

A Digital Twin Feasibility Study (Part II): Non-Deterministic Predictions of Fatigue Life Using In-Situ Diagnostics and Prognostics

Patrick E. Leser, James E. Warner, William P. Leser,
Geoffrey F. Bomarito, John A. Newman, and Jacob D. Hochhalter

May 23, 2019

Abstract

The Digital Twin (DT) concept has the potential to revolutionize the way systems and their components are designed, managed, maintained, and operated across a vast number of fields from engineering to healthcare. The focus of this work is the implementation of DT for the health management of fatigue critical structures. This paper is the second part of a two-part series. The first of the series demonstrated the use of multi-scale, initiation-to-failure crack growth modeling to form non-deterministic predictions of fatigue life. In this second part, a general method for reducing uncertainty in fatigue life predictions is presented that couples in-situ diagnostics and prognostics in a probabilistic framework. Monte Carlo methods and high-fidelity finite element models are used to (i) generate probabilistic estimates of crack state throughout the life of the same geometrically-complex test specimen and (ii) predict fatigue life with decreasing uncertainty as more of these diagnoses are obtained. The ability to predict accurately and in the presence of uncertainty is demonstrated, suggesting that the proposed DT method is feasible for fatigue life prognosis and should be pursued further with a focus on increasing application realism.

1 Introduction

Digital Twin (DT) is an emerging concept that has steadily gained popularity among industry, academia, and government agencies since it was first introduced in the early 2000s [1]. Borrowing from [2], the Defense Acquisition Glossary defines Digital Twin as “an integrated multiphysics, multi-scale, probabilistic simulation of an as-built system . . . that uses the best available models, sensor information, and input data to mirror and predict activities/performance over the life of its corresponding physical twin” [3]. By establishing an inherent requirement to operate at the leading edge of engineering capabilities, DT is

intended to be applicable to problems of arbitrary complexity. Research in the field focuses first and foremost on this requirement and uses enabling technologies such as high-performance computing to achieve these goals. A common question during implementation is thus not what model can be used, but how can the *best* model be used.

The idealistic goal of DT is to provide an exact digital replica of a real-world component, structure, or system – a digital twin. Through monitoring of the real-world counterpart, or physical twin, the digital twin can simulate both operational and environmental conditions that are identical to those that have been or may be experienced by the physical twin. Changes to the state of the physical twin (e.g. incurred damage) are detected, quantified, and added to the digital twin model. Assuming a perfect match between physical and digital, the digital twin can then simulate future missions, allowing for prediction of performance, risk, and maintenance requirements. In this ideal case, the only uncertainties are in the future state of the world in which the physical twin will operate. In reality, uncertainties are rife in the digital twin itself, present in both the models and the measurements as well as in the expected future operating conditions. The model is an approximation of the physical twin, and the loads applied to, environment around, and damage to the physical twin cannot be precisely measured. This missing knowledge manifests itself as uncertainty; uncertainty that in turn must be quantified and propagated through all predictive simulations.

While DT has become a common term, particularly in the last five years, its implementation for aerospace structures has proven difficult. In 2009, the first of a series of papers was published by Emery et al. on their concept, the multi-scale Damage and Durability Simulation strategy (DDSim) [4]. The work was an early step toward the realization of DT and embraced high-fidelity, physics-based modeling and the inherent probabilistic requirements of predicting damage evolution. The approach addressed issues associated with multiple length scales and computational expense through a variable-fidelity, hierarchical approach. In 2011, the DT concept was applied to fighter aircraft health management by the Air Force [5] and, in 2012 and 2015, the concept was included in NASA technology road maps [6, 7]. Work has also been published on DT applied to fleet management [8]. More recently, industry interest in the concept has grown, with GE being one of the most invested companies [9].

Building on preliminary NASA studies [10, 11, 12, 13, 14], the work presented in this two-part series was started in 2016 and focused on DT as a prognostics and health management approach for detecting, quantifying, and predicting fatigue crack growth. New methods were developed for characterizing damage and calibrating predictive models while quantifying uncertainty from a variety of sources. Crack growth was modeled from initiation to failure using high-fidelity finite element models. The multi-scale nature of fatigue crack growth was addressed in two phases, defined and separated by the minimum detectable flaw size in the component of interest. Prior to the regime of detectable damage, information about crack growth rate and material response were derived from empirical databases or microstructure characterization and simulation. This

phase is discussed in detail in Part I of this series [15].

The work presented herein focuses on the second phase of the DT predictive process developed under the project. In this context, fatigue cracks are now considered detectable. Current detection capabilities link this regime with that of microstructurally-large crack growth. The advantage of the second phase is that material and crack growth rate parameters can be calibrated based on observations of actual crack growth, dramatically reducing the uncertainty in the predictive model. One of the early and most complete outlines of this concept in its contemporary form was provided in [16]. The approach is straightforward and generally applicable. As temporal snapshots of crack state are gathered from a diagnostic system, the data is fed through an uncertainty quantification framework that calibrates the parameters of a high-fidelity fatigue crack growth simulator. Once calibrated, the uncertainty in model parameters and measurements is propagated forward through the model to a future time of interest (e.g., next maintenance appointment) or until the crack reaches a critical length (e.g., impending failure). The result is a probabilistic prognosis of the time to failure or the remaining useful life.

Realization of this iterative diagnosis-to-prognosis process is the subject of abundant foundational research in both the structural health monitoring and prognostics and health management fields. Early works involved calibration of analytical crack growth models using observations of damage evolution, either in-situ or from tests conducted offline [17, 18, 19, 20]. Methods for quantifying uncertainty varied in early contributions. Orchard et al. provided an accessible introduction to particle filters as a natural framework for state-space-based fatigue crack growth prognostics [21]. A number of notable studies incorporating particle filters followed, including [22, 23, 24]. Difficulties estimating static model parameters, such as those defining crack growth rate laws, motivated the use of Markov chain Monte Carlo (MCMC) methods as an alternative to particle filters [25, 26, 27, 28].

DT is predicated on the use of high-fidelity models, which conflicts with the use of quick-to-evaluate, analytical models in the aforementioned contributions. Reduction of model fidelity alleviated the computational burden associated with traditional uncertainty quantification methods. Frameworks like DDSim [4] and DARWIN [29, 30, 31], the probabilistic reliability software developed by Southwest Research Institute, enable the use of high-fidelity damage models and are more similar in scope to that of DT. Grell et al. demonstrated fatigue prognostics centered around the Air Force damage tolerance framework AFGROW [32]. Bazilevs et al. demonstrated the use of the Dynamic Data-Driven Application System (DDDAS) to predict fatigue crack growth in a wind turbine blade [33]. While these studies exhibited more DT-like fidelities, they were either deterministic or primarily propagated uncertainty (i.e., preclude parameter estimation), falling short of a continuously updated digital replica of a physical system.

The aim of this paper was to demonstrate the feasibility of DT as a complete health management philosophy, specifically in the regime of detectable damage. Building upon the literature, a rigorous DT framework was developed with a focus on non-deterministic calibration of high-fidelity models and uncer-

tainty propagation for the purpose of forming a continuously-updated fatigue life prognosis. While the scope of the work was limited to fatigue crack growth, the means for quantifying and propagating both diagnostic and prognostic uncertainties are intended to be generally applicable to damage in any structure. The only requirements for the approach are observations of material response that are sensitive to damage evolution and a reliable model that can predict that response. The remainder of this paper is organized as follows. Section 2 introduces the methods used to diagnose and prognose fatigue crack growth, including methods for uncertainty quantification. In Section 3, relevant information about the DT study, namely the specimen and test used to demonstrate feasibility, is provided. Details regarding the implementation of the methods are presented in Section 4, followed by results and discussion in Section 5 and concluding remarks.

2 Method

Fatigue crack growth prognostics were implemented in two phases as part of the DT framework. The separation of the two phases is defined by the minimum detectable flaw size (MDFS). Transition to the second phase commences once the actual flaw can be diagnosed. The phases are referred to as the pre- and post-diagnosis regimes herein. Due to the impracticality of micro-scale material characterization at every point in the structural volume, the pre-diagnosis phase generally yields predictions with higher uncertainty than its post-diagnosis counterpart. The latter benefits from direct observations of damage growth. However, a significant portion of fatigue life is generally spent in the pre-diagnosis regime, especially for high-cycle fatigue, necessitating a DT approach comprising both phases.

The pre-diagnosis regime involves simulating crack growth from initiation to failure with all uncertainty quantified empirically from microstructure statistics and assumptions regarding loads and operational environment. The process, outlined in [15], involves a two-scale model of microstructurally small crack (MSC) and microstructurally large crack (MLC) growth. First, crack initiation and microstructurally small, grain-by-grain crack growth are simulated until the crack reaches a characteristic transition length. Referred to as the handshake crack length in that paper, this transition defines the shift to the MLC model. Linear elastic MLC growth is then simulated from the handshake crack length, through a homogeneous material, and to satisfaction of a predefined failure criteria. The final predicted number of cycles to failure is then computed as

$$N_f = N_{MSC} + N_{MLC} \quad (1)$$

where N_{MSC} and N_{MLC} are random variables representing the number of cycles required to initiate and grow a crack to the handshake crack length and the number of cycles to grow from the handshake crack length to failure, respectively.

The post-diagnosis regime differs in two ways from its counterpart. First, since actual observations of the crack (e.g., crack position, length, orientation) are available by definition, a model calibration step can be added to the prediction process. These observations inform and reduce uncertainty in the crack state and crack growth rate model parameters. The reduced parameter uncertainties and quantified measurement uncertainties are then propagated forward through the simulation until the crack growth reaches the predefined stopping condition. The model calibration step significantly reduces uncertainty in the final quantity of interest, N_f , when compared to the pre-diagnosis regime prediction.

The second difference between regimes is that the number of cycles required to reach the first diagnosis is now known. The crack measurement at that cycle count is a random variable, but the cycle count itself is deterministic. This point is significant as it reframes how N_f is calculated. Now,

$$N_f = n_D + N_f^R \quad (2)$$

where n_D is the deterministic number of cycles sustained by the monitored component until the first diagnosis of crack length is made. There is no need to compute the number of cycles to initiate and grow to a handshake length any more, as this is a fixed quantity and included in n_D (plus the number of cycles from handshake to the first diagnosis).

The role of prognostics in the post-diagnosis regime is thus the estimation of the random variable N_f^R , or the number of cycles between the first diagnosis and failure. By definition, diagnoses are available with which to calibrate the large crack model, leading to a reduction of uncertainty in N_f^R over time. As part of this prediction, the initial crack length for the MLC simulation is now a random variable (in contrast to the deterministic handshake length in the pre-diagnosis regime). This variable is unknown, but can be estimated as part of the model calibration.

The remainder of this section is divided into three parts, all pertaining to the topic of this paper – the post-diagnosis regime of the DT approach. In the first subsection, a general overview of uncertainty quantification methods are discussed. The same general methods were used by the diagnostic and prognostic algorithms, which are discussed in further detail in the remaining two subsections, respectively.

2.1 Uncertainty quantification

Both the diagnostic and prognostic approaches used herein involve calibration of a model, \mathcal{M} , with random parameters, X , given observations of the physical response being modeled. Mathematically,

$$Y_i = \mathcal{M}_i(X) + \varepsilon_i \text{ for } k = 1, \dots, K \quad (3)$$

where Y_i and ε_i are both random variables and represent the observations of the system and the measurement noise, respectively. K denotes the total number

of observations available.

The calibration is the solution to the inverse problem posed by Equation 3. Given realizations y_i of Y_i for $k = 1, \dots, K$, determine X . To accomplish this, the solution to the inverse problem can be formulated using Bayes' Theorem as

$$\pi(x|y) = \frac{\pi(y|x)\pi(x)}{\int_{\mathbb{R}^q} \pi(y|x)\pi(x)dx} \quad (4)$$

with $\pi(x|y)$ referred to as the parameter posterior probability density function (PDF). The numerator is the multiplication of two densities, the likelihood function, $\pi(y|x)$, and the prior distribution, $\pi(x)$. The denominator is a normalizing constant ensuring that the integral of the posterior PDF over \mathbb{R}^q equals unity. For most engineering applications, the evaluation of this normalizing constant is intractable. Instead, the posterior is typically approximated using sampling methods such as Markov chain Monte Carlo (MCMC), which was the approach used in this work.

Detailing the inner-workings of MCMC is beyond the scope of this article. In summary, MCMC involves the construction of a Markov chain through the parameter space that has a stationary distribution equal to the targeted posterior distribution [34]. Drawing a finite set of samples $\{x^{(\lambda)}\}_{\lambda=1}^{\Lambda}$ from the chain in this stationary condition is thus equivalent to drawing samples from the posterior. These samples can be used to estimate statistics or approximate the PDF of the posterior. As with other Monte Carlo methods, these approximations converge asymptotically as $\Lambda \rightarrow \infty$.

The formation of the Markov chain takes advantage of the invariance of the normalizing constant with respect to x . By assuming that the measurement errors are independent and identically distributed (iid) and $\varepsilon_i \sim N(0, \sigma^2)$, the likelihood becomes,

$$\pi(y|x) = \frac{1}{(2\pi\sigma^2)^{K/2}} \exp\left(-\frac{1}{2\sigma^2} \sum_{k=1}^K [y_i - \mathcal{M}_k(x)]^2\right), \quad (5)$$

a function of the sum of squared errors between the model response and the corresponding observations. The only other dependence of Equation 5 is the variance term, σ^2 . If this value is unknown, it can be added to the random parameter vector, X , and estimated along with the model parameters. The prior distribution is assumed to be known, derived from expert knowledge, empirical data, or previous calibrations. Therefore, the posterior can be evaluated up to some proportionality for any $x \in \mathbb{R}^q$.

The Markov chain is constructed sample-by-sample. A new sample, $x^{(*)}$, is proposed from a proposal distribution, J , that is dependent only on the previous sample; i.e., $x^{(*)} \sim J(x^{(*)}|x^{(\lambda-1)})$. While a variety of algorithms exist, acceptance or rejection of the proposed sample generally depend on the acceptance ratio, defined as

$$\mathcal{R} = \frac{\pi(x^{(*)}|y)}{\pi(x^{(\lambda-1)}|y)} = \frac{\pi(y|x^{(*)})\pi(x^{(*)})}{\pi(y|x^{(\lambda-1)})\pi(x^{(\lambda-1)})}. \quad (6)$$

Samples which have a higher posterior probability than the previous sample ($\mathcal{R} \geq 1$) are accepted into the chain while samples with lower posterior probability ($\mathcal{R} < 1$) are accepted with probability \mathcal{R} and rejected with probability $1 - \mathcal{R}$. If the proposed sample is accepted, $x^{(\lambda)} = x^{(*)}$ and if rejected, $x^{(\lambda)} = x^{(\lambda-1)}$. The process of proposing and accepting or rejecting new samples is iterated until the number of samples in the chain is equal to Λ .

A few challenges remain, such as how to initialize the Markov chain. Randomly sampling $x^{(1)} \sim \pi(x)$, for example, means that the chain initially is non-stationary as it searches for the region of high posterior probability. Additionally, the proposal distribution, J , is typically tuned throughout sampling to encourage chain mixing. Both of these details result in a period of non-stationarity at the beginning of the sampling process. To remove this effect when generating the sample-based estimate of the posterior, an initial portion of the chain is typically discarded. This is referred to as chain burn-in. In this work, deterministic optimization was performed prior to sampling such that the chain was always initialized at the optimal parameter vector in a least squares sense. For single-mode posterior densities, this step largely removes the need for a global search and reduces the number of samples required for burn-in. Finally, convergence of the Markov chain to its stationary condition must be achieved in order to generate samples from the target posterior. Ensuring chain convergence is important and methods for monitoring chain convergence have been researched for over two decades [35, 36]. Readers are referred to [37, 34, 38] for more information regarding MCMC sampling.

2.2 Diagnostics

The role of diagnostics in the DT framework is to periodically assess the state of damage (e.g., a fatigue crack in this work) after it has exceeded the MDFS using in-situ material response measurements. The crack state estimates over time then represent observation data that facilitate prognostics. By continually generating new diagnoses and re-calibrating the prognostics model based on these observations, the DT framework can significantly reduce the uncertainty in predictions of the final number of cycles (Equation 2).

Diagnostic methods operate under the assumption that the physical response of a structural component is altered in the presence of damage. To this end, the goal of diagnostics here is to use a set of measured responses at time t , $\{\hat{e}_{t,q}\}_{q=1}^Q$, to estimate parameters that characterize the crack (location, size, etc.) at that time, \mathcal{C}_t . A computational model, \mathcal{M}^D , capable of predicting the material response for an arbitrary crack configuration is also needed.

The observed response data and predicted responses using the model are then related by

$$E_{t,q} = \mathcal{M}_q^D(\mathcal{C}_t) + \varepsilon_q^D \text{ for } q = 1, \dots, Q \quad (7)$$

where ε_q^D is a random variable representing the measurement noise associated with the sensors used. Note that it is assumed that any additional free parameters of \mathcal{M}^D (material properties, boundary conditions, etc.) have been properly

calibrated beforehand and that necessary post-processing is performed within the model to return the material response at the Q sensor locations.

Equation 7 forms the basis of the diagnostics inverse problem that can be solved by leveraging the approach described in Section 2.1. The result is a sample-based estimate of $\pi(\mathbf{c}_t|e_t)$, the posterior PDF of the crack parameters given a set of material response observations at a particular time. For the purposes of DT, $\pi(\mathbf{c}_t|e_t)$ is used to produce a scalar estimate of the crack state that can serve as a single observation for prognostics. In this work, the mean value of \mathbf{c}_t is used

$$\hat{\mathbf{c}}_t = E[\mathbf{C}_t|e_t] = \int \pi(\mathbf{c}_t|e_t)\mathbf{c}_t d\mathbf{c}_t. \quad (8)$$

Diagnostics is performed at T instances throughout time to obtain the set of crack state observations, $\{\hat{\mathbf{c}}_t\}_{t=1}^T$, needed to perform calibration and prediction with prognostics. For more details on the diagnostic algorithm, see [13] for an in depth formulation.

2.3 Prognostics

The goal of the DT implementation herein is the formation of a prognosis of fatigue cycles to failure, N_f . This article focuses on the post-diagnosis phase in which the crack size has exceeded the MDFS. From Equation 2, the primary task is the estimation of N_f^R given n_D and an accumulating series of diagnoses of crack state.

Formation of this prognosis requires a predictive model, \mathcal{M}^P , capable of simulating the observed crack growth and dependent on a set of random parameters, Θ . Uncertain model parameters can be calibrated given a set of observations of the crack state, $\{\mathbf{c}_i\}_{i=1}^T$ using the approaches discussed in Section 2.1. The relationship between the predictive model and observations is given by

$$\mathbf{C}_t = \mathcal{M}_t^P(\Theta) + \varepsilon_t^P \text{ for } t = 1, \dots, T \quad (9)$$

where ε_t^P is a random variable representing the measurement noise.

The output of the calibration process is a sample-based estimate of the parameter posterior PDF, $\pi_t(\theta|\mathbf{c}_t)$, that represents the uncertainty in the model. The uncertainty in the measurements is defined by σ^2 , which can be fixed if known or included in θ and estimated as discussed in Section 2.1. With uncertainties quantified, the next step is to propagate them through the model to estimate the quantity of interest, N_f^R .

Monte Carlo simulation (MCS) was used to propagate the quantified uncertainties back through the model and arrive at a probabilistic prediction of the final quantity of interest. In the context of the proposed approach, MCS involves drawing $\Gamma \leq \Lambda$ samples from the Markov chain constructed as part of the model calibration step. Each sample is used to simulate crack growth to a failure criterion. The number of fatigue cycles to reach failure is stored for each simulation. The resulting set, $\{[n_f^R]_\gamma\}_{\gamma=1}^\Gamma$, is a sample-based approximation of the random quantity of interest, N_f^R , conditioned on all available data at time

t. It can be shown that this approximation converges asymptotically as $\Gamma \rightarrow \infty$. Substituting this result into Equation 2 yields the prognosis of interest.

3 Test Setup

The objective of the feasibility study was to develop a DT approach that could increase reliability in novel designs and accelerate their certification. A motivating example was the Boeing truss-braced wing concept [39, 40]. The high-aspect ratio, truss-supported wing is considered an N+3 design, or one that is three generations ahead of current state-of-the-art. Stringent certification requirements resulted in over-design of fatigue-critical components and degradation of aircraft performance. By implementing a DT-based health management system, a reduction in uncertainty in remaining fatigue life could be leveraged to relax these requirements and realize the full potential of the design. The use case was ideal as it would require state-of-the-art modeling and uncertainty quantification techniques, staying true to the DT concept.

Working directly with the truss was impractical, and so a custom specimen was designed and fabricated that maintained the primary features of the more complex part (e.g., thickness changes, counter sunk bolt holes, and a variety of other crack initiators). The primary advantage of this representative specimen was the ease with which it could be tested in a standard uniaxial servo-hydraulic test machine. A photograph of the specimen, which was machined from aluminum alloy (AA) 7075-T651 plate, is shown in Figure 1(a), with dimensions displayed on the diagram in Figure 1(b). As shown, there were six evenly-spaced, counter sunk holes running down a raised rib-like feature that was slightly offset from center. The regions on either side of the rib were of equal thickness. A large notch was machined on the right side of the specimen to add additional stress concentrators. These features culminated in a complex, asymmetric geometry that had fourteen stress concentrations – two at each hole and one at each fillet of the notch. The complexity of this structure ensured that crack growth would be non-planar with non-constant crack front dimensions.

To initiate a fatigue crack, the test specimen was subjected to constant-amplitude fatigue loading with a maximum applied force of $\mathcal{F} = 53.38$ kN and a minimum to maximum load ratio, $R = 0.1$. Load cycles were applied at a frequency of 10 Hz. The choice of loading was based on previous tests that had resulted in reasonable crack initiation times for prototype specimens. The final geometry was chosen based on finite element analysis of the specimen. The location and size of the six holes, horizontal rib location and width, and dimensions of the notch were varied until the strains at each concentrator were equal and near yield.

Crack growth was measured with an optical microscope on a traveling stage from the rear face, or flat side, of the specimen. A pattern was applied to the front face of the specimen for digital image correlation (DIC), as shown in Figure 1(a). DIC is a method for three-dimensional, optical tracking of point-wise motion of a random pattern. When that random pattern is applied to the

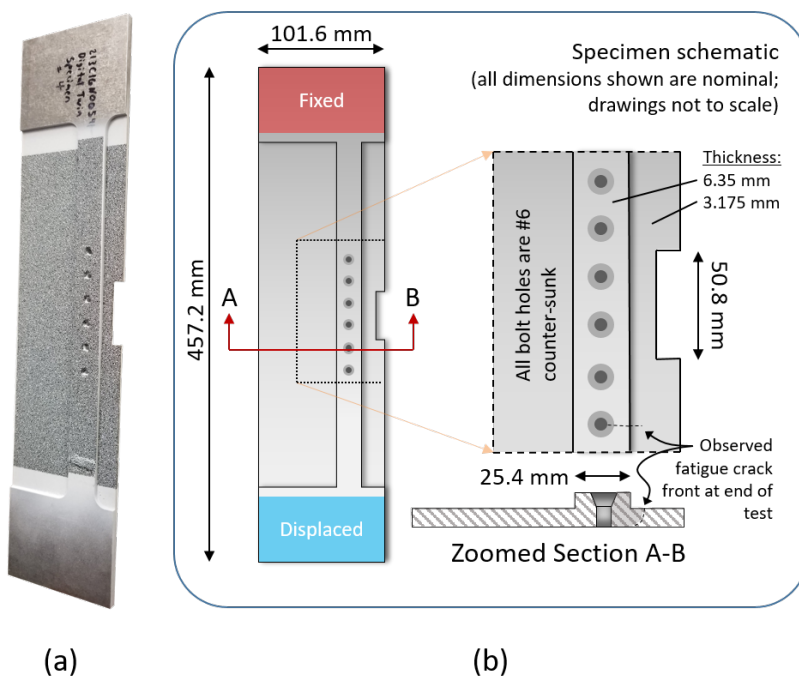


Figure 1: (a) Photograph and (b) diagram of the six-hole, DT specimen. The dimensions shown are measurements of the as-built specimen.

surface of a test specimen, DIC is capable of measuring quantities such as full-field deformation, vibration, and strain [41, 42, 43]. In this work, an optimized pattern was applied to the specimen surface for improved accuracy [44], and DIC was used to measure full-field strain on the front surface of the specimen during the test.

Under fatigue loading, a crack initiated at the bottom bolt hole as shown in Figure 1(b). Once the crack initiated, the load was decreased to a maximum force of $\mathcal{F} = 44.48$ kN with $R = 0.1$. Measurements of crack length and full-field strain were captured every 1000 cycles. At each measurement, the strain was used to form a diagnosis of crack length while the visual measurements were used for validation. Additionally, at each measurement, a prognosis of future life was generated. The cyclic process of diagnosis and prognosis was carried out postmortem; that is, the DT was not demonstrated in real-time. Instead, the experiment was replayed using the saved data, which aided in the development and validation of the proposed method without having to conduct new tests.

4 Implementation

By design, the complexity of the test specimen geometry demanded a high-fidelity model capable of simulating the displacement response under applied loads. The same finite element model was used for both the diagnosis of fatigue cracks and the prognosis of fatigue crack growth, although the implementation of the model was slightly different for each. The purpose of the diagnostic model was to return strains at designated locations for an arbitrary crack configuration. To enable this capability, the model was coupled with `gmsh` [45], a publicly available software for generating three-dimensional meshes, for rapid insertion of cracks at prescribed locations and lengths. The prognostic model was coupled with an additional fracture mechanics software, `FRANC3D` [46], which was used to simulate fatigue crack growth. Due to computational expense, surrogate models¹ were used when performing MCMC sampling in both the diagnostic and prognostic steps.

The remainder of this section details the finite element model, the crack insertion approach used for diagnostics, and the crack growth simulation approach used for prognostics. The surrogate models for diagnostics and prognostics are discussed in the respective subsections.

4.1 Finite Element Model

A finite element model was developed based on the as-built dimensions of the specimen shown in Figure 1. Material properties were considered deterministic and were fixed at literature reported values for aluminum alloy 7075-T651.

¹It should be noted that advancements in future computing capability could obviate the need for surrogate modeling. In general, using full-order models is preferred. However, it should also be noted that researchers tend to maximize the capabilities available to them, and surrogate models may still be needed in the future as a means for extending beyond the state-of-the-art in computing.

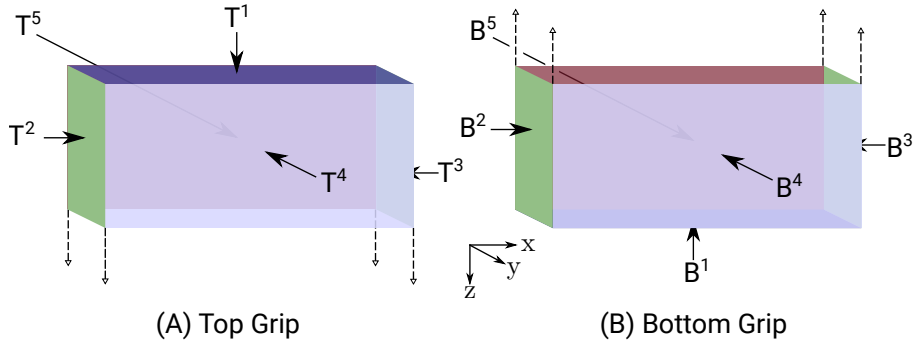


Figure 2: Diagram of the faces to which boundary conditions were applied on the (a) top grip and (b) bottom grip.

| Geometry Set | x -displ. (mm) | y -displ. (mm) | z -displ. (mm) | z -traction (MPa) |
|-------------------|---------------------|---------------------|---------------------|------------------------|
| T^2 | - | 0.6350 | - | - |
| T^3 | - | 0.4064 | - | - |
| $\{T^i\}_{i=1}^5$ | 0.1905 | - | 0.0 | - |
| B^1 | - | - | - | 78.05 |
| B^2 | - | 0.0 | - | - |
| B^3 | - | 0.4064 | - | - |
| $\{B^i\}_{i=1}^5$ | -0.4445 | - | - | - |

Table 1: Calibrated boundary conditions applied to the modeled grips.

Specifically, Young’s modulus and Poisson’s ratio were set to 10,400 ksi and 0.33, respectively [47]. The model consisted of 138,729 ten-node tetrahedral elements.

Displacement boundary conditions were defined on the surface nodes of the grip regions; i.e., the highlighted fixed and displaced regions in the schematic of Figure 1. Ten node sets were defined on the grips as illustrated in Figure 2. Due to stand misalignment, the boundary conditions on each set were unknown. To remedy this, DIC was used to capture full-field displacements on the surface of the component at multiple loads during the first fatigue cycle. This data was used to calibrate the node set displacement boundary conditions such that the sum of squared errors between modeled displacements and observed displacements was minimized. The resulting boundary conditions are shown in Table 1.

4.2 Strain Prediction for Diagnosis

In order to implement the approach outlined in Section 2.2, the diagnostic model, $\mathcal{M}^{\mathcal{D}}$, needed to be capable of predicting strains at prescribed locations in the domain for a given crack state. The calibrated finite element model described above was augmented with a `gmsh` pre-processor that allowed for insertion of arbitrary crack configurations, \mathbf{c} , into the mesh for this purpose. The crack was assumed to be a flat through-crack oriented such that its surfaces were in the x - y plane. The crack front was straight, aligned parallel with the y -axis. Cracks could be inserted at any of the 14 crack initiation sites, and changes in thickness of the modeled specimen were automatically compensated for by the `gmsh` script. With these assumptions, the crack state for the diagnostic model could be characterized with just two parameters

$$\mathbf{c} = [l, a], \quad (10)$$

where $l \in [1, 14]$ is an index parameter specifying the crack location and a is the crack length.

After crack insertion and meshing with `gmsh`, the finite element model could be executed under the calibrated boundary conditions to yield the strain field throughout the test specimen geometry in Figure 1. A post-processing script was then written in Python to extract strain values at discrete points in the domain to accommodate an arbitrary set of sensor locations for diagnosis. The calibrated finite element model along with these pre- and post-processing routines represent the complete implementation of the diagnostic model, $\mathcal{M}^{\mathcal{D}}$, for strain prediction as required in Equation 7.

In general, a diagnostic framework should provide results quickly enough to enable an actionable prognosis. The execution time for a single evaluation of the diagnostic model in this study was approximately one minute using a single core of a 2.8GHz Intel Core i7 processor. Since $\sim O(10^4)$ samples (and model evaluations) are typically required with MCMC for the approach described in Section 2.2, the resulting time to complete a diagnostic analysis was deemed impractical. To reduce the computation times for DT, a surrogate model, $\widetilde{\mathcal{M}}^{\mathcal{D}}$, was introduced to approximate the strain predictions using the finite element model in a computationally efficient manner.

More specifically, an individual surrogate model was generated for each crack initiation site that could return the estimated strain at each sensor location for a given crack size. That way, the surrogate modeling approach was reduced to a simple interpolation problem in one dimension given a particular site l^*

$$\widetilde{\mathcal{M}}_q^{\mathcal{D}}([l^*, a]) = \mathcal{I}_q^*(a) \quad (11)$$

where \mathcal{I}_q^* is a linear interpolant for the strain at sensor q for a crack at that site. Data used to construct each interpolant was obtained by executing the finite element model for a set of crack lengths, $\{\bar{a}^p\}_{p=1}^P$, and storing the resulting strains at each sensor location

$$\mathcal{M}_q^{\mathcal{D}}([l^*, \bar{a}^p], \text{ for } p = 1, \dots, P. \quad (12)$$

For this work, $P = 56$ and the crack lengths varied in even increments from 1.27 mm to 16.256 mm, the latter of which was approximately 50% greater than the failure crack length.

The surrogate models were generated in this manner and stored so that they could be leveraged to accelerate subsequent diagnostic analyses. Here, the speedup for a strain prediction using the linear interpolation approach rather than an FEA solve was approximately $300,000\times$. Furthermore, the strain values calculated with the finite element model and surrogate models were found to match within 1% on average.

4.3 Fatigue Crack Growth Simulation for Prognosis

For prognostics, the base finite element model was wrapped in a fatigue crack growth simulation framework. The linear elastic fracture mechanics code FRANC3D (FRacture ANalysis Code, 3D) [46] was used for this purpose. FRANC3D enables the simulation of arbitrary, non-planar, three-dimensional crack growth. It interfaces with multiple commercial codes as well as ScIFEN (Scalable Implementation of Finite Elements by NASA) [48], NASA’s in-house, parallel finite element code. FRANC3D also has a Python Application Programming Interface (API), making it easy to incorporate into the overall prognostic workflow.

4.3.1 Deterministic Fatigue Crack Growth Simulation

Using FRANC3D and the finite element model described in Section 4.1, the deterministic simulation of growth for a single crack was carried out as follows. An initial crack of a given shape and size was explicitly inserted into the mesh topology at a specified location. The continuous, three-dimensional crack front was discretized, comprising a subset of J nodes in the finite element mesh. The shape used in this work was a circular sector of the form shown in Figure 3. This initial crack was planar with crack surfaces parallel to the x - y plane. The radius of the initial crack for a given simulation is denoted by the crack length $a^{(0)}$, where the superscript indicates the index of the crack step in the simulation. In general, the central angle ϕ of the sector was defined such that the crack always extended to free surfaces, regardless of whether it was a corner crack or surface crack. For the DT demonstration described in this work, the crack initiation site was assumed known and fixed at a corner. The initial crack plane was vertically aligned with the midpoint of the bottom countersunk hole with center of the the sector at the countersink elbow, as shown in Figure 4.

Once initialized, Abaqus [49] was used to solve the finite element equations derived from the cracked mesh and specified boundary conditions. Nodal displacements were used to compute all three stress-intensity factor modes along the discretized crack front using the M-integral [50]. Growth of the arbitrary crack front was then simulated by first estimating the crack growth rate at each node along the crack front,

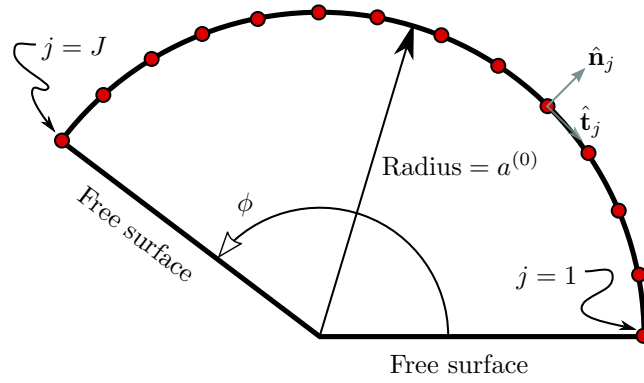


Figure 3: Arbitrary initial crack used for crack growth simulations.




Figure 4: Top view of the crack path simulated by FRANC3D.

$$\frac{da}{dn_j} = F(\mathbf{K}_j, \dots), \quad (13)$$

where the function F defines the pseudo-instantaneous relationship between crack growth rate and crack driving forces (e.g., the stress-intensity factors, $\mathbf{K} = [K_I, K_{II}, K_{III}]$). In this work, the crack growth rate behavior was expected to be log-linear, so Paris' Law was adopted such that

$$F = \beta(\Delta K)^m, \quad (14)$$

where β and m are material constants and invariant over the crack front and $\Delta K = (1 - R) \times K_{I,\max}$ for a single fatigue cycle.

Crack extension was approximated using a quasi-static median extension approach. Point-wise extension was defined as

$$\Delta a_j = \Delta a_{\text{median}} \left[\frac{K_{I,j}}{K_{I,\text{median}}} \right]^2, \quad (15)$$

where the user-defined median extension $\Delta a_{\text{median}} = 0.1016$ mm and $K_{I,\text{median}}$ was the median stress-intensity factor in the set $\{K_{I,j}\}_{j=1}^J$. The crack is advanced by applying the extension Δa_j in the crack propagation direction, which is the direction of the crack front normal unit vector at each point, $\hat{\mathbf{n}}_j$, rotated about the tangent, $\hat{\mathbf{t}}_j$, by a computed kink angle computed using Maximum Tangential Stress Theory [51].

The number of cycles consumed by the growth between steps s and $s + 1$ were calculated through integration of Equation 13,

$$\Delta n_j^{(s)} = \int_{a_j^{(s)}}^{a_j^{(s+1)}} [F^{-1}(a_j)] da = \frac{1}{\beta} \int_{a_j^{(s)}}^{a_j^{(s+1)}} [\Delta K_j(a_j)]^{-m} da. \quad (16)$$

Evaluating the integral over crack extensions requires F to be defined as a function of a_j over the crack length extension. In this work, $\Delta K(a_j)$ was assumed to be a monotonic, linear function passing through $\Delta K(a_j^{(s)})$ and $\Delta K(a_j^{(s+1)})$. While the true relationship is known to be nonlinear, this approach improves accuracy over an assumption that stress-intensity factors remain constant for each step. It should also be noted that the result of Equation 16 over the crack front is a set $\{\Delta n_j^{(s)}\}_{j=1}^J$. The elements of this set are not guaranteed to be equal, although they are expected to be similar if an appropriate value of Δa_{median} was chosen. The final number of cycles consumed by a given step should be a scalar. To enforce this constraint, the mean value was used,

$$\Delta n^{(s)} = \frac{\sum_{j=1}^J \Delta n_j^{(s)}}{J}. \quad (17)$$

Following crack extension from step s to $s+1$, the new geometry was meshed. For efficiency, a local/global remeshing scheme was adopted where only a pre-defined subset of the model geometry was re-meshed at each step, as shown in Figure 4.

The step-by-step growth process was iterated until the crack reached the length $\hat{c}_{t=8}$, or the expected value of the last crack length diagnosis generated over the course of the experiment. This was chosen as the failure criterion to avoid violation of linear elastic assumptions as the crack approached the edge of the specimen. It was also beneficial as the cycle at which the last diagnosis occurred was known, making this a useful point for validation. Note that $\hat{c}_{t=8}$ was left out of the calibration data set for this reason.

Upon completion of the crack growth simulation, the relative prediction of cycles from the point at which the first diagnosis was made to failure was

$$n_f^R = \sum_{s=1}^S \Delta n^{(s)}, \quad (18)$$

where S is the total number of crack growth steps required to reach the failure condition. The quantity of interest for the DT feasibility study was the total number of fatigue cycles until this failure condition, or end of life (EOL). Substituting Equation 18 into a deterministic version of Equation 2 yields the deterministic EOL,

$$n_f = \sum_{s=1}^S \Delta n^{(s)} + n_D, \quad (19)$$

where n_D is the number of cycles consumed at the time of the first diagnosis.

4.3.2 Incorporating Uncertainty Into the Predictive Model

A single evaluation of the crack growth model, \mathcal{M}^P , comprising a crack insertion step followed by coupled iterations of finite element analysis and crack growth, resulted in a single, deterministic prediction of n_f^R . To provide a probabilistic prediction, uncertainty had to be quantified via calibration as described in Section 2 and given a set of observations, $\{\mathbf{c}_t\}_{t=1}^T$, where T is the current time index. The objective of the calibration was to estimate the posterior probability density function $\pi_t(\boldsymbol{\theta}|\mathbf{c}_t)$ with each new diagnosis. However, in the case where the crack location, l , was fixed, this set of observations was simply a set of mean estimates of crack length $\{\hat{a}_t\}_{t=1}^T$. In that case, the model response, \mathcal{M}_t^P , consisted of T estimates of crack length at time indices $t = 1, \dots, T$.

The number of cycles consumed prior to the zeroth step of the simulation, n_D , was fixed, but the starting crack length shown in Figure 3 was an unknown random variable, $A^{(0)}$, and had to be estimated along with the other model parameters. The parameter vector was thus $\boldsymbol{\theta} = [\beta, m, a^{(0)}]$. The prior distributions for the crack growth rate model parameters β and m were assumed to be noninformative uniform distributions as shown in Table 2. The prior distribution for the initial crack length was set equal to the first diagnosis obtained at $t = 0$, where $A^{(0)} \approx \pi(\mathbf{c}_0|e_0)$. Note that \hat{a}_0 was omitted from the observations for this reason. This coupling was not required but did serve to regularize the inverse problem, aiding in calibration when the number of available observations, T , was small relative to the dimension of Θ .

| Parameter | Prior Distribution |
|-----------|--------------------|
| a_0 | $A^{(0)}$ |
| β | $U(0, 1)$ |
| m | $U(0, 50)$ |

Table 2: Predictive model parameter prior distributions

Following calibration via MCMC, as described in Section 2, the quantified uncertainties were propagated using MCS to form a probabilistic prediction of N_f^R . Each sample in both MCMC and MC required a single deterministic simulation, leading to significant computational expense. Timing statistics were gathered for over 130 steps of the deterministic simulation. The mean time required to complete each step on a single core of a 2.4GHz Quad-Core AMD Opteron[®] processor was approximately 16.75 minutes. The median time was 15.5 minutes, and the minimum time required was 7.5 minutes. With the number of total steps for a typical simulation, S , on the order of 100, a conservative estimate of the time to complete a single simulation was 12.5 hours. Thousands to millions of samples were required for parameter estimation via MCMC and propagation via MCS, rendering the full-order analysis intractable. To combat this, a reduced-order model, or surrogate model, was developed that reduced mean simulation times by 6 orders of magnitude. The surrogate modeling approach for prognosis is briefly outlined in the next subsection.

4.3.3 Prognostic Surrogate Model

Like the full-order model, the surrogate model, $\widetilde{\mathcal{M}}^{\mathcal{P}}$, had input parameters $\boldsymbol{\theta} = [\beta, m, a^{(0)}]$ and the response was a set of predicted crack lengths corresponding to $t = 1, \dots, T$. To accomplish this, a reduced-order relationship between stress-intensity factor and a characteristic crack length, $\Delta\tilde{K}(a^*)$, was approximated by running a single deterministic crack growth simulation. A third-order polynomial was fit to the median crack front points of each step, resulting in the curve shown in Figure 4. A $K_{I,\max}$ was extracted at each point along the fitted path, resulting in the relationship of interest shown in Figure 5. Once obtained, $\Delta\tilde{K}(a^*)$ replaced $\Delta K(a_j)$ in Equation 16, yielding

$$\Delta\tilde{n}^{(s)} = \frac{1}{\beta} \int_{a^{*(s)}}^{a^{*(s+1)}} [\Delta\tilde{K}(a^*)]^{-m} da^*. \quad (20)$$

Given a realization of input parameters, $\boldsymbol{\theta}$, the surrogate model performed a step-by-step integration over the reduced-order path according to Equation 20. In this way, the parameters could be varied freely by the MCMC and MCS sampling algorithms without needing to run a new, full-order simulation. The use of surrogate modeling yielded a conservative speedup of 937,500 \times . While the surrogate model resulted in a loss of generality, it maintained accuracy for the specific application of interest. Extending the surrogate to a random starting




Figure 5: Relationship between mode I stress-intensity factor and the characteristic crack length along the reduced-order crack path.

location (e.g., one of the other geometric features likely to initiate fatigue cracks) would involve 13 additional full-order simulations and estimation of the discrete random variable L with realization $l \in [1, 14]$.

5 Results & Discussion

The pre-diagnosis regime, discussed in [15], involved the propagation of uncertainties from the microscale to the macroscale to form a probabilistic prediction of fatigue life. The uncertainty in the estimate was undesirably high, as no observations were available to inform the modeling of the primary crack growth drivers. In contrast, the post-diagnosis regime discussed herein involved systematic reduction of uncertainty in the pre-diagnosis prediction through information gained from in-situ diagnostics. Bayesian inference was used to convert strain measured on the surface of the test specimen into a stochastic estimate of crack length at seven time steps; i.e., $t = 1, \dots, 7$. The diagnoses were used to quantify uncertainty in the fatigue crack growth model through a calibration step. Each time the set of diagnoses was updated, the model was re-calibrated and used to generate a non-deterministic prediction of the number of fatigue cycles to failure. The remainder of this section presents results and discussion for (i) the crack length diagnoses, and (ii) the resulting prognosis of total fatigue life generated after each diagnosis.




Figure 6: Diagram of the effective sensor arrays used to estimate crack length using the probabilistic damage diagnosis framework. The markers indicate where strain data was sampled from the DIC data. Sensor arrays 1, 2, 3, and 4 are 7.62 mm, 12.7 mm, 17.78 mm, and 22.86 mm away from the bottom bolt hole, respectively. The sensors are evenly distributed in the horizontal direction, spaced 0.254 mm apart.

5.1 Crack Length Diagnosis

Diagnoses of crack length were obtained using the probabilistic crack characterization framework described in Section 2.2 and the strain prediction model described in Section 4.2. The goal of this analysis was twofold: (i) to provide crack length estimates as inputs to the damage prognosis method described in Section 2.3, and (ii) to study the effectiveness of the diagnostic framework as the distance between the damage and measurements was increased. Access to full field strain data from DIC meant that the latter was performed with ease by selectively sampling strain values at discrete locations throughout the specimen. Figure 6 shows the four effective strain sensor arrays that were sampled from DIC and used to estimate crack length.

The diagnostic algorithm was executed at seven different stages of the experiment for each of the four sensor arrays. For each analysis, 10,000 samples were drawn using MCMC after an initial burn-in period of 1,000 samples to estimate the posterior PDF of the crack length $\pi(a_t|e_t)$. A uniform prior distribution was used for the crack length and the noise level (variance, σ^2) was treated as unknown during the sampling. The surrogate modeling approach described in Section 4.2 was used to accelerate the analyses.

The resulting crack length estimates are displayed in Figure 7. Here, the predicted probability distributions of crack length using strain data from each of the four sensor arrays is compared to a reference value from visual inspection




Figure 7: Diagnosis results estimating the crack length probability distributions at a) $t = 1$ (234,000 cycles), b) $t = 3$ (237,000 cycles), c) $t = 5$ (239,000 cycles), and d) $t = 7$ (241,000 cycles) with each of the sensor arrays from Figure 6.

at different stages of the experiment (a) 234,000 cycles, b) 237,000 cycles, c) 239,000 cycles, and d) 241,000 cycles, or $t = 1, 3, 5,$ and 7 , respectively. As expected, strain data from Sensor Array 1 provides the most precise (i.e., the PDF is most concentrated) estimate of crack length and the uncertainty grows in the predictions as the measurements are taken farther away from the crack. Despite the impact of sensor proximity on the results, it can be seen that the diagnostics generated for all cases provide useful information about the actual crack size. The crack length observations, $\{\hat{a}_t\}_{t=1}^T$, used for the prognostics results to follow were generated using strain data from Sensor Array 2.

5.2 Fatigue Life Prognosis

To calibrate the crack growth surrogate model, 20,000 samples were drawn using MCMC with a burn-in period of 10,000 samples. The Delayed Rejection Metropolis (DRAM) algorithm was used with an adaptation interval of 1,000 samples [52]. The noise variance, σ^2 , was estimated along with the other model parameters. After each calibration, the uncertainty was propagated via MCS using all 10,000 samples remaining after burn-in. The model evaluation times for the MCMC-driven surrogate model and the MCS-driven surrogate model varied

slightly. The former only had to generate estimates up to the crack length a_T while the latter had to simulate crack growth up to failure. In general, this can yield additional speedup for an MCMC-driven model if the number of simulation steps required to reach failure from the most recent diagnosis is large.

The diagnoses obtained over time are shown in the left column of Figure 8, while the corresponding predictions are shown on the right. The first diagnosis was used to form the prior distribution $A^{(0)}$, and the first prediction was made at $T = 3$, meaning the first set of available diagnoses was $\{\hat{a}_t\}_{t=1}^3$. Overall, in-situ diagnostics led to a dramatic decrease in variance when compared with the pre-diagnosis regime prediction of [15]. The precision generally improved with additional diagnoses, as expected. The percent error of the mean predictions at times $t = 4$, $t = 5$, $t = 6$, and $t = 7$ compared to the true cycles to failure observed during the test were 0.08%, 0.11%, 0.01%, 0.02%, respectively.

While MCS led to a probabilistic prediction of the absolute fatigue life, N_f , a more useful measure of time to failure might be the remaining useful life (RUL). RUL is a linear function of time,

$$RUL = N_f - n_T, \quad (21)$$

where n_T is the cycle count at the time the prognosis was made. In a real-world scenario, RUL provides a measure of the time remaining for maintainers, operators, or mission planners to take action prior to failure. A plot of RUL vs. cycles is shown in Figure 9. This plot shows true RUL, based on the actual fatigue life observed at the conclusion of the test, $n_{f,\text{true}}$, plotted as a straight black line. RUL predictions and error bounds can be drawn about this true value for qualitative assessment of the prognostic algorithm's performance as $t \rightarrow t_{\text{failure}}$.

The sequential RUL predictions are represented by box and whisker plots. Here, the predicted mean and median are denoted by a small square symbol and a line dividing the box, respectively. The upper and lower quartiles of the data are represented by the extent of the box, and the whiskers, or the capped dotted lines, denote the range of the data. A single set of error bounds based on the prognostic metrics proposed by [53] are plotted for reference. These bounds are a percent error metric referred to as the α -bounds, where α^+ and α^- represent the upper and lower error bounds, respectively, and $\alpha \in [0, 1]$. The α -bounded region for the time-varying RUL decreases as $t \rightarrow t_{\text{failure}}$ and is defined as

$$\alpha^\pm = RUL \times (1 \pm \alpha). \quad (22)$$

These bounds allow for easy visualization of prognostics performance and are drawn in Figure 9 as a grey, triangular region.

The percent error in the mean predictions of RUL at times $t = 4$, $t = 5$, $t = 6$, and $t = 7$ were 5.19%, 8.51%, 0.85%, and 5.93%, respectively. The accuracy of the predicted mean fluctuated but remained within the 10% bounds. While the first prediction in Figure 9 was fairly accurate, the lack of precision in the prediction indicates a possibility of immediate failure. It was not until

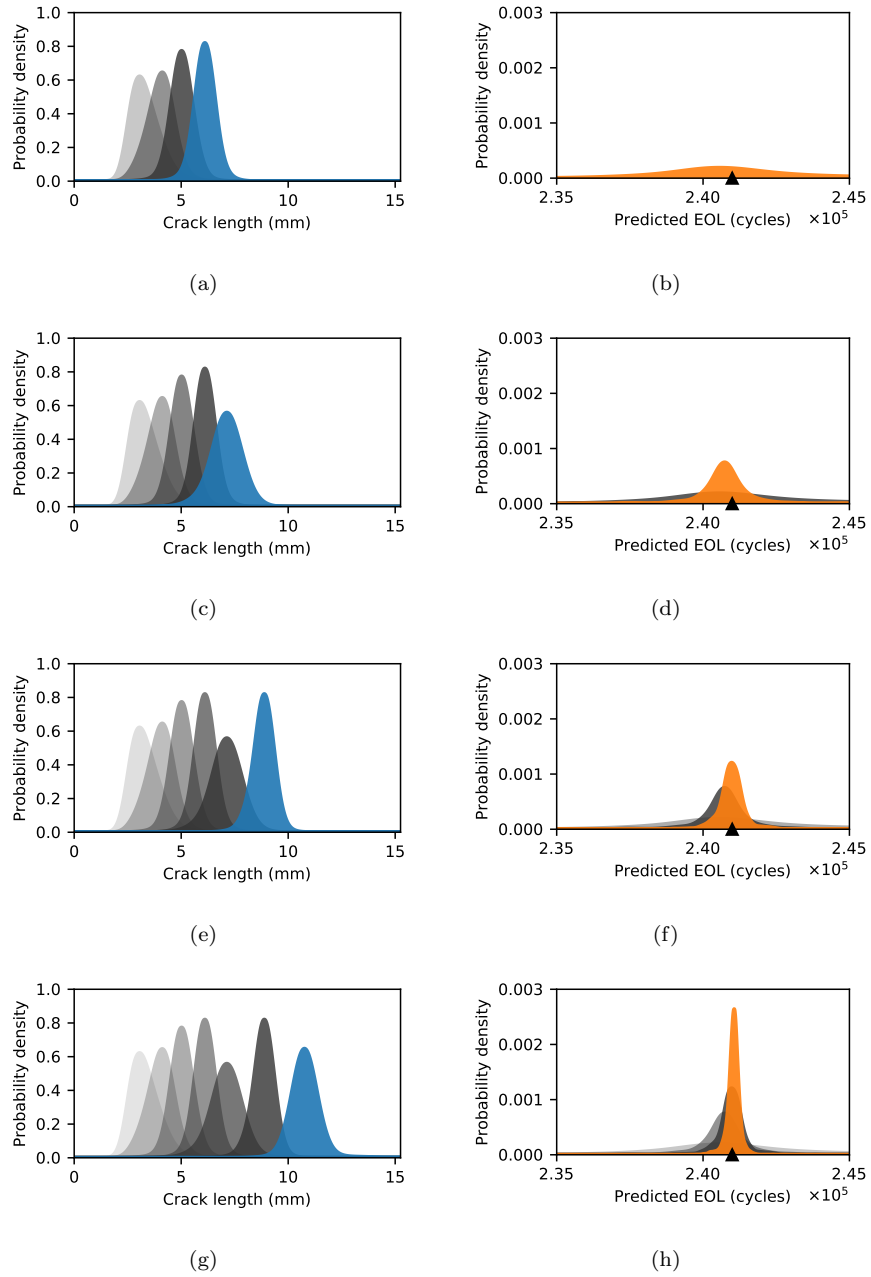


Figure 8: Available diagnoses and the corresponding prognosis over time. The left column displays the diagnoses up to times $t = 4$, $t = 5$, $t = 6$, and $t = 7$, respectively, while the right column shows the corresponding prediction made with available diagnoses at those times.




Figure 9: Remaining useful life (RUL) prognoses. True RUL is represented by the solid black line. The grey cone is the 10% α -bounded region. The red box represents the mean prediction made at the number of elapsed cycles, and the blue box represents the upper and lower quartiles of the RUL samples, and the dotted lines represent the range of the samples.

the fourth diagnosis that a prognostic horizon² was established, approximately equal to 3,000 fatigue cycles.

In practice, a lower limit will exist for the prognostic horizon below which there is not enough time left to take corrective or preventative action. In this study, the fatigue crack growth was being driven by relatively high loads causing the crack to grow fast enough that only seven measurements were taken. This experiment was just for demonstration but serves as a good example of the need to define monitoring intervals according to the lower limit of detection capabilities and the expected rate of damage progression once that limit has been surpassed.

6 Conclusion

A general method for coupling in-situ diagnostics and prognostics in a non-deterministic, Digital Twin (DT) framework was presented. The framework was

²The prognostic horizon is defined mathematically in [53]. In short, it defines the number of cycles prior to failure at which a pre-defined percentage of prediction probability lies within the true RUL α -bounds. Note that this metric can only be used postmortem to evaluate performance.

demonstrated using a geometrically complex, metallic test specimen exhibiting life-limiting fatigue crack growth. The complexity of the test specimen studied was geometrically-complex, necessitating the use of high-fidelity, finite element analysis to model material response due to the presence of a fatigue crack and to simulate fatigue crack growth. Using Markov chain Monte Carlo sampling, a limited number of measurements of strain on the surface of the specimen were used to generate a temporal series of probabilistic crack state estimates. These observations were used to calibrate the prognostic model with quantified uncertainty. These uncertainties were propagated via Monte Carlo simulation to failure, yielding probabilistic predictions of fatigue life.

The results of this study demonstrate (i) that the coupling of the proposed diagnostic and prognostic methods enabled successful realization of DT for predicting fatigue life in a laboratory environment and (ii) that the use of damage diagnostics can systematically reduce uncertainty in fatigue life predictions over time. Continuation of the work should focus on increasing application realism, which will serve to test the generality of the DT approach and guide its development. Along with the first part of this two-part paper series, the DT feasibility study laid the foundation for an initiation-to-failure, microscale-to-macroscale, non-deterministic DT methodology for life prediction of fatigue critical structures.

References

- [1] M. Grieves, J. Vickers, Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems, in: *Transdisciplinary Perspectives on Complex Systems*, Springer, 2017, pp. 85–113.
- [2] E. Glaessgen, D. Stargel, The digital twin paradigm for future NASA and US Air Force vehicles, in: *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2012, pp. 1–14.
- [3] Defense Acquisition University, Digital Twin, Defense Acquisition Glossary, <https://www.dau.mil/glossary/pages/3386.aspx>, accessed: 2017-10-22 (2017).
- [4] J. Emery, J. Hochhalter, P. Wawrzynek, G. Heber, A. Ingrassia, DDSim: A hierarchical, probabilistic, multiscale damage and durability simulation system—part I: methodology and level I, *Engineering Fracture Mechanics* 76 (10) (2009) 1500–1530.
- [5] E. J. Tuegel, A. R. Ingrassia, T. G. Eason, S. M. Spottswood, Reengineering aircraft structural life prediction using a digital twin, *International Journal of Aerospace Engineering* (2011).
- [6] B. Piascik, J. Vickers, D. Lowry, S. Scotti, J. Stewart, A. Calomino, NASA TA12: Materials, structures, mechanical systems, and manufactur-

- ing roadmap, NASA Space Technology Roadmaps and Priorities (2012) 1–34.
- [7] E. H. Glaessgen, B. Biegel, F. Chandler, D. Chrichton, J. LeMoigne, M. Little, C. Null, W. Peters, J. B. Ransom, L. Wang, NASA TA11: Modeling, Simulation, Information Technology, and Processing, NASA Space Technology Roadmaps (2015) 1–170.
 - [8] K. Reifsnider, P. Majumdar, Multiphysics stimulated simulation digital twin methods for fleet management, in: 54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2013, p. 1578.
 - [9] A. Gilchrist, Introduction to the Industrial Internet, in: Industry 4.0, Springer, 2016, pp. 1–12.
 - [10] J. D. Hochhalter, W. P. Leser, J. A. Newman, V. K. Gupta, V. Yamakov, S. R. Cornell, S. A. Willard, G. Heber, Coupling damage-sensing particles to the digital twin concept, NASA/TM-2014-218257 (2014).
 - [11] P. E. Leser, J. D. Hochhalter, J. E. Warner, J. A. Newman, W. P. Leser, P. A. Wawrzynek, F.-G. Yuan, IWSHM 2015: Probabilistic fatigue damage prognosis using surrogate models trained via three-dimensional finite element analysis, Structural Health Monitoring (2016) 291–308.
 - [12] P. E. Leser, J. A. Newman, J. E. Warner, W. P. Leser, J. D. Hochhalter, F.-G. Yuan, Probabilistic prognosis of non-planar fatigue crack growth, in: M. J. Daigle, A. Bregon (Eds.), Proceedings of the Annual Conference of the Prognostics and Health Management Society, Denver, Colorado, 3-6 October 2016, Prognostics and Health Management Society, New York, 2016.
 - [13] J. E. Warner, G. F. Bomarito, J. D. Hochhalter, W. P. Leser, P. E. Leser, J. A. Newman, A computationally-efficiency probabilistic approach to model-based damage diagnosis, International Journal of Prognostics and Health Management 8 (2) (2017) 1–17.
 - [14] S. R. Yeratapally, M. G. Glavicic, M. Hardy, M. D. Sangid, Microstructure based fatigue life prediction framework for polycrystalline nickel-base superalloys with emphasis on the role played by twin boundaries in crack initiation, Acta Materialia 107 (C) (2016).
 - [15] S. R. Yeratapally, P. E. Leser, J. D. Hochhalter, W. P. Leser, T. J. Ruggles, Digital Twin Feasibility Study (Part I): Non-Deterministic Predictions of Fatigue Life in Aluminum Alloy 7075-T651 Using a Microstructure-based Multi-scale Model, Engineering Fracture Mechanics (this volume) (2019).
 - [16] S. J. Engel, B. J. Gilmartin, K. Bongort, A. Hess, Prognostics, the real issues involved with predicting life remaining, in: S. Sepahban (Ed.),

Aerospace Conference Proceedings, Vol. 6, Big Sky, Montana, 18-25 March 2000, IEEE, 2000, pp. 457–469.

- [17] A. Ray, S. Tangirala, Stochastic modeling of fatigue crack dynamics for on-line failure prognostics, *IEEE Transactions on Control Systems Technology* 4 (4) (1996) 443–451.
- [18] A. Makeev, Y. Nikishkov, E. Armanios, A concept for quantifying equivalent initial flaw size distribution in fracture mechanics based life prediction models, *International Journal of Fatigue* 29 (1) (2007) 141–145.
- [19] R. Cross, A. Makeev, E. Armanios, Simultaneous uncertainty quantification of fracture mechanics based life prediction model parameters, *International journal of fatigue* 29 (8) (2007) 1510–1515.
- [20] F. Perrin, B. Sudret, M. Pendola, Bayesian updating of mechanical models-application in fracture mechanics, 18ème Congrès Français de Mécanique (2007).
- [21] M. E. Orchard, G. J. Vachtsevanos, A particle-filtering approach for on-line fault diagnosis and failure prognosis, *Transactions of the Institute of Measurement and Control* 31 (3-4) (2009) 221–246.
- [22] F. Cadini, E. Zio, D. Avram, Monte Carlo-based filtering for fatigue crack growth estimation, *Probabilistic Engineering Mechanics* 24 (3) (2009) 367–373.
- [23] E. Zio, G. Pelsoni, Particle filtering prognostic estimation of the remaining useful life of nonlinear components, *Reliability Engineering & System Safety* 96 (3) (2011) 403–409.
- [24] J. Chiachio, M. Chiachio, A. Saxena, G. Rus, K. Goebel, An energy-based prognostics framework to predict fatigue damage evolution in composites, in: S. Sankararaman, I. Roychoudhury (Eds.), *Proceedings of the Annual Conference of the Prognostics and Health Management Society, New Orleans, Louisiana, 14-17 October 2013*, Prognostics and Health Management Society, New York, 2013, pp. 363–371.
- [25] X. Guan, R. Jha, Y. Liu, Model selection, updating, and averaging for probabilistic fatigue damage prognosis, *Structural Safety* 33 (3) (2011) 242–249.
- [26] X. Guan, R. Jha, Y. Liu, Probabilistic fatigue damage prognosis using maximum entropy approach, *Journal of intelligent manufacturing* 23 (2) (2012) 163–171.
- [27] T. Peng, J. He, Y. Xiang, Y. Liu, A. Saxena, J. Celaya, K. Goebel, Probabilistic fatigue damage prognosis of lap joint using Bayesian updating, *Journal of Intelligent Material Systems and Structures* 26 (8) (2015) 965–979.

- [28] T. Peng, Y. Liu, A. Saxena, K. Goebel, In-situ fatigue life prognosis for composite laminates based on stiffness degradation, *Composite Structures* 132 (2015) 155–165.
- [29] R. C. McClung, M. P. Enright, H. R. Millwater, G. R. Leverant, S. Hudak, A software framework for probabilistic fatigue life assessment of gas turbine engine rotors, *Journal of ASTM International* 1 (8) (2004) 1–16.
- [30] M. P. Enright, S. J. Hudak, R. C. McClung, H. R. Millwater, Application of probabilistic fracture mechanics to prognosis of aircraft engine components, *AIAA Journal* 44 (2) (2006) 311–316.
- [31] R. C. McClung, Integrating fracture mechanics into the material and structural design process, *Materials Performance and Characterization* 5 (3) (2016) 171–193.
- [32] W. A. Grell, P. J. Laz, Probabilistic fatigue life prediction using afgrow and accounting for material variability, *International Journal of Fatigue* 32 (7) (2010) 1042–1049.
- [33] Y. Bazilevs, X. Deng, A. Korobenko, F. L. di Scalea, M. Todd, S. Taylor, Isogeometric fatigue damage prediction in large-scale composite structures driven by dynamic sensor data, *Journal of Applied Mechanics* 82 (9) (2015) 091008.
- [34] R. C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*, SIAM, 2014.
- [35] J. Geweke, Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, *Bayesian Statistics* 4 (1992) 169–193.
- [36] M. K. Cowles, B. P. Carlin, Markov chain Monte Carlo convergence diagnostics: a comparative review, *Journal of the American Statistical Association* 91 (434) (1996) 883–904.
- [37] J. Kaipio, E. Somersalo, *Statistical and computational inverse problems*, Springer Science & Business Media, 2006.
- [38] S. Theodoridis, *Machine Learning: a Bayesian and Optimization Perspective*, Academic Press, 2015.
- [39] R. E. Bartels, R. C. Scott, C. Funk, T. J. Allen, B. W. Sexton, Computed and experimental flutter/LCO onset for the Boeing truss-braced wing wind-tunnel model, in: *Proceedings of the AIAA Aviation Conference, Atlanta, Georgia, 16-20 June 2014, AIAA, 2014*, pp. 16–20.
- [40] O. Gur, M. Bhatia, J. A. Schetz, W. H. Mason, R. K. Kapania, D. N. Mavris, Design optimization of a truss-braced-wing transonic transport aircraft, *Journal of Aircraft* 47 (6) (2010) 1907–1917.

- [41] T. Chu, W. Ranson, M. A. Sutton, Applications of digital-image-correlation techniques to experimental mechanics, *Experimental Mechanics* 25 (3) (1985) 232–244.
- [42] H. Bruck, S. McNeill, M. A. Sutton, W. H. Peters, Digital image correlation using newton-raphson method of partial differential correction, *Experimental Mechanics* 29 (3) (1989) 261–267.
- [43] B. Pan, K. Qian, H. Xie, A. Asundi, Two-dimensional digital image correlation for in-plane displacement and strain measurement: a review, *Measurement Science and Technology* 20 (6) (2009) 062001.
- [44] G. F. Bomarito, J. D. Hochhalter, T. J. Ruggles, A. H. Cannon, Increasing accuracy and precision of digital image correlation through pattern optimization, *Optics and Lasers in Engineering* 91 (2017) 73–85.
- [45] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities, *International Journal for Numerical Methods in Engineering* 79 (11) (2009) 1309–1331.
- [46] Fracture Analysis Consultants Inc, FRANC3D Reference Manual, Version 7 (2016).
- [47] American Society for Metals, Metals Handbook, Vol. 2, Properties and Selection: Nonferrous Alloys and Pure Metals, 9th Edition, 1979.
- [48] J. E. Warner, G. F. Bomarito, G. Heber, J. D. Hochhalter, Scalable Implementation of Finite Elements by NASA - Implicit (ScIFEi), Tech. rep. (2016).
- [49] Dassault Systemes, Abaqus/CAE User’s Guide, Version 6.13 (2016).
- [50] J. Yau, S. Wang, H. Corten, A mixed-mode crack analysis of isotropic solids using conservation laws of elasticity, *Journal of Applied Mechanics* 47 (2) (1980) 335–341.
- [51] F. Erdogan, G. Sih, On the crack extension in plates under plane loading and transverse shear, *Journal of Basic Engineering* 85 (4) (1963) 519–525.
- [52] H. Haario, M. Laine, A. Mira, E. Saksman, DRAM: efficient adaptive MCMC, *Statistics and Computing* 16 (4) (2006) 339–354.
- [53] A. Saxena, J. Celaya, B. Saha, S. Saha, K. Goebel, Metrics for offline evaluation of prognostic performance, *International Journal of Prognostics and Health Management* 1 (1) (2010) 4–23.