the damage description adopted. Second, the model, after proper calibration, must have adequate predictive capability of the quantity being measured for the range of possible damage states. The validity of these two assumptions can be further investigated prior to conducting diagnosis through sensitivity analysis (Saltelli, Chan, & Scott, 2000; *Global sensitivity analysis: the primer*, n.d.) and model validation (Roy & Oberkampf, 2011) studies, respectively.

In some simple scenarios, the model, \mathcal{M} , may be efficient enough to use directly in the framework. Otherwise, the sur-

rogate modeling strategy in Section 2.4 is a viable approach to alleviate the computational burden associated with executing \mathcal{M} repeatedly during MCMC sampling. Successful application of surrogate modeling is contingent upon thorough tuning and testing of the regression algorithm used and the ability to generate a sufficient amount of training data. To this end, it is important in practice to choose the simplest, low dimensional description of damage possible for the application, as regression becomes more difficult in high dimensions and the amount of training data required grows exponentially with the number of input parameters. Furthermore, higher order descriptions of damage are limited in the first place by the sensitivity of the measurement data to them for diagnosis.

To summarize, the necessary steps that must be taken both offline (prior to putting the component in service) and online (while operating) to implement the diagnosis framework are provided below.

Offline:

1. Create model, $\mathcal{M}(\mathbf{c}, \mathbf{f})$, of component and calibrate it
 - Perform initial measurement
 - Determine optimal parameters, $\hat{\mathbf{f}}$, to produce calibrated model, $\mathcal{M}(\mathbf{c})$, (Equation 2)
 - Estimate noise level, σ^2 , (Eqs. 8, 9)
2. Train surrogate models
 - Generate input-output training dataset, \mathcal{S} , (Eqs. 12, 13)
 - Train surrogate models for each measurement (Equation 14)
 - Validate surrogate models, add training data and re-train if necessary

Online:

1. Acquire measurement data, \mathbf{d}^{meas} , from sensors
2. Compute initial guess, $\mathbf{c}^{(0)}$, (Equation 17)
3. Perform MCMC sampling (Algorithm 1) utilizing trained surrogate models

The next section will demonstrate how this general framework can be applied to solve a specific, practical damage diagnosis problem.

3. EXPERIMENTAL VALIDATION

3.1. Application: Strain-Based Crack Characterization

The general model-based diagnosis framework presented in the previous section is now applied to the specific problem of crack characterization in thin plates using strain data. A schematic illustrating the application can be seen in Figure 1. Here, the damage is represented by a four-dimensional array

$$\mathbf{c} = [x, y, a, \theta], \quad (18)$$




Figure 1. Boundary conditions and damage parameterization for the crack characterization application.

describing the center location, (x, y) , length (a) , and orientation (θ) of the crack.

The panel is subject to general prescribed displacement boundary conditions along the top and bottom surfaces, comprising the additional model parameters (\mathbf{f}) for this application

$$\mathbf{f} = [U_x^1, U_z^1, U_x^2, U_z^2, U_y^{12}, U_x^3, U_z^3, U_x^4, U_z^4, U_y^{34}]. \quad (19)$$

The measurement data used to estimate the crack parameters are an array of m strains recorded throughout the domain

$$\mathbf{d}^{\text{meas}} = \{\hat{S}_i\}_{i=1}^m, \quad (20)$$

where \hat{S}_i is the i^{th} strain measurement. Note that \hat{S}_i can generally represent any one of the surface strain components $[\hat{\epsilon}_{xx}, \hat{\epsilon}_{yy}, \hat{\gamma}_{xy}]$. However, in this study it will be limited to a single component to mimic the capabilities of fiber optic strain sensing (Meltz & Snitzer, 1981), as described in the next section. For this application, $\mathcal{M}(\mathbf{c}; \mathbf{f})$ (Equation 1) is a FE model capable of computing and extracting strains at the measurement locations for a given set of damage parameters and prescribed displacements.

The implementation of the method, including surrogate model development and MCMC sampling, was carried out in Python (Python Software Foundation, 2016). All FE modeling was performed using the Scalable Implementation of Finite Elements by NASA (SciFEN) (Warner, Bomarito, Heber, & Hochhalter, 2016) software. Both the FE and Python software for this study was executed on a machine with quad-core 2.4GHz AMD Opteron processors. The remainder of the section will detail each aspect of the experimental vali-

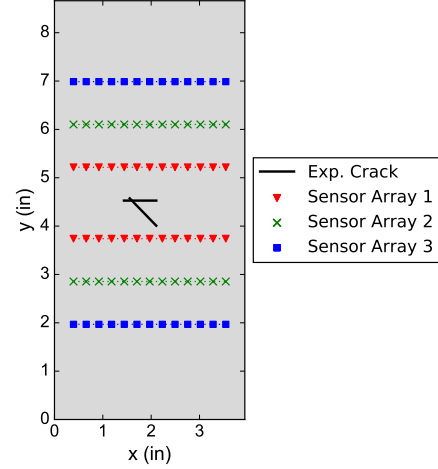


Figure 2. Diagram of the effective sensor arrays tested and the cracks from the flat and angled crack specimens.

ation effort. First, a description of the strain measurement data used for diagnosis is provided followed by an overview of the FE model calibration that was performed. Next, the development and verification of the surrogate models used for accelerated diagnosis is then presented. Finally, the performance of the framework is demonstrated on two examples: 1) damage localization and 2) general crack characterization (location, size, and orientation) in thin metal plates.

3.2. DIC Strain Data

Two cracked thin sheet specimens of Aluminum Alloy 2024 (AA2024) were considered for experimental validation of the diagnosis framework, one with a flat crack (i.e., oriented 0° from the x axis) used to test damage localization and the other with an angled crack to test full crack characterization. The width and height of the specimens were 3.93in and 8.73in, respectively. The crack parameters were $[x_{\text{flat}}, y_{\text{flat}}, a_{\text{flat}}, \theta_{\text{flat}}] = [1.81\text{in}, 4.53\text{in}, 0.67\text{in}, 0\text{rad}]$ for the flat crack specimen and $[x_{\text{angled}}, y_{\text{angled}}, a_{\text{angled}}, \theta_{\text{angled}}] = [1.83\text{in}, 4.29\text{in}, 0.78\text{in}, -0.82\text{rad}]$ for the angled crack specimen.

Full-field strain data was acquired with DIC using the VIC3D (Correlated Solutions Inc., 2012) software. DIC is an optical measurement technique that takes a sequence of digital photographs of a component and uses computer vision algorithms to track blocks of pixels and build up full 2D and 3D deformation and strain fields (Sutton, Orteu, & Schreier, 2009). A speckle pattern was first applied to the test specimens using spray paint to facilitate pixel tracking with DIC. Each specimen was then loaded individually in tension and strain fields were obtained using VIC3D.

The motivation behind gathering full-field strain data using DIC was that measurement (“sensor”) locations could be freely chosen in order to test the diagnosis accuracy and uncertainty

as the distance between the measured data and damage was varied. Three different sensor arrangements were tested in the diagnosis examples to follow, shown in Figure 2, along with the two crack configurations considered. Each arrangement was composed of two separate horizontal arrays of sensors with increasing distance between them (1.47in, 3.25in, and 5.02in). Thirteen measurement locations were recorded along each array. Only the ϵ_{XX} component of strain was used at each location, so that there were $m = 26$ measurements for each arrangement. This setup was chosen to mimic a practical scenario where fiber optic strain sensors (Meltz & Snitzer, 1981; Li et al., 2004) were utilized to collect monitoring data.

The resulting strain fields captured with DIC can be seen in Figures 3(a) and 3(b), showing the ϵ_{XX} strain component for the flat and angled crack specimens, respectively. Dashed lines are overlaid across the strain fields to denote where data will be extracted for the three sensor arrays in Figure 2. The localized nature of the strain field in the presence of damage is apparent, indicating the importance of sensor proximity for practical monitoring applications. It is worth pointing out that while four sensor arrays were tested in the previous work (Warner, Hochhalter, et al., 2016), the fourth and furthest array from the damage was removed from consideration in this work because it resulted in diagnoses with little useful information about the cracks. Hence, only diagnosis results using the three sensor arrays in Figure 2 will be presented in the examples to follow.

3.3. Model Calibration

A FE model representing $\mathcal{M}(\mathbf{c}; \mathbf{f})$ for each test specimen based on the measured dimensions given above was first created in ScIFEN to facilitate model calibration. A Young's Modulus, $E = 10.6\text{Msi}$, and a Poisson's ratio, $\nu = 0.33$, for AA2024 were considered known and deterministic. Thus, the model calibration problem in Equation 2 was solved only for an appropriate set of boundary condition parameters in Equation 19 for each test specimen. A more detailed, rigorous calibration is performed here with respect to the crude single parameter approach in (Warner, Hochhalter, et al., 2016) in an attempt to decrease the model discrepancy observed due to misalignment in the test stand used for uniaxial loading.

While in the ideal case initial measurements would be taken prior to the introduction of damage to use for calibration, no DIC measurements were available from the undamaged state in this study. Instead, \mathbf{d}^{cal} was chosen as displacement data in the damaged states (sampled at 500 random locations in the domain) to ensure that the model was still calibrated with a different set of data than was used for diagnosis later. The 2-norm was used for the objective function in Equation 2 to

quantify the error between the model and calibration data, i.e.,

$$g(\mathbf{a} - \mathbf{b}) = \sqrt{\sum_{l=1}^n (a_l - b_l)^2} \quad (21)$$

for two arbitrary n -dimensional vectors, \mathbf{a} and \mathbf{b} . This optimization was performed using the Nelder-Mead algorithm (Wilde & Beightler, 1967) as implemented by the open source Python package `SciPy` (Jones, Oliphant, Peterson, et al., 2001). The model calibration problem was solved individually for the flat and angled crack specimens, resulting in two sets of optimal boundary condition parameters, $\hat{\mathbf{f}}_{\text{flat}}$ and $\hat{\mathbf{f}}_{\text{angled}}$, respectively.

To assess the effectiveness of the calibration and estimate the noise levels, σ^2 , needed for Bayesian inference, the strains from DIC and the calibrated model were then compared on the randomly sampled grid of points. The resulting distribution of errors (Equation 8) is shown in Figure 4 for the flat (a) and angled (b) crack specimens along with a fitted Gaussian distribution to the errors. The estimated noise levels (Equation 9) for Bayesian inference are depicted in the plots and will be used for the damage diagnosis results to follow. The mean values of error here, $\hat{\mu}_{\text{flat}}$ and $\hat{\mu}_{\text{angled}}$, can be interpreted as model bias and would be 0 in the ideal case. While non-zero, these values are substantially smaller using the more rigorous calibration approach here than in the preliminary work (Warner, Hochhalter, et al., 2016): $\hat{\mu}_{\text{flat}}$ was reduced from 1.87×10^{-5} to 1.09×10^{-5} and $\hat{\mu}_{\text{angled}}$ was reduced from 1.01×10^{-4} to 8.40×10^{-6} .

3.4. Surrogate Model Development

Two separate sets of surrogate models were developed for the flat crack and angled crack specimen since the former was used to demonstrate damage localization ($\mathbf{c}_{\text{flat}} = [x, y] \in \mathbb{R}^2$, with $a = 0.67\text{in}$ and $\theta = 0\text{rad}$), while the latter was used for full crack characterization ($\mathbf{c}_{\text{angled}} = [x, y, a, \theta] \in \mathbb{R}^4$). The bounds for the parameters for training and testing were specified as

$$x \in [0.64, 3.18]\text{in}, \quad (22)$$

$$y \in [0.64, 7.94]\text{in}, \quad (23)$$

$$a \in [0.20, 1.19]\text{in}, \quad (24)$$

$$\theta \in [-\pi/2, \pi/2]\text{rad}, \quad (25)$$

where the bounds for x , y , and a were chosen such that the entire crack would always be contained within the geometry (i.e., edge cracks were not considered). For both specimens, surrogate models were constructed and stored for each measurement in the sensor arrays considered (Figure 2) that were capable of mapping new values of \mathbf{c}_{flat} and $\mathbf{c}_{\text{angled}}$ directly to the resulting strain (Equation 14). Several different machine learning algorithms from the `scikit-learn` (Buitinck et

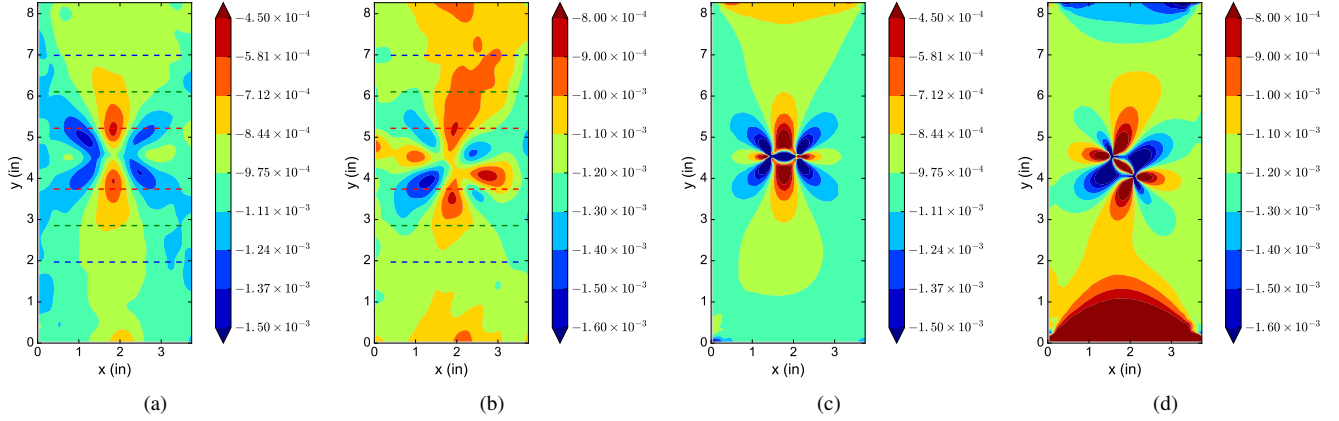


Figure 3. A comparison of the ϵ_{XX} strain field obtained by DIC ((a) flat crack specimen, (b) angled crack specimen) versus the calibrated FE model ((c) flat crack specimen, (d) angled crack specimen). The dashed lines in (a) and (b) indicate the locations of the sensor arrays.




Figure 4. Histogram of errors between DIC and FE strains (Equation 8) with fitted Gaussian distributions, for the a) flat and b) angled crack specimen.

al., 2013) and SciPy (Jones et al., 2001) Python modules were compared to obtain a surrogate model with an optimal balance of prediction accuracy and efficiency.

First, training datasets, \mathcal{S} , (Equation 13) were generated using the ScIFEN FE code according to Equation 12. For the case of damage localization in the flat crack specimen, five uniform training grids, $\{\mathbf{c}_{\text{flat}}^{(k)}\}_{k=1}^T$, were considered from $T = 200$ to $T = 5000$ to study the accuracy and efficiency of the machine learning algorithms for increasing training data size. The grids were created by choosing T_x and T_y equally-spaced points across the range of x and y in Equations 22 and 23, respectively, where $[T_x, T_y] = [10g, 20g]$ for $g = 1, \dots, 5$. Only one training grid with $T = 32076$ was generated for the angled crack specimen due to the added computational expense of the increased dimension of the input space, $d = 2$ to $d = 4$. Here, the number of grid points for each variable was $[T_x, T_y, T_a, T_\theta] = [12, 33, 9, 9]$, again selected to be equally-spaced across the ranges in Equations 22-25. Additionally, two test datasets, $\hat{\mathcal{S}}$, (Equation 15) were generated from 1000 randomly selected values of \mathbf{c}_{flat} and $\mathbf{c}_{\text{angled}}$ to verify the accuracy of the trained surrogate models for each case.

Note that the number of training grid points for each variable were selected based on the size of its range and on the expected sensitivity of computed strains to changes in its value. For example, more grid points were used for x and y versus a and θ for damage characterization due to the larger bounds in Equations 22 and 23 and because strain values at the sensors are more affected by damage proximity than changes in damage size and orientation. While no attempt was made to optimize the selection of the training grid in this study, more sophisticated adaptive and selective refinement schemes could be used to increase accuracy with fewer training data points. It is shown later in this section, however, that the simple uniform grids selected resulted in surrogate models with a satis-




Figure 5. Performance comparison of four different regression algorithms for surrogate modeling in the damage localization problem in terms of a) relative error and b) prediction time.

factory level of accuracy.

Generating all the necessary surrogate training and test data took approximately 2000 CPU-hours for this study. However, exploiting the independent nature of the computations (as mentioned in Section 2.4) and utilizing parallel processing resulted in about 5 days of total run time. Specifically, four separate FE simulations were executed simultaneously with each simulation running on four processors with ScIFEN. The CPU times for each FE model execution varied between 25 and 45 seconds depending on the crack geometry and computational mesh.

Surrogate models for the flat crack specimen were generated using linear regression, nearest neighbors, and Gaussian process algorithms from `scikit-learn` and a multi-linear interpolation algorithm from `SciPy`. Free parameters for the nearest neighbors and Gaussian process models were tuned using cross-validation. A comparison of algorithm performance for increasing training dataset sizes is shown in Figure 5. Figure 5(a) shows the average relative error over the testing dataset (Equation 16), while Figure 5(b) compares prediction times for the different models tested.

It is clear that the nearest neighbors and linear interpolation models provide the most accurate predictions, nearing 1% error for the larger training datasets. In terms of prediction speeds, the linear interpolation model is significantly faster in this case, and was chosen for surrogate modeling for the damage localization study for this reason. It is important to point out, however, that the nearest neighbors model demonstrates a near-constant scaling in prediction time versus the size of the dataset, which will be highlighted in the crack characterization surrogate performance to follow.

The disparity in performance seen with the Gaussian process and linear regression models may be due in part to the global nature of the approximations provided by these meth-

ods (Bishop, 2006) in contrast with the localized behavior of the strains being predicted. Nearest neighbors and linear interpolation benefit in this situation by basing their approximations locally from the data points in closest proximity to the one being predicted. It is also important to point out that these results are not indicative of the performance of these regression algorithms in general, but are specific to this particular application.

For the angled crack specimen, only the nearest neighbor and linear interpolation approaches were considered. Here, training Gaussian process models on a dataset of this size was infeasible due to the memory consumption imposed by the method, while linear regression was omitted due to poor accuracy. The performance of the algorithms in terms of accuracy and efficiency is displayed in Table 1. In this case, choosing the more superior method for surrogate modeling is not as straightforward. The linear interpolation models have nearly three times less error but are two orders of magnitude slower in terms of prediction time with respect to the nearest neighbors regressor. Based on the larger disparity in prediction speeds, the nearest neighbor models are selected over linear interpolation for the general crack characterization study in the angled crack specimen. However, the tradeoff in accuracy and speed for the two methods is investigated further in their application to diagnosis in Section 4.

3.5. Damage Localization

The Bayesian damage diagnosis framework was first applied to the problem of damage localization in the flat crack specimen. The MCMC algorithm described in Section 2.3 was used to sample the posterior probability distribution $p(c|\mathbf{d}^{\text{meas}})$ and estimate c_{flat} . A uniform distribution was used for the prior probability $p(c)$, simply enforcing the bounds in Equations 22 - 23 (i.e., 0 probability if either parameter falls outside the bounds in a given sample). The linear interpolation

Method	Mean Relative Error	Prediction Time
Nearest Neighbor	8.22×10^{-2}	8.43×10^{-4} sec
Linear Interpolation	2.86×10^{-2}	1.39×10^{-1} sec

Table 1. Performance comparison for the crack characterization surrogate models. The models were generated using 32076 training data points.

surrogate models described in the previous section were used to accelerate the evaluation of $p(\mathbf{c}|\mathbf{d}^{\text{meas}})$ during sampling. DIC strain data (Figure 3(a)) extracted at each of the three sensor arrays in Figure 2 were tested individually to compare the impact of measurement location on the resulting damage location estimates. The noise level, σ^2 , was prescribed according to Figure 4(a).

For all of the damage diagnosis results presented in this study, 11000 total samples were drawn using the MCMC algorithm. The first 1000 samples were discarded for the burn-in period after which a thinning interval of 10 was applied to reduce autocorrelation, yielding 1000 samples to produce estimates of damage location probability. The initial guess for sampling in each case was generated automatically using the approach in Equation 17. The covariance matrix Σ_q for the proposal distribution (Equation 11) was chosen such that the variance for each parameter was 10% of the size of its corresponding bounds in Equations 22 - 25, which resulted in sample acceptance rates of 9%, 56%, and 80% for sensor arrays 1, 2, and 3, respectively. For sensor array 1, Σ_q was scaled by 0.5 from here to increase the acceptance rate to 20% and reduce autocorrelation for the results to follow. The average solution time for the three cases was just 23.4 seconds.

Figure 6 shows the resulting crack location probability contours for each of the sensor arrays. It can be seen that sensor array 1, which is closest to the crack, provides an estimate that is nearly coincident with the true value with a high degree of certainty. Sensor array 2 is also highly accurate, in the sense that the highest predicted probability agrees with the true crack location, while the diagnosis using sensor array 3 is slightly less accurate. Of equal importance, it can be seen that Bayesian diagnosis framework effectively captures the increasing uncertainty in estimates with increasing distance between the damage and sensors, as indicated by the growing spread in the probability distributions. While it is clear that sensor proximity has a significant impact on the diagnoses provided, even sensor array 3, which is farthest from the crack, provides useful information about the nature of the damage present.

3.6. General Crack Characterization

The performance of the proposed diagnosis framework is now illustrated for general crack characterization in the angled crack specimen. That is, the probability distribution, $p(\mathbf{c}_{\text{angled}}|\mathbf{d}^{\text{meas}}) = p(x, y, a, \theta|\mathbf{d}^{\text{meas}})$, for unknown crack location, size, and ori-

entation was estimated using DIC strains (Figure 3(b)) at each of the sensor arrays in Figure 2. Nearest neighbor surrogate models were used to accelerate the sampling process for crack characterization following the developments in Section 3.4. The MCMC parameters remained unchanged from those provided for the damage localization results in the previous section and a uniform distribution was used again for the prior probability, $p(\mathbf{c})$. The noise level, σ^2 , was prescribed according to Figure 4(b). The sample acceptance rates were 24%, 39%, and 52% for sensor arrays 1, 2, and 3, respectively, while the average solution time was 356 seconds. Note the slower execution time versus damage localization is a result of an increase in prediction time of the nearest neighbor surrogate models trained with a larger training dataset and higher input dimension.

The results for general crack characterization for each of the three sensor arrays considered are shown in Figure 7. Here, the estimated marginal probability distributions for x , y , a , and θ are displayed along with the true values of these parameters from the angled crack specimen. Again, the trend of decreasing accuracy and increasing uncertainty is observed with increasing distance between sensor arrays used for diagnosis. The one outlier in this trend is the distributions of crack length in Figure 7(c), where the predictions for arrays 1 and 3 appear to be more similar than those for arrays 1 and 2 (which are closer in proximity). The exact cause of this anomaly is likely only explained by a more detailed investigation of the individual measurement errors at each sensor location and the correlations between the crack parameters under the joint posterior probability distribution, which is omitted here for brevity.

Generally speaking, it can be seen that each sensor array provides reasonably accurate predictions of the true values in that the points of maximum probability are near the true damage parameters in all cases. Predictions with sensor array 1, in particular, have a high degree of accuracy and relatively high precision while even sensor array 3 provides useful crack characterization information, more so in terms of the location of damage. Furthermore, it appears to be harder to predict the extent and nature of the damage with a high degree of certainty as compared with the location, as evidenced by the relative spread in the predicted distributions of $p(a)$ and $p(\theta)$ versus $p(x)$ and $p(y)$.




Figure 6. Crack location probability contours using the different sensor arrays from Figure 2 (denoted by white circles).

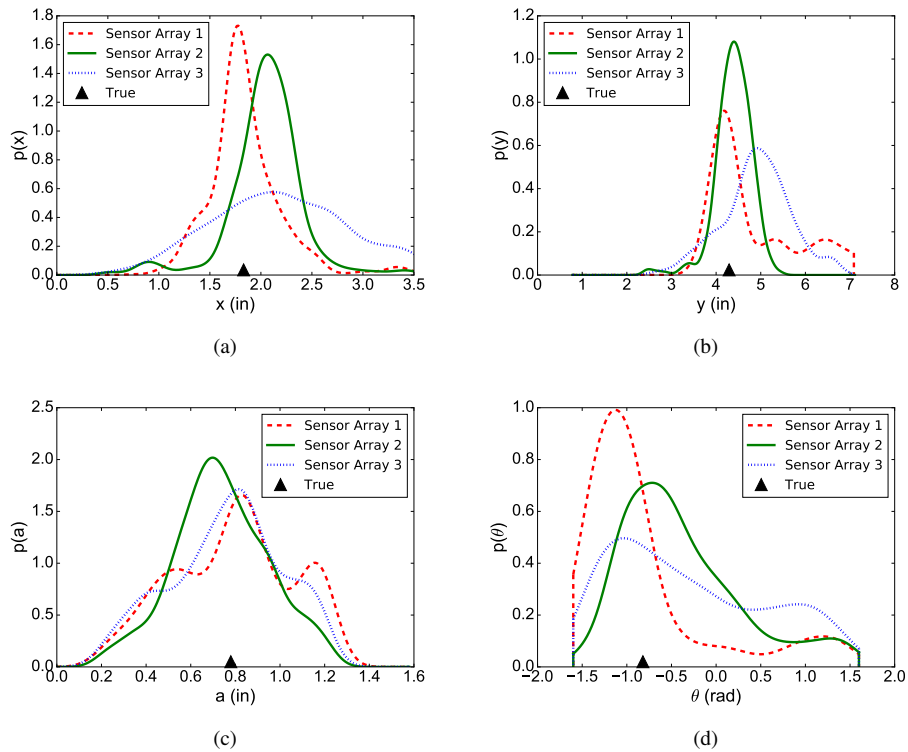


Figure 7. Crack characterization results for each sensor array. The marginal distributions of the (a) x -coordinate, (b) y -coordinate, (c) length, and (d) orientation of the crack compared to the true measured values.

4. COMPUTATIONAL EFFICIENCY

The computational efficiency of the surrogate-accelerated damage diagnosis framework is now illustrated in detail. First, the tradeoff in terms of computational speed and accuracy associated with surrogate modeling is presented. Then, the impact of generating an informed initial guess for MCMC according to Equation (17) on sampling convergence as opposed to doing so randomly is illustrated.

4.1. Surrogate Modeling vs FEM

In order to study the computational speedup and accuracy provided by surrogate modeling, the damage localization and crack characterization problems were also solved using the original FE model, \mathcal{M} , to evaluate $p(\mathbf{c}|\mathbf{d}^{\text{meas}})$ during sampling instead of $\widetilde{\mathcal{M}}$. All MCMC parameters for these analyses matched those used in the previous examples. Each evaluation of the FE models during sampling was executed in parallel on four CPUs. The run times and estimated probability distributions were stored for each sensor array for comparison. Furthermore, the crack characterization example was solved using linear interpolation models to assess the impact of choosing the less accurate, but faster nearest neighbor models in this case (Table 1).

A comparison of the run times and computational speedup using each model is displayed in Table 2 for damage localization and characterization. Note that the run times reported are average values of the analyses for each of the sensor arrays. Diagnosis using the FE model took over four days for each case. The execution time here (in contrast with using surrogate models) is independent of number of unknowns; the slight disparity observed between localization and characterization times is due to random variation in the crack parameters and computational meshes generated during sampling. It can be seen that tremendous computational speedup was provided by linear interpolation (14299X) and nearest neighbors (955X) for crack localization and characterization, respectively. If linear interpolation was used for crack characterization instead, only a 10X speedup would have been provided.

To assess any loss in accuracy associated with replacing the FE model with a surrogate, a comparison of the resulting probability distributions using each model is displayed in Figures 8 and 9 for damage localization and crack characterization, respectively. The crack location probabilities using linear interpolation show excellent agreement with the FE model in Figure 8, a trend that was seen previously in Figure 5(a).

More variation is seen among the surrogate model and FE model distributions in Figure 9 for crack characterization, illustrating the increased challenge of performing accurate regression in higher dimensions. Note that although there was a clear advantage in terms of accuracy for linear interpolation

over nearest neighbors in Table 1, the resulting diagnoses provided here by each is comparable. For practical situations, it is likely that the substantial benefit in computational speedup from surrogate modeling (Table 2) outweighs any relatively small inaccuracies in the diagnoses that were observed (Figures 8 and 9). Additionally, some of the disagreement observed can likely be attributed to minor statistical variations between the randomly-drawn, finite sample sets with MCMC used to generate solutions.

4.2. Initial Guess Effects

In this section, the benefit of generating an informed initial guess for MCMC using the surrogate training data (Equation 17) is briefly illustrated. Initializing sampling in a high probability region this way can decrease the required burn-in period and reduce the need for more advanced adaptive algorithms (Smith, 2013). In fact, the basic MCMC implementation presented in Section 2.3 was specifically utilized in this study to emphasize the latter.

In an attempt to further substantiate this claim, the difficulty of converging to the true damage probability distribution from a purely random guess is highlighted. To this end, a convergence study was performed where the damage localization problem was solved for each sensor array using 100 different random initial guesses. MCMC was performed for 2500 iterations with the same parameters that were used in Section 3.5. For each initial guess, it was recorded whether or not any samples had been drawn within a 0.5in radius of the true crack location during sampling. This criteria was selected to roughly reflect a point where the effect of the initial guess had been overcome and the burn-in period could be terminated. Note that random initial guesses that fell within the target area around the true location were not considered.

The results from this study are displayed in Figure 10 for (a) sensor array 1, (b) sensor array 2, and (c) sensor array 3. Here, the true crack location and target area are denoted by the black triangle and surrounding dashed circle, respectively. The green dots represent initial samples that ultimately reached this target area within 2500 iterations and the red \times 's are initial samples that did not reach the target area.

In Figure 10, only 9% of the random starting points led to samples within 0.5in of the true location for sensor array 1, 46% for sensor array 2, and 68% for sensor array 3. The low success rate is likely due, at least in part, to the existence of local maxima in the probability distribution where the Markov chain can become temporarily trapped during sampling. Furthermore, Figure 10 shows a clear trend between distance of the initial guess from the true solution and the likelihood that it will reach the area of interest. Here, the proximity of the initial guess is most critical in the case of sensor array 1, where the target distribution is most localized (Figure 6). In all cases, it is clear that a substantial burn-in pe-

Model	Damage Localization		Crack Characterization	
	Time (sec)	Speedup	Time (sec)	Speedup
Finite Element	3.64×10^5	1	3.40×10^5	1
Linear Interpolation	2.34×10^1	14299	3.56×10^4	10
Nearest Neighbor	-	-	3.56×10^2	955

Table 2. Computational Speedup comparison for surrogate modeling versus the original FE model.




Figure 8. Comparison of the damage localization results using the linear interpolation surrogate model versus the original FE model for sensor array 2.

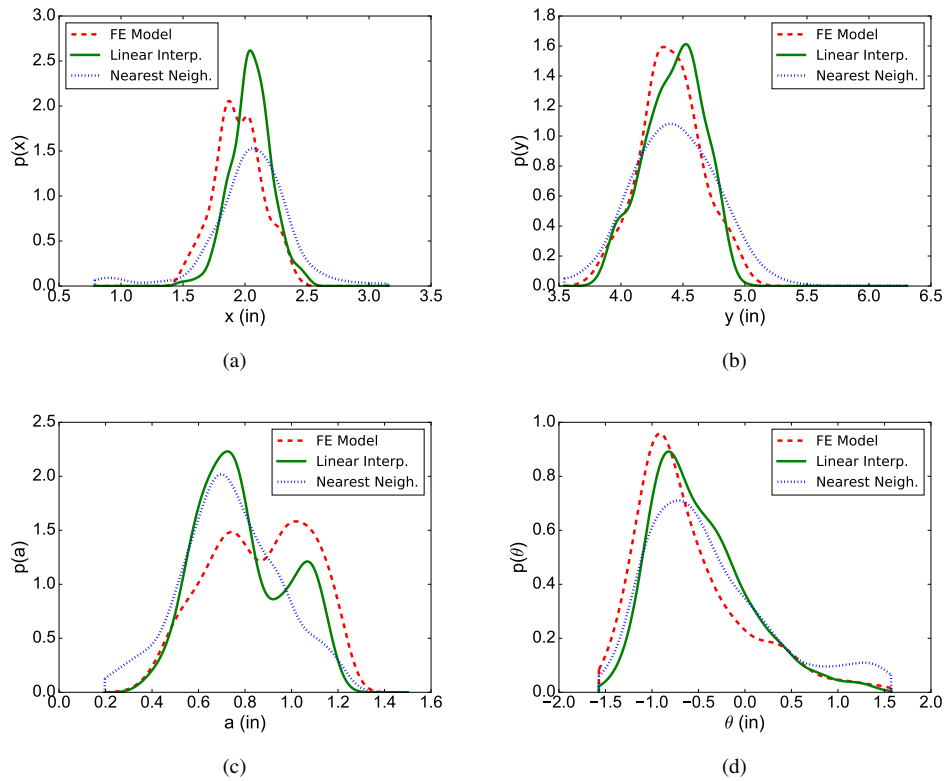


Figure 9. Crack characterization results using linear interpolation and nearest neighbor surrogate models compared with the original FE model for sensor array 2.

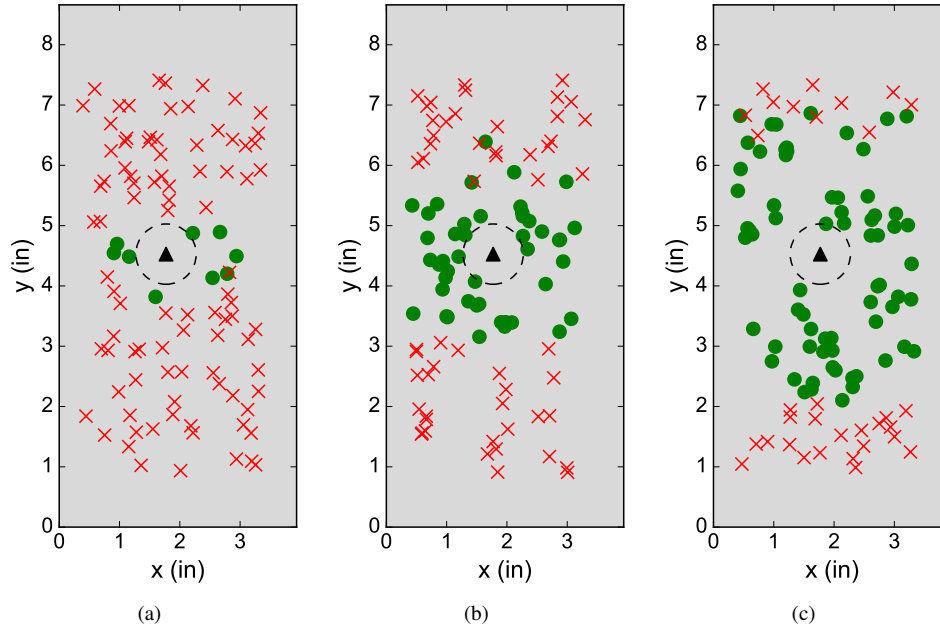


Figure 10. Convergence of sampling for 100 random initial guesses with (a) sensor array 1, (b) sensor array 2, and (c) sensor array 3. The green dots represent initial guesses that reached the target area (dashed circle) around the true solution (black triangle) within 2500 iterations and the red \times 's are those that did not reach the target area.

riod will often be required when randomly selecting a starting point for MCMC. By contrast, the initial guess generated via Equation 17 was 0.09in from the true location with sensor array 1, 0.04in away with sensor array 2, and 0.86in with sensor array 3, virtually eliminating the need for a burn-in period in these cases.

It is important to note that the observed effects of a poor initial guess could be reduced with extensive tuning of the MCMC algorithm used here or the introduction of a more advanced sampling approach. However, this study advocates the use of Equation 17 to initialize sampling as a simpler alternative when surrogate modeling is utilized. Note that while the convergence study was performed for localization and not characterization to aid in visualization, issues with convergence and poor initial guesses are exacerbated in higher dimensional cases making the scheme even more valuable in these cases.

5. CONCLUSION

In this study, a computationally-efficient, probabilistic damage diagnosis framework was presented and experimentally validated. Given measurement data, probability distributions of unknown damage parameters were estimated using Bayesian inference and MCMC sampling. Substantial computational speedup was obtained by replacing a three-dimensional FE model with an efficient surrogate model. While the proposed formulation is general for arbitrary component geometry, damage type, and sensor data, it was demonstrated on the problem of panel crack characterization using strain data determined

from DIC. Subsets of data were extracted in a fashion that mimicked monitoring with fiber optic strain gauges (a single component of strain along a linear array) for a more practically relevant application. The effectiveness of strain-based diagnosis was tested as the distance between the damage and these measurement locations increased.

The ability of the framework to efficiently perform probabilistic damage localization and characterization while capturing the uncertainty in the predictions as the measurement locations were varied was demonstrated. Furthermore, the use of a surrogate model to replace a 3D FE model was shown to yield average analysis times of 23.4 and 356.0 seconds for damage localization and full crack characterization, respectively, representing a 14299X and 955X computational speedup. Furthermore, a simple scheme to generate a highly probable initial guess for MCMC sampling using the surrogate model training data was shown to improve convergence and reduce the burn-in period needed. While the accuracy and certainty of the diagnosis results naturally degraded as measurement locations were moved further from the damage, this study reinforced the potential for strain sensors to allow for effective local SHM of hot spots in components. Additionally, the framework, capable of providing full crack characterization with UQ and computational efficiency, encompasses the necessary characteristics to enable subsequent damage prognosis.

Since the proposed diagnosis framework was presented here in a relatively basic form to emphasize the practicality of the

approach, several enhancements could be made to accommodate more complex applications. First, a sensitivity analysis could be performed to assess the relative identifiability of the damage parameters chosen for the available sensor measurements. This would be done both to discover and remove any higher order damage variables that could not be recovered during diagnosis and to help guide the selection of a training data grid for surrogate modeling. In regards to the latter, adaptive and selectively-refined training data grids could also be adopted in the framework to yield more efficient and accurate surrogate models with fewer evaluations of the original computational model. Furthermore, an explicit model discrepancy term could be included in the formulation to account for any bias or inability of the computational model to predict the measurement data after calibration. Finally, a variety of advanced MCMC algorithms can be deployed when the basic Metropolis algorithm presented here is too inefficient for sampling the resulting posterior probability distribution. This may be necessary for extensions of the diagnosis framework that seek more unknown variables (e.g., considering higher dimensional damage parameterizations or cases of unknown sensor noise levels and variable boundary conditions).

REFERENCES

- Barthorpe, R. J. (2010). *On model- and data-based approaches to structural health monitoring* (Unpublished doctoral dissertation). University of Sheffield.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. New York, NY: Springer.
- Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F., Mueller, A., Grisel, O., ... Varoquaux, G. (2013). API design for machine learning software: experiences from the scikit-learn project. In *Ecml pkdd workshop: Languages for data mining and machine learning* (pp. 108–122).
- Correlated Solutions Inc. (2012). *VIC-3D*. Retrieved from www.correlatedsolutions.com
- Farrar, C. R., & Worden, K. (2013). *Structural health monitoring: A machine learning perspective*. Wiley.
- Gamerman, D., & Lopes, H. F. (2006). *Markov chain monte carlo: Stochastic simulation for bayesian inference* (Second ed.). Boca Raton, Florida: Chapman and Hall/CRC.
- Global sensitivity analysis: the primer*. (n.d.).
- Haario, H., Laine, M., & Mira, A. (2006). Dram: Efficient adaptive MCMC. *Statistics and Computing*, 16(4), 339-354.
- Hochhalter, J. D., Krishnamurthy, T., Aguilo, M. A., & Gallegos, A. M. (2016). Strain-based damage determination using finite element analysis for structural health management. NASA/TM-2016-219186.
- Huhtala, A., & Bossuyt, S. (2011). A bayesian approach to vibration based structural health monitoring with experimental verification. *Journal of Structural Mechanics*, 44(4), 330-344.
- Isakov, V. (1998). *Inverse problems for partial differential equations*. New York: Springer.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001). *SciPy: Open source scientific tools for Python*.
- Kaipio, J., & Somersalo, E. (2004). *Statistical and computational inverse problems*. Springer.
- Katsikeros, C., & Labeas, G. (2009). Development and validation of a strain-based structural health monitoring system. *Mechanical Systems and Signal Processing*, 23(2), 372 - 383.
- Kehlenbach, M., & Hanselka, H. (2003, April). Automated structural integrity monitoring based on broadband lamb wave excitation and matched filtering. In *Proceedings of the 44th AIAA, ASME, ASCE, AHS, ASC Structures, Structural Dynamics, and Materials Conference*. Norfolk, VA.
- Kennedy, M. C., & O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3), 425–464.
- Kim, J. T., & Stubbs, N. (2002). Improved damage identification method based on modal information. *Journal of Sound and Vibration*, 252, 223-238.
- Krishnamurthy, T., & Gallegos, A. M. (2011, April). Damage characterization using the extended finite element method for structural health management. In *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 13th AIAA Non-Deterministic Approaches Conference*. Denver, CO.
- Leser, P. E., & Warner, J. E. (2017, January). A diagnosis-prognosis feedback loop for improved performance under uncertainties. In *19th aiaa non-deterministic approaches conference, aiaa scitech forum*. Grapevine, TX.
- Li, H.-N., Li, D.-S., & Song, G.-B. (2004). Recent applications of fiber optic sensors to health monitoring in civil engineering. *Engineering Structures*, 26(11), 1647 -

1657.

- Mal, A. K., Ricci, F., Banerjee, S., & Shih, F. (2005). A conceptual structural health monitoring system based on vibration and wave propagation. *Structural Health Monitoring*, 4, 283-293.
- Marzouk, Y. M., Najm, H. N., & Rahn, L. A. (2006). Stochastic spectral methods for efficient bayesian solution of inverse problems. *Journal of Computational Physics*, 224, 339-354.
- Meeds, E., & Welling, M. (2014). GPS-ABC: gaussian process surrogate approximate bayesian computation. *CoRR*, abs/1401.2838. Retrieved from <http://arxiv.org/abs/1401.2838>
- Meltz, G., & Snitzer, E. (1981, October 20). *Fiber optic strain sensor*. Google Patents. (US Patent 4,295,738)
- Moore, E. Z., Murphy, K. D., & Nichols, J. M. (2011). Crack identification in a freely vibrating plate using bayesian parameter estimation. *Mechanical Systems and Signal Processing*, 25, 2125-2134.
- Neiswanger, W., Wang, C., & Xing, E. (2013). Asymptotically exact, embarrassingly parallel MCMC. *arXiv preprint arXiv:1311.4780*.
- Nichols, J. M., Link, W. A., Murphy, K. D., & Olson, C. C. (2010). A bayesian approach to identifying structural nonlinearity using free-decay response: Application to damage detection in composites. *Journal of Sound and Vibration*, 329, 2995-3007.
- Nichols, J. M., Moore, E. Z., & Murphy, K. D. (2011). Bayesian identification of a cracked plate using a population-based markov chain monte carlo method. *Computers and Structures*, 89, 1323-1332.
- Peng, T., Saxena, A., Goebel, K., Xiang, Y., & Liu, Y. (2014). Probabilistic damage diagnosis of composite laminates using bayesian inference. In *16th AIAA Non-Deterministic Approaches Conference*.
- Peters, W. H., & Ranson, W. F. (1982). Digital imaging techniques in experimental stress analysis. *Optical Engineering*, 21(3), 427-431.
- Prudencio, E., Bauman, P. T., Faghihi, D., Ravi-Chandar, K., & Oden, J. T. (2015). A computational framework for dynamic data-driven material damage control, based on bayesian inference and model selection. *International Journal for Numerical Methods in Engineering*, 102(3-4), 379-403.
- Prudencio, E., & Cheung, S. H. (2012). Parallel adaptive multilevel sampling algorithms for the bayesian analysis of mathematical models. *International Journal for Uncertainty Quantification*, 2(3), 215-237.
- Python Software Foundation. (2016). *Python language reference, version 2.7*. Retrieved from www.python.org
- Roy, C. J., & Oberkampf, W. L. (2011). A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Computer Methods in Applied Mechanics and Engineering*, 200(2528), 2131 - 2144.
- Saltelli, A., Chan, K., & Scott, E. M. (2000). *Sensitivity analysis* (Vol. 1). Wiley New York.
- Sbarufatti, C., Manes, A., & Giglio, M. (2013). Performance optimization of a diagnostic system based upon a simulated strain field for fatigue damage characterization. *Mechanical Systems and Signal Processing*, 40(2), 667 - 690. doi: <http://dx.doi.org/10.1016/j.ymssp.2013.06.003>.
- Smith, R. C. (2013). *Uncertainty quantification: Theory, implementation, and applications*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
- Sutton, M. A., Orteu, J.-J., & Schreier, H. (2009). *Image correlation for shape, motion and deformation measurements: Basic concepts, theory and applications* (1st ed.). Springer Publishing Company, Incorporated.
- Vrugt, J. A., ter Braak, C. F., Diks, C. H., Higdon, D., Robinson, B. A., & Hyman, J. M. (2009). Accelerating markov chain monte carlo simulation by differential evolution with self-adaptive randomized subspace sampling. *International Journal of Nonlinear Science and Numerical Simulation*, 10(3), 273-290.
- Wang, J., & Zabaras, N. (2014). A bayesian approach to the inverse heat conduction problem. *International Journal of Heat and Mass Transfer*, 47, 3927-3941.
- Wang, L., & Yuan, F. G. (2007). Active damage localization technique based on energy propagation of lamb waves. *Smart Structures and Systems*, 3, 201-217.
- Warner, J. E., Bomarito, G. B., Heber, G., & Hochhalter, J. D. (2016). Scalable implementation of finite elements by NASA - implicit (SciFEi). NASA/TM-2016-219180.
- Warner, J. E., & Hochhalter, J. D. (2016). Probabilistic damage characterization using a computationally-efficient bayesian approach. NASA/TP-2016-219169.
- Warner, J. E., Hochhalter, J. D., Leser, W. P., Leser, P. E., &

- Newman, J. A. (2016, October). A computationally-efficient inverse approach to probabilistic strain-based damage diagnosis. In *Annual conference of the prognostics and health management society*. Denver, CO.
- Warner, J. E., Zubair, M., & Ranjan, D. (2017, January). Near real time damage diagnosis using surrogate modeling and high performance computing. In *19th AIAA non-deterministic approaches conference, AIAA scitech forum*. Grapevine, TX.
- Wilde, D. J., & Beightler, C. S. (1967). *Foundations of optimization*. Prentice-Hall Englewood Cliffs, N.J.
- Yan, G. (2012). A bayesian approach for identification of structural crack using strain measurements. In *Sixth European Workshop on Structural Health Monitoring*. Dresden, Germany.