Title

A probabilistic estimation approach for the failure forecast method using Bayesian inference

Authors

O'Dowd, Niall M
Madarshahian, Ramin
Leung, Michael Siu Hey
et al.

Publication Date

2021

DOI

10.1016/j.ijfatigue.2020.105943

Peer reviewed

A PROBABILISTIC ESTIMATION APPROACH FOR THE FAILURE FORECAST METHOD USING BAYESIAN INFERENCE

Niall M. O'Dowd[1], Ramin Madarshahian[1], Michael Siu Hey Leung[2], Joseph Corcoran[2], and Michael D. Todd[1,*]

[1] Department of Structural Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla CA 92093-0085 USA; *email: mdtodd@eng.ucsd.edu

[2] Department of Mechanical Engineering, Imperial College London, UK

Key words: failure forecast method; Bayesian inference; fatigue life; remaining life estimation; structural health monitoring

Abstract. Positive-feedback mechanisms such as fatigue induce a self-accelerating behavior, captured by models displaying infinite limit-state asymptotics, collectively known as the failure forecast method (FFM). This paper presents a Bayesian model parameter estimation approach to the fully nonlinear FFM implementation and compares the results to the classic linear regression formulation, including a regression uncertainty model. This process is demonstrated in a cyclic loading fatigue crack propagation application, both on a synthetic data set and on a full fatigue experiment. A novel "switch point" parameter is included in the Bayesian formulation to account for nonstationary changes in the growth parameter.

1 Introduction

Information provided by structural health monitoring (SHM) data is generally used to assess the current diagnostic state of components or systems [1]. Such SHM assessments may be subsequently used to inform predictive models that estimate remaining useful life (RUL) or some related prognostic condition [2,3]. RUL is most commonly defined as the amount of time, subject to some assumed usage profile, that a structure has before achieving some limit state that prevents its ability to perform its intended design functions safely and reliably [4]. The achievement of this limit state (measured as time, response, or load cycles, or a similar metric) will be defined as the time-of-failure, denoted by the variable t_f.

In particular, this paper focuses on estimating t_f of specimens subjected to ultimate failure induced by fatigue cracking. A common method employed for predicting t_f in a fatigue scenario is the Paris-Ergodan crack growth law [5]. This law expresses the rate of crack growth as a function of parameters such as the stress intensity factor from the load around the crack tip, material properties which may be obtained experimentally, and stress information particular to the specific cyclic loading. Although motivated by linear elastic fracture mechanics, the parameters which affect t_f predictions using Paris' law have intrinsic uncertainty, and in fact they may be impossible to estimate or measure accurately in cases such as complex geometries, materials, or load states. Varying environmental effects may also infuse substantive uncertainty into the Paris' law approach.

Such uncertainties in Paris' law (and many other classes of such crack growth models) necessitated development of more probabilistic methods for estimating RUL in fatigue loading applications. Some of the sources of variability or uncertainty that affect prediction of t_f, generally regardless of modeling

approach, include (but aren't limited to) stochastic environmental or loading effects, uncertainty in initial material state, and measurement process noise. This inherently makes the prediction problem probabilistic, since typically these influences cannot be measured or estimated completely [6, 7]. One subsequent approach addressed this with a modified version of the Paris-Ergodan crack growth law known as the NASGRO equation, which quantified the material uncertainties and loading conditions in fatigue behavior using Monte Carlo (MC) analysis [8, 9]. Mallor et al. recently presented a study on the probabilistic formulation for fatigue crack propagation based on the NASGRO equation and provided a stochastic approach for predicting statistical moments of fatigue lifetime for components subject to non-uniform loading patterns [10]. In this study, the RUL expected value and variance of numerically simulated fatigue loading experiments are approximated and verified using MC simulation. Artificial neural networks are also being used to estimate RUL, without the need for an explicit physical fatigue mathematical model form. Studies by Jimenez-Martinez et al. and Barbosa et al. show the results of using machine learning capabilities with only two inputs (component strain and fatigue cycle) to estimate failure [11, 12]. Both studies reported a higher prediction capability at some load sequences in comparison to traditional modeling techniques.

For failure modes expected to possess self-accelerating observable behavior (such as unarrested crack growth), a method for estimating t_f that has been shown to have broad application is widely-known as the failure forecast method (FFM). Early origins of the FFM began with Fukuzono, who observed that landslide mechanics could be explained by using an empirical positive-feedback model of the inverse rate of ground strain [13]. Voight later formalized this observation into a more comprehensive predictive model and first coined the term "failure forecast method" [14, 15]. The FFM has been further widely implemented in modeling a diverse variety of self-accelerating failure mechanisms, ranging from material-level failures (e.g., creep, fatigue) to volcanic eruptions [16–22]. This paper will focus on the use of the FFM to sample the distribution of t_f for fracture in fatigue loading applications. For fatigue cracking (and many other applications, for that matter), the FFM has traditionally been implemented in a linearized way by linearly regressing the inverse rate of crack growth against time, as detailed later in Section 2.1. When the inverse rate of the monitored feature approaches zero (the time axis intercept), the feature rate approaches infinity, and the failure event is defined to have occurred [16–18]. However, the regression implementation of the FFM provides a single estimated point for t_f, which does not consider explicitly the propagation of uncertainty in the estimation process. A recent study proposed a model to approximate this propagated uncertainty in the FFM linear regression process [23]; this same model will be used in this study to estimate probability density functions for t_f for the "classical" linear FFM approach, in comparison to the Bayesian model that will be proposed and implemented in this work. More recent work has shown that the linearity assumption made in the regression implementation of FFM has been observed to be false for both early and late stage crack growth [24]. This work shows that imposing linearity may cause a positive bias in t_f estimation for late stage crack growth, where the inverse feature rate has been observed to slope downward; this implies that the failure time occurs consistently sooner than predicted, introducing a non-conservative bias that could be catastrophic. While this study focuses on the problem of fatigue, a range of different damage mechanisms exhibit multiple distinct phases of progression determined by the underlying mechanics, and indeed the inclusion of pre-damage data will lead to non-constant trends, and so a solution to this problem will be widely applicable.

The Bayesian model introduced in this work allows both for relaxation of the linearity assumption and for sampling the posterior distribution of t_f. We shall compare t_f distribution models from the fully-Bayesian approach to t_f estimations made from the linear regression approach, including distributions from the linear regression uncertainty model. Approaching model parameter estimation from a Bayesian statistical perspective is advantageous when only a few realizations (or even a single realization) can be obtained for analysis, such as data from a single cyclic fatigue loading experiment. The operation of the Bayesian modeler consists of formulating and continuously sampling the distributions of model parameters as new data become available, intrinsically allowing for the generation of uncertainty distributions for each parameter, including t_f. We thus relax the linearity constraint by including a model parameter in the Bayesian-inferred parameter set which accounts for nonlinearity, rather than setting it to an assumed value.

Recent published works have begun to propose various uncertainty quantification approaches to the FFM. Commonly, published works operate on data from historic earthquakes, but the FFM is agnostic to the particular failure mechanism which makes these studies related to the current work. Bell et al. [25] evaluated the FFM using synthetic strain and earthquake sequences. A small amount of Gaussian noise was added to the strain rate data to simulate the effect of measurement error, and earthquake data was simulated as a Poisson process with mean rate and variance. The article synthesized 200 realizations of data for each process, and generated probability density functions based on t_f estimations of the FFM for each realization. Bevilacqua et al. [26] proposed a doubly stochastic enhancement of the FFM by introducing a formulation similar to the Hull-White model in financial mathematics. By including stochastic noise terms in the original FFM governing equation, t_f estimation uncertainty can be systematically characterized. Their method provides complete posterior probability distributions, allowing for worst case scenario estimation with a specified level of confidence. Bevilacqua's method is applied to earthquake eruption prediction, with small sets of data compared to the data presented in this study. In the domain of fatigue prediction, Leung et al. [24] compared the performance between conventional periodic inspections informing Paris' law parameters (themselves updated via Bayesian inference) and continuous monitoring for use by the linearized formulation of the FFM. In that study, random uncertainty in the damage accumulation rate measurements is the only source of uncertainty for the FFM, which in turn results in uncertainties in the regression fit and the extrapolated time axis intercept (t_f). The data are assumed to fit the model form, and it is assumed that data realizations are readily obtained.

While uncertainty modeling within the FFM has indeed progressed, the use of Bayesian inference is quite limited. Boue et al. in [20] and [27] implemented a Bayesian model to estimate the eruption time of volcanoes, specifically trained on data from Volcán de Colima. Their Bayesian model operated on the empirical power law which contained parameters including the failure time coinciding with eruption onset. Their Bayesian model generated confidence levels associated with the failure time, which is especially useful for crisis management and decision making. They showed that Bayesian model sampling methods were beneficial to exploring each measured feature, and using a linear regression for the FFM was not always relevant for every measured feature set. The present work provides a significantly broader look into the use of Bayesian reasoning in the FFM, achieved by evaluating the performance of the Bayesian model on synthetic data with varying parameters, allowing for many realizations to be generated. We also provide advancement by developing a parameter switching model

which determines a discrete point where our Bayesian parameters can switch (detailed in section 4).

Within the FFM framework, this paper compares two probabilistic models for estimating t_f: the linear regression with an uncertainty model, and the proposed fully-Bayesian approach. The probability density functions (PDFs) of the linear regression model $p(\hat{t_f})$ and the t_f posterior belief distributions of the Bayesian model are compared on a synthesized dataset (in section 3), and a real fatigue dataset (in section 4).

2 Predictive Models

The most common general FFM model form proposes that the time rate-of-change R of some time-dependent measured feature Ω, $R = \dot{\Omega}$, obeys the following evolution equation

$$\dot{R} = kR^{\alpha}, \tag{1}$$

where $k > 0$ and $\alpha > 1$ are empirical constants relevant to the specific physical process. As these constants are not derived from a mechanics consideration, one attractive property of the FFM model formulation is that, unlike Paris' law, no knowledge of application-specific loading parameters and material properties is required. The solution to Eq. (1), assuming that the rate R at the time of failure t_f is R_f (often assumed infinite in the original literature, since the model form admits an infinite asymptote at t_f) is given by

$$R(t) = [R_f^{1-\alpha} + k(\alpha - 1)(t_f - t)]^{\frac{1}{1-\alpha}}. \tag{2}$$

The FFM literature have shown it to be more convenient to consider the *inverse* of the rate R, since this facilitates easier definition of the failure criterion, i.e., the inverse rate tending to zero rather than the rate itself tending to infinity at the time of failure. Defining the inverse rate $P = R^{-1}$, the solution Eq. (2) becomes

$$P(t) = [P_f^{\alpha-1} + k(\alpha - 1)(t_f - t)]^{\frac{1}{\alpha-1}}, \tag{3}$$

where t_f may be solved for as

$$t_f = t + \frac{P^{\alpha-1} - P_f^{\alpha-1}}{k(\alpha - 1)}. \tag{4}$$

This form of the FFM predicts that failure will occur at some inverse feature rate-informed amount of time after present time t, encapsulated in the second term on the right-hand side of Equation 4.

2.1 The Linearized FFM with Linear Regression Method

The linearized FFM approach is to assume that α is approximately 2 (based on a number of empirical studies done over a variety of failure mechanisms, as evidenced in the literature already cited), which permits a linear regression on Eq. (4), after rearrangement, as

$$P - P_f = kt_f - kt. \tag{5}$$

Clearly, as $t \to t_f$, then $P \to P_f$, where the mathematically idealized target failure criterion is that $P_f = 0$, corresponding to $R_f \to \infty$, as alluded to above. By setting it to any positive value, a degree of conservatism is introduced into the approach. Of course, in most practical applications, "failure" occurs at a point prior to an infinite data rate-of-change observation, but to be consistent with general implementation in the literature and for the purposes of parametric studies in this paper, we will employ $P_f = 0$ as the failure criterion, which won't change the basic nature of this study. With that, the two regression coefficients obtained from a time/data linear regression (over some given window of time) are given by $\beta_0 = kt_f$ (intercept) and $\beta_1 = -k$ (slope) such that the regression-estimated time to failure is $\hat{t}_f = -\beta_0/\beta_1$, the negative of the ratio of the intercept to the slope. If k is not known, it could be estimated via a maximum likelihood technique or via Markov Chain Monte Carlo (MCMC) sampling methods [22, 28].

We will review the uncertainty model recently developed for this linear regression process presented in detail in [23, 24]. Any given regression on some time-stamped feature set represents a single "block" observation, which is presumed representative of an ensemble population of regressions over the same time frame. Thus, the regression coefficient vector $\boldsymbol{\beta} = [\hat{\beta}_0, \hat{\beta}_1]^T$ is an estimate from populations of regression coefficients. A visualization of this process in seen in Figure (1).

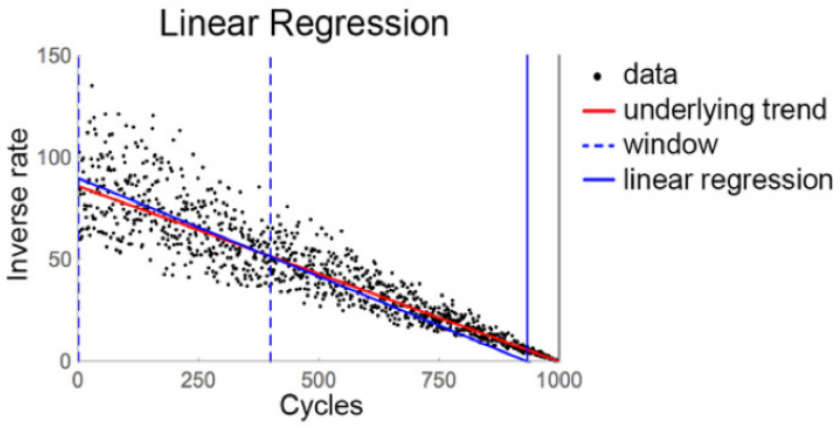

Figure 1: An example of a regression block occurring at 300 cycles. The vertical solid line represents the failure point.

For a given linear regression model $\boldsymbol{P} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{e}$, where $\boldsymbol{P}$ are observations of data, $\boldsymbol{X}$ is the design matrix, and $\boldsymbol{e}$ are the residuals between the data the linear regression model. It is assumed that the regression process yields residuals that are unbiased, uncorrelated, and normally-distributed $\boldsymbol{e} \equiv N(\boldsymbol{0}, \sigma^2\boldsymbol{I})$, under typical central limit theorem assumptions (regardless of the distribution of the regressed data, $\boldsymbol{P}$). Under this assumption, the theory follows that the regression coefficients themselves have normal distributions $\hat{\beta}_j \equiv N(\beta_j, \sigma^2(\boldsymbol{X}^T\boldsymbol{X})^{-1})$, $j = 0, 1$, and the superscript "T" indicates the matrix transpose. An unbiased estimate of the population error variance is $\sigma^2 = \|\boldsymbol{P} - \boldsymbol{X}\hat{\boldsymbol{\beta}}\|^2/(n-2)$, where n is the number of data points used in the regression design, reduced by two since two regression coefficients were estimated in the process. Given the uncertainty models in the regression coefficients, the time-of-failure is predicted by $\hat{t}_f = -\hat{\beta}_0/\hat{\beta}_1$, as stated previously. The distribution of this ratio of

correlated normal variables may be computed as

$$p(\hat{t}_f) = \int_{-\infty}^{\infty} \frac{p(-\hat{t}_f\hat{\beta}_1, \hat{\beta}_1)}{|-1/\hat{\beta}_1|} d\hat{\beta}_1 \tag{6}$$

where $p(\hat{\beta}_0, \hat{\beta}_1)$ is the bivariate normal density function. The computation of Eq. (7) yields the result

$$p(\hat{t}_f) = \frac{\sqrt{1-\rho^2}\sigma_0\sigma_1 e^{\frac{-\mu_1^2\sigma_0^2+2\rho\mu_0\mu_1\sigma_0\sigma_1-\mu_0^2\sigma_1^2}{2\sigma_0^2\sigma_1^2(1-\rho^2)}}}{\pi(\sigma_1^2+2\rho\sigma_0\sigma_1\hat{t}_f+\sigma_0^2\hat{t}_f^2)} + \frac{e^{-\frac{(\mu_1+\mu_0\hat{t}_f)^2}{2(\sigma_1^2+2\rho\sigma_0\sigma_1\hat{t_f}+\sigma_0^2\hat{t}_f^2)}}\text{erf}(\frac{\mu_1\sigma_0(\rho\sigma_1+\sigma_0\hat{t}_f)-\mu_0\sigma_1(\sigma_1+\rho\sigma_0\hat{t}_f)}{\sqrt{2-2\rho^2}\sigma_0\sigma_1\sqrt{\sigma_1^2+2\rho\sigma_0\sigma_1\hat{t}_f+\sigma_0^2\hat{t}_f^2}})}{\sqrt{2\pi}(\sigma_1^2+2\rho\sigma_0\sigma_1\hat{t}_f+\sigma_0^2\hat{t}_f^2)^{3/2}}, \tag{7}$$

where erf(*) is the error function, $\mu_j = \hat{\beta}_j$, $\sigma_j = \sqrt{\|\boldsymbol{P}-\boldsymbol{X}\hat{\boldsymbol{\beta}}\|^2(\boldsymbol{X}^T\boldsymbol{X})^{-1}_{j+1,j+1}/(n-2)}$, and $\rho = \|\boldsymbol{P}-\boldsymbol{X}\hat{\boldsymbol{\beta}}\|^2(\boldsymbol{X}^T\boldsymbol{X})^{-1}_{12}/(\sigma_0\sigma_1(n-2))$, with $j = 0, 1$; the double subscript 1,2 refers to the row-column selection of the subscripted matrix. Strictly speaking, since the exact population μ_j and σ_j are not known *a priori* and must be estimated from the data as presented above, a sampling distribution for the ratio mean and standard deviation should be derived, but here the population ratio distribution, Eq. (7) will be used as a surrogate. It should be noted that this PDF has no analytically-calculable order statistics [29]. However, Refs. [23] and [24] used Eq. (7) and showed it was sufficient to describe the distribution of the failure time, and it will be used here.

2.2 Bayesian Model

At its core, the Bayesian technique operates by continuously updating posterior beliefs (distributions) of parameters as new data become available. The general form of Bayes' equation is

$$p(\Theta|D) = \frac{P(D|\Theta)\times P(\Theta)}{P(D)}, \tag{8}$$

where D represents the data (in this case we are using inverse crack growth rate data, or P) and Θ represents the model parameters to be estimated (akin to the regression coefficients of the previous section). After obtaining new data, denoted by $D = \{d_1, ..., d_n\}$, model parameter beliefs are updated, $P(\Theta|D)$, influenced by the likelihood $P(D|\Theta)$ and a prior distribution, $P(\Theta)$. This prior distribution describes the modeler's degree of belief about the parameter values before observing any data, which may also be based on past experience, if such experience exists. One can also use uninformed prior distributions, e.g., a Jeffry's prior, if no such prior belief or information exists.

In the current work, the parameter vector Θ consists of the parameters k, α, t_f, and the standard deviation σ_l (noise parameter) that will be a part of our likelihood function (discussed more later). Because t_f should be positive and greater than the current time of data collection t, the prior is assumed to be uniformly distributed from current time, t, to ∞. The prior for k is also assumed to be an improper (it doesn't integrate to unity) uniform distribution, from 0 to ∞ because it is known to be positive. We chose the prior for α to be uniform between 1.5-2.5 because the literature shows that this parameter is close to 2 (as per the discussion above leading to the linear regression implementation)

and fluctuates in accordance to physical properties of failure mechanisms [22]. We chose the prior for σ_l to be half normal, a weakly informative distribution as suggested by Gelman et al. [30,31]. Table (1) shows a summary of the priors selected for the parameters.

Table 1: Parameters of the Bayesian model (synthetic data)

Parameter	**Prior**
t_f	Improper uniform distribution with lower bound of t
k	Improper uniform distribution with lower bound of 0
α	Uniform distribution between 1.5 and 2.5
σ_l	Half normal distribution, standard deviation of 1

If a specific forward measurement model were proposed or developed, a tailored likelihood function could be derived from that. The current work, however, seeks to maintain a Bayesian model that is more agnostic to the specific kind of data/features D that are used, so a normal model for the likelihood function $N(\mu_l, \sigma_l)$ was chosen. In order to capture the heteroscedastic converging effect of the noise structure observed in real fatigue experiments [22, 24], inference is performed in the logarithmic space, i.e., $P(D'|\Theta) = N(\mu_l, \sigma_l)$ where $D' = \log(D)$. In the logarithmic space, the noise statistics observed on inverse feature rate data can be approximated as stationary which allows the Bayesian model to sample the noise standard deviation consistently. By using inference in the logarithm domain, and maintaining the sampling of the likelihood's standard deviation σ_l outside the logarithm domain, we were able to capture posterior predictions matching the noise structure, seen in Figure (4). The mean of our likelihood function μ_l takes the form of the logarithm of Eq. (7), containing sampled parameters α, k, and t_f, provided by

$$\mu_l = \log([P_f^{\alpha-1} + k(\alpha-1)(t_f - t)]^{\frac{1}{\alpha-1}}), \tag{9}$$

where t is current time, and $P_f = 0$ remains the failure criterion; thus our likelihood function may be written

$$P(D'|\Theta) = \frac{1}{\sqrt{2\pi}\sigma_l} e^{\left(-\frac{\left(D' - \log([P_f^{\alpha-1} + k(\alpha-1)(t_f-t)]^{\frac{1}{\alpha-1}}) \right)^2}{2\sigma_l^2} \right)}, \tag{10}$$

Obtaining the analytical expression for the joint posterior of parameters requires calculating high dimensional integrals, which is not often feasible. In the present case, the Bayesian model is of dimension 4, making sampling-based algorithms like Markov chain Monte Carlo (MCMC) viable to explore the posterior belief space. Sampling-based methods such as MCMC are an absolutely fundamental part of Bayesian inference, as they allow the design of more flexible and complex models of higher dimension. In the current work, we used the No-U-Turn Sampler (NUTS) within the PyMC3 python package [32] to sample the joint posterior of parameters. The NUTS sampler is an extension to the Hamiltonian Monte Carlo (HMC) algorithm which eliminates the need to manually select a desired number of steps and their size. NUTS works by building a set of likely candidate points spanning a wide range of the target distribution and stopping when the selection begins to double back on itself. We have made this selection because NUTS retains (and in some cases improves upon) HMC's ability to generate effectively independent samples efficiently [33]. The initialization method of the sampler was selected as the "jitter+adapt_diag" option built into the PYMC3 package, which worked by setting the starting point of the Bayesian sampler with a identity mass matrix, adapting a diagonal based on

the variance of the tuning samples, and adding a uniform "jitter" in $[-1, 1]$ to the starting point of each chain, detailed in [32]. We used 2 chains for each simulation to verify convergence.

To generate posterior predictions using observable data, our model utilizes predictive inference, which is derived from the general form of the Bayesian model. After observations have been recorded in D' ($D' = \log(D)$), we can predict an unknown observable, $\tilde{D}'$, using similar Bayesian logic. The distribution of $\tilde{D}'$ is called the posterior predictive distribution, and is shown Eq. (11), where the last step follows the assumed conditional independence of $\tilde{D}'$ and D' given Θ. Figure (2) shows examples of the posterior predictive distributions at two separate cycle instances [30].

$$\begin{aligned} p(\tilde{D}'|D') &= \int p(\tilde{D}'|\Theta, D')p(\Theta|D')d\Theta \\ &= \int p(\tilde{D}'|\Theta)p(\Theta|D')d\Theta \end{aligned} \tag{11}$$

To mimic the nature of failure prognosis, parameter estimations were performed by using data from a time window containing all sample points until the present cycle. This is seen in Figure (2) where data projections from the Bayesian model are generated based on data from the first cycle to the current cycle. The posterior predictions seen in figure (2) are made from synthesized data described in a subsequent section.

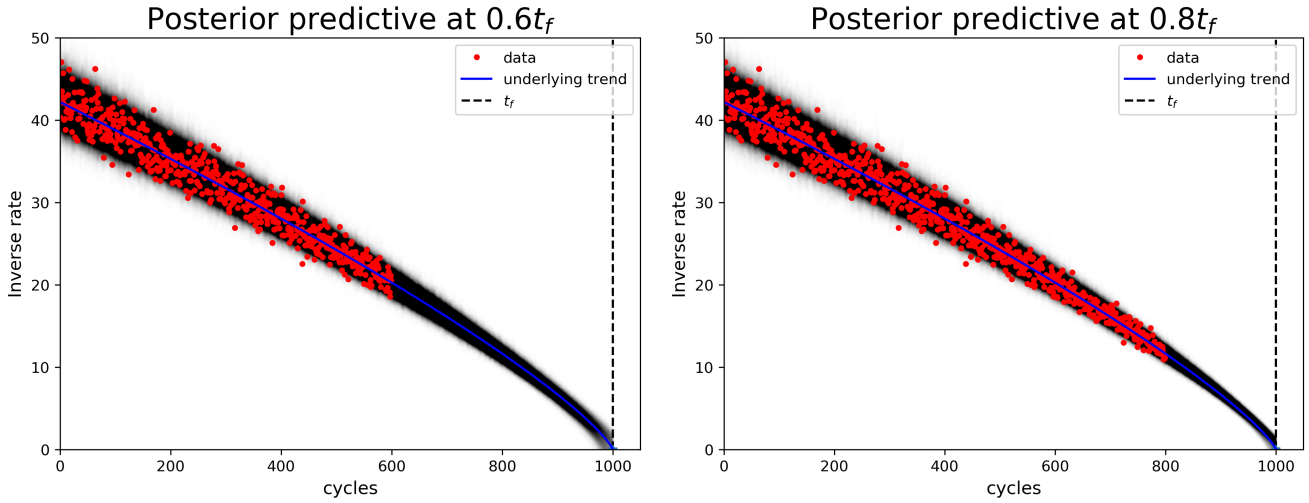

Figure 2: Predictions generated by the Bayesian model at different time windows. Figures show prediction at $0.6t_f$ (left) and $0.8t_f$ (right). Red dots show the inverse feature rate data, the black shaded regions show the posterior prediction.

Figure (2) shows a progression of the predictive capabilities of both models. The data on each plot (red points) show the data D provided to the Bayesian model (before transformation into log-space D'). The posterior prediction made by the model is shown in the shaded regions. Darker shaded regions on the plot show areas of where the posterior predictive generated more samples, corresponding to areas of higher probability of data. The heteroscedastic (appears to converge) nature of the noise of the prediction areas in the plots is a result of the logarithm space mean likelihood μ_l sampling with a "non-logarithm" space sampling of the likelihood standard deviation σ_l.

We verified the Bayesian model's posterior prediction and the linear regression model's uncertainty PDF against a Monte-Carlo brute force simulation taken from simulated data at different cycles,

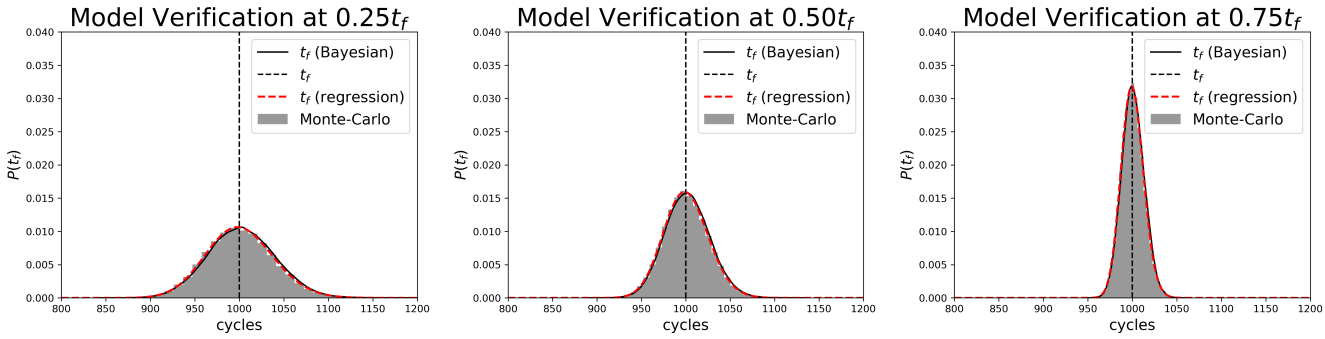

Figure 3: Verification of the agreement of the Bayesian posterior predictions of t_f, the linear regression uncertainty model, and a brute-force Monte-Carlo simulation, taken at different cycles.

corresponding with $0.25t_f$, $0.50t_f$, and $0.75t_f$. For the simplest case, in each verification in Figure (3) we assumed $\alpha = 2$ (thus not initially considered part of the Bayesian hyperparameter set), and sampling occurring at a known, prescribed time t. Figure (3) shows $p(t_f)$ provided by the linear regression uncertainty model in Eq. (7), the posterior distribution of t_f obtained from the Bayesian model, and a Monte-Carlo brute force simulation using Eq. (3) with noise added to the P_f term. We observed excellent agreement in this baseline verification process, which allows us to directly compare the two models' estimation capabilities.

We will next look at comparisons between the linear regression t_f and the Bayesian model posterior predictions of t_f using synthesized data based a fatigue experiment and then unaltered data from a real fatigue experiment.

3 Synthesized Fatigue Data Experiment

A fatigue test was simulated based on a laboratory-scale accelerated fatigue test as described in Corcoran [22]. Inverse strain rate data and time (fatigue cycles) were scaled to facilitate easier evaluation of Eq. (7), as the actual values from the test made numerical evaluation of Eq. (7) challenging. Data were generated at each fatigue cycle N, $0 < N < 1000$, where the actual time of failure from [22] was scaled to $t_f = 1000$ cycles, using Eq. (3). Such simulations were conducted at three different uncertainty/noise levels (0.02, 0.05, and 0.10), where the noise on the data P were assumed unbiased lognormal, i.e., the logarithm of the noise was normally-distributed with stationary zero mean and variance σ^2. Practically, the noise-free data was generated using Eq. (3) and normally distributed noise was added to the logarithm of the noise-free data. The resulting data was then transformed back into "non-logarithm" space to be used as the available data for both the linear regression model and the Bayesian model. This procedure was used to generate a noise structure which mimics the trend of the real accelerated fatigue test from [22]. Figure (4 top left) shows the feature rate, *R(N)* from Eq. (2), and figure (4 top right) shows inverse feature rate, *P(N)* from Eq. (3) for a synthetic data realization where $\alpha = 1.75$. Figure (4 bottom left) shows the feature rate, *R(N)* from Eq. (2), and figure (4 bottom right) shows inverse feature rate, *P(N)* from Eq. (3) for a synthetic data realization where $\alpha = 2.25$. It is apparent that the α parameter is responsible for the convexity/concavity of the inverse feature rate, and can cause substantial errors in t_f (x-intercept) estimation due to nonlinearity when $\alpha \neq 2$.

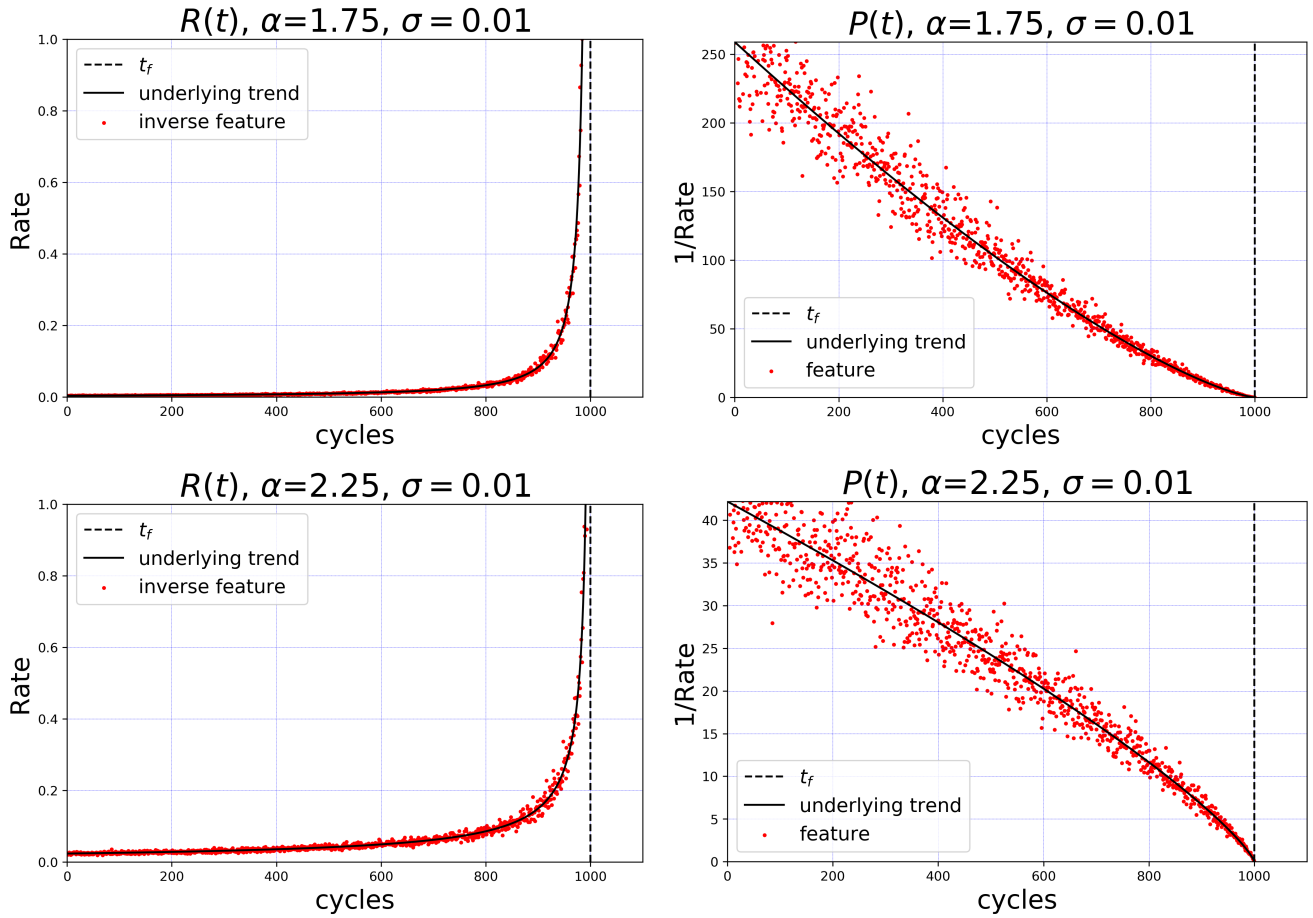

Figure 4: Realizations of simulated data with $\alpha = 1.75$ (top) and $\alpha = 2.25$ (bottom). Left plots show the feature rate (Eq. (2)), and right plots show the inverse feature rate (Eq. (3)). The linear regression and Bayesian model operated on the inverse feature rate (right plots).

3.1 Results

This section includes the results of t_f (regression) and t_f (Bayesian) on synthetic data with a variety of σ and α values: $\alpha = 1.75 - 2.25$, $\sigma = 0.02 - 0.10$. Estimations are taken with the entirety of the data until the current time without windowing, i.e. the $0.4t_f$ distribution represents using the first 40% of the data before failure. Each plot in figure (5) shows the evolution of t_f (Bayesian) and t_f (regression), for cycles $0.4t_f$ to $0.9t_f$. Figure (6) shows $t_f - \alpha$ joint distributions made by the Bayesian model.

When $\alpha < 2$, as seen in figure (5, top row), the shape of the inverse feature rate is convex, or curving upward along the horizontal axis, as seen in figure (4), as often observed in stage 1 crack growth [22]. Figure (5, top row) shows t_f (Bayesian) and t_f (regression) from data generated where $\alpha = 1.75$. The nonlinear inverse feature rate governed by $\alpha < 2$ causes t_f (regression) to greatly underestimate t_f for all levels of noise (with tight distributions), due to the linear regression model's forced linearity assumption of the inverse feature rate. Distributions of t_f (Bayesian) are able to include nonlinear inverse feature rate effects, which allows the distributions to converge onto the true t_f value, shown in t_f distributions at cycles $0.8t_f$ and $0.9t_f$ for all noise levels. In all noise cases, t_f (Bayesian) followed a trend of negatively biased, conservative estimations in prediction cycles $0.4t_f$ to $0.6t_f$, then providing a positive bias in the estimation at cycle $0.7t_f$, before centering onto the true

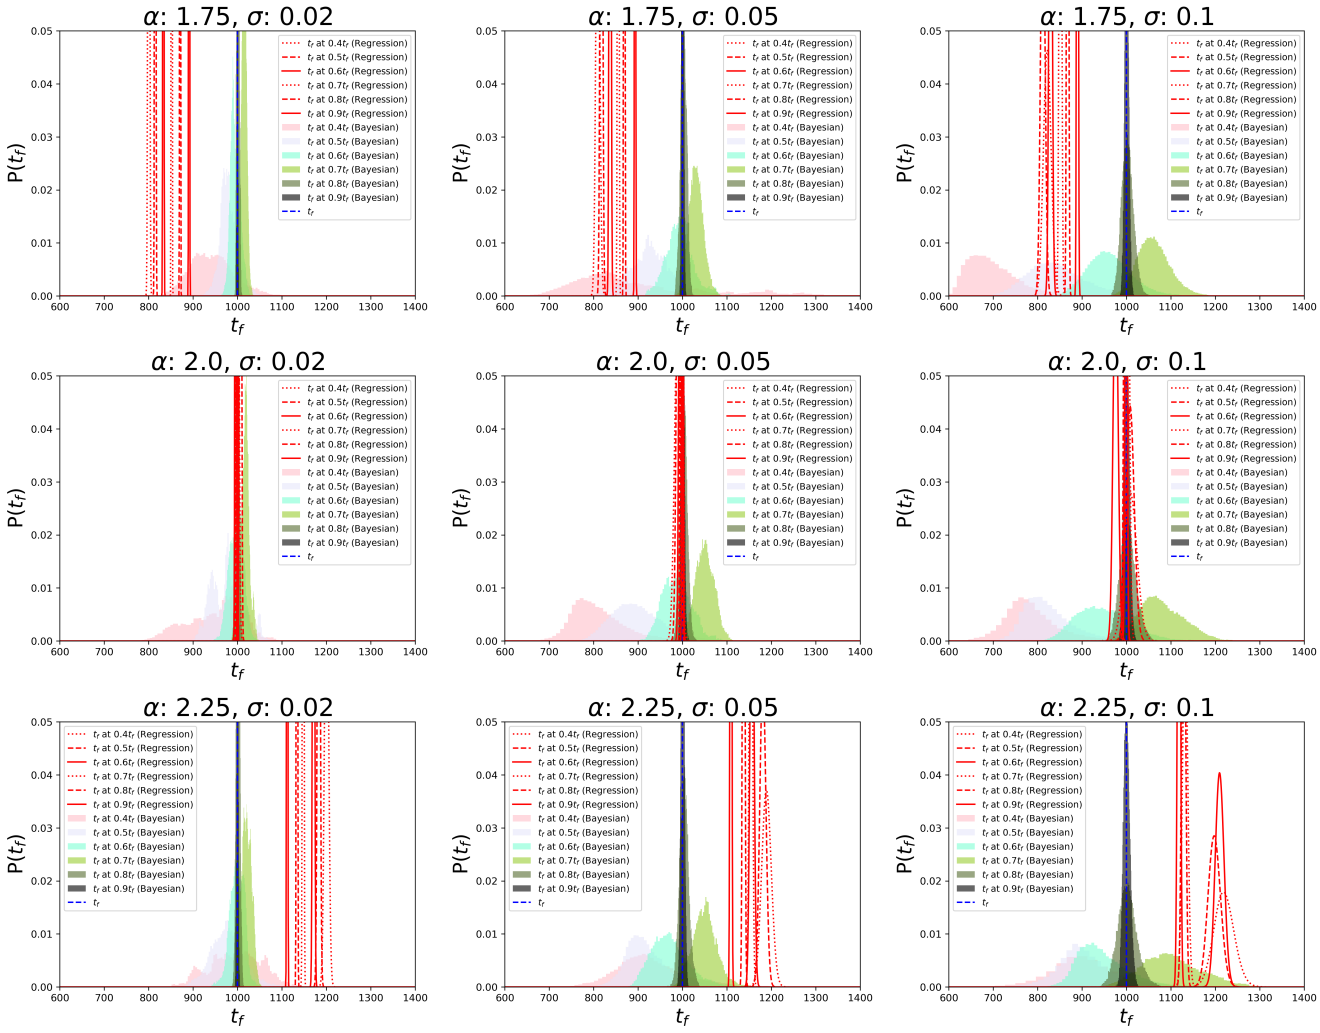

Figure 5: t_f (Bayesian) compared to t_f (regression).

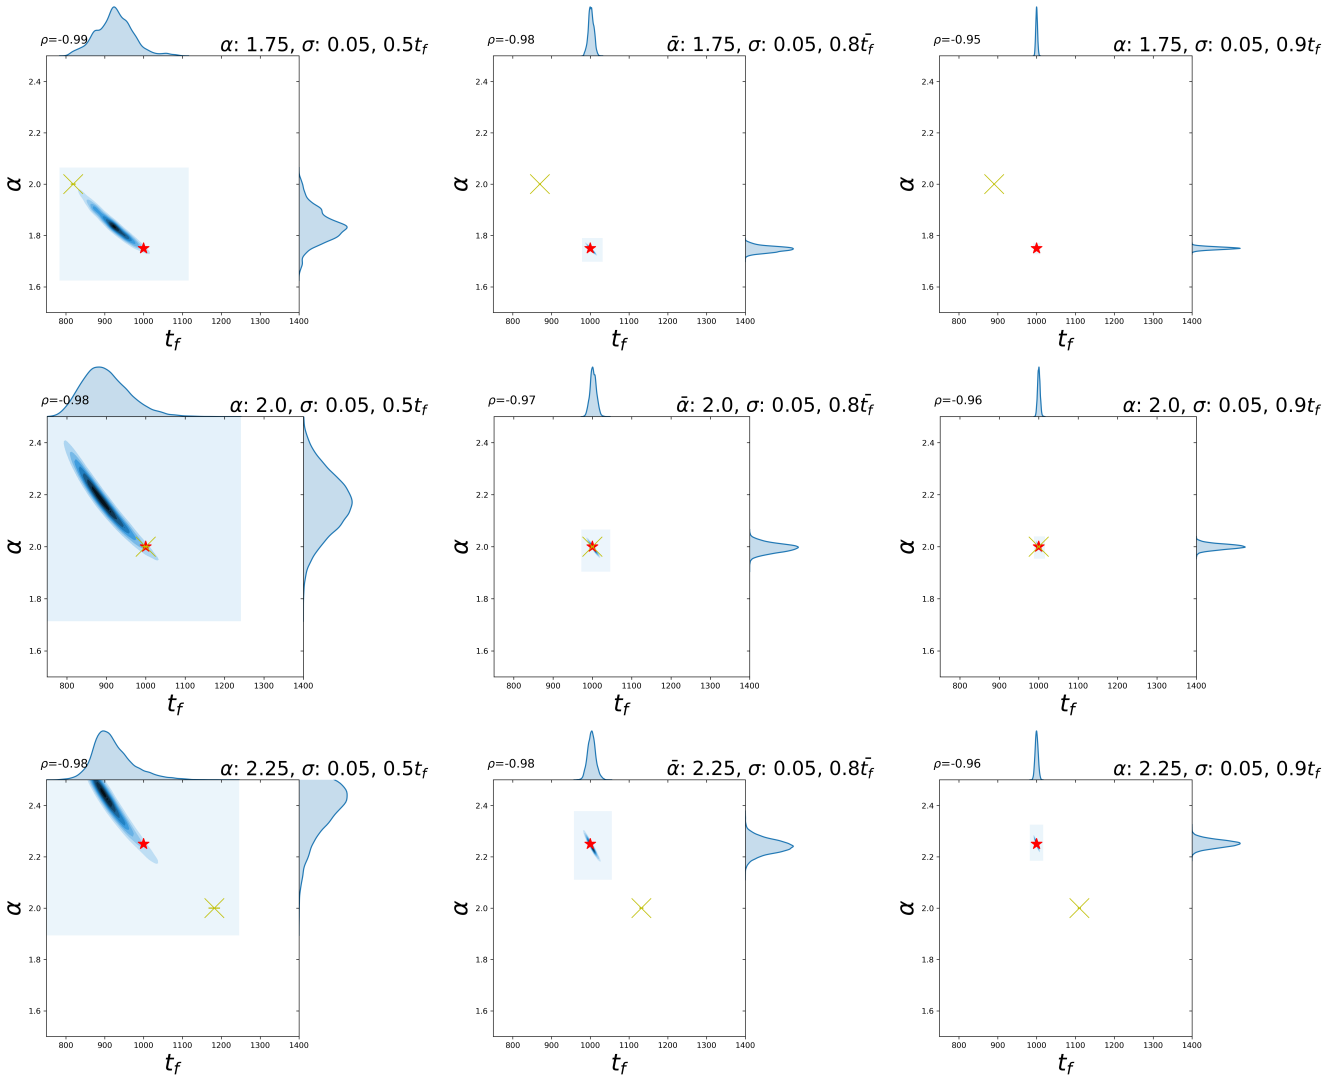

Figure 6: Joint Distributions of t_f and α distributions sampled by the Bayesian model. The yellow $\times$ is the median prediction of the linear regression model, with an error bar extending to $+/-$ the standard deviation of a best-fit normal distribution (not always visible in the figure due to their very small extent from the prediction), the blue shaded regions correspond to the Bayesian model's joint traces, and the red $\star$ is the true value of α and t_f. The linear regression marker is not estimating the value of α, and is intrinsically set at 2.0.

t_f for prediction cycles $0.8t_f$ and $0.9t_f$. In the smallest noise case $\sigma = 0.02$, all of the sampled t_f distributions included the true t_f. As expected, as more data became available to the Bayesian model, t_f (Bayesian) distributions showed increased confidence (reduced variance). In the medium and large noise cases, $\sigma = 0.05$ and $\sigma = 0.1$, early cycle predictions were too negatively biased to incorporate the true t_f until prediction cycles at $0.5t_f$ and $0.6t_f$, respectively.

One circumstance where the linear regression model outperformed the Bayesian model is when $\alpha = 2$, as seen in figure (5, middle row) as observed in stage 2 crack growth; this is because the linear regression model assumes the inverse feature rate is linear, while the Bayesian model samples α, (nonlinearity measure), unlike during the model verification in figure (3). For noise level $\sigma = 0.02$, t_f (Bayesian) incorporated t_f at each prediction cycle, but with considerably less confidence than t_f (regression). The Bayesian model, performed similarly as the $\alpha = \{1.75, 2.25\}$ cases, with increasing levels of negative bias as noise σ level increases, and being positively biased at cycle $0.7t_f$. The value of t_f (Bayesian) converged to the true t_f value at cycles $0.8t_f$ and $0.9t_f$ for all σ levels.

Simulations run using data realized with $\alpha > 2$, also caused considerable bias in t_f (regression), seen in figure (5, bottom row). As observed in stage 3 crack growth, when $\alpha > 2$ the shape of the inverse feature rate is concave, or curving downward along the x-axis, as seen in figure (4). The linear regression is based on assuming the linearity of the inverse feature rate causing t_f (regression) to be positively biased for all levels of noise. This equates to earlier-than-predicted failure, which may be undesirable from a practical standpoint. The Bayesian model was able to adapt to the nonlinear inverse feature rate, and t_f (Bayesian) distributions behaved similarly to the distributions on data realized with $\alpha = 1.75$ and $\alpha = 2.0$. The Bayesian t_f predictions had increasing levels of negative bias as noise σ level increased, but t_f (Bayesian) shown to converge onto the true t_f value by $0.8t_f$ for all cases.

In summary, in all noise levels σ and α parameters, The Bayesian model was observed to converge to include the true t_f as the prediction cycle increased. The Bayesian model performed agnostic to the true value of α, while the bias of the linear regression estimator was determined by α.

A sample of α-t_f parameter joint distribution plots generated with the Bayesian model are shown in figure (6). Darker regions in each plot correspond to areas containing more samples meaning higher probability estimates from the Bayesian model. The distributions on the outside of the plot grid represent the sampled distribution of t_f (bayesian) (top of each plot) and α (right of each plot). The star on each plot shows the exact t_f and α values. The $\times$ shows t_f (regression), which is constrained to $\alpha = 2$, with an error bar extending to $+/-$ the standard deviation of a best-fit normal distribution. It should be noted that no correlation was assumed for α and t_f in the prior beliefs, but a highly negative correlation is observed in the joint posterior distribution. For every simulation, we observed a strong negative correlation ($\rho \approx -0.95$ to -0.99) between the sampled α and t_f parameters, suggesting a next-iteration Bayesian model could incorporate a joint α-t_f distribution, or the possibility of combining terms to simplify the average term in the likelihood taken from Eq. (3) through principal component analysis.

Of course, the bias evident in the linear regression may be addressed by non-linear regression and solving for α, in addition to k and t_f, as shown in [22]. The benefit of using the Bayesian approach over non-linear regression is that it provides a full probabilistic analysis, is useful for exploring the distributions of all parameters (seen in figure 6), and may be built on to include more complexity as shown in the following section.

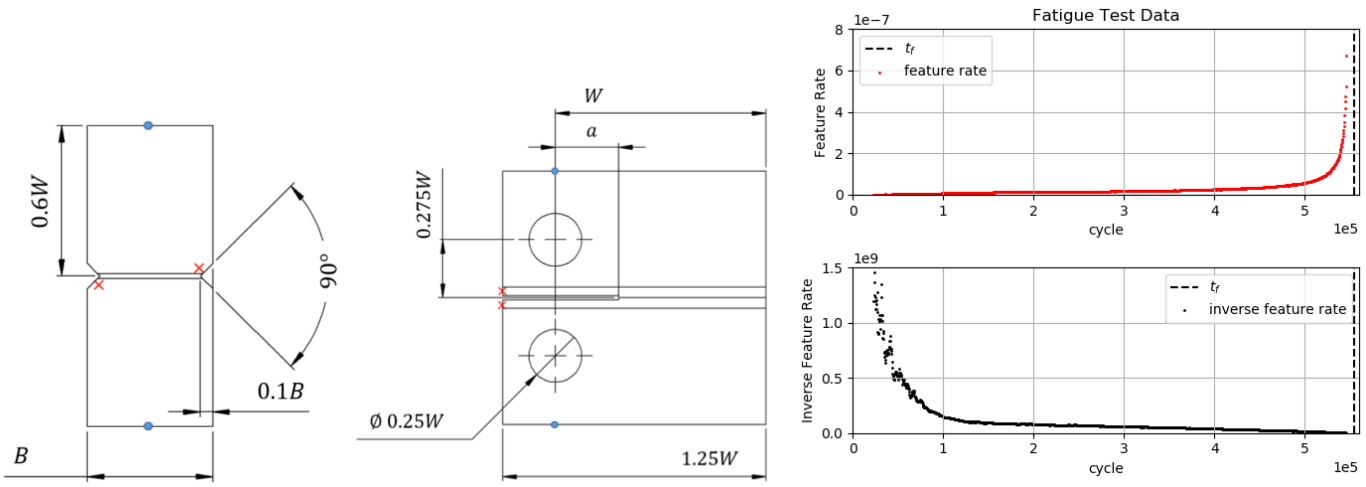

Figure 7: Left and center show geometry of the specimen used in the fatigue experiment. Right shows feature (crack growth rate) and its inverse with respect to cycle.

4 Fatigue Data Experiment

To further evaluate the Bayesian model, we performed failure prediction on another set of data from a fatigue experiment, first published in [24], again using an input feature related to inverse crack growth. This experiment differed from the first experiment presented in Section 3 because the fatigue data was used directly here, without bootstrapping or adjusting levels of noise or the α parameter. The motivation for this experiment was to test our Bayesian model on unaltered data and fully challenge the usual FFM assumptions on stationarity. The feature data was taken from a fatigue experiment using a standard 316 stainless steel compact tension test specimen, detailed in figure (7) and Tables (2) and (3). The feature was monitored using a permanently-installed potential drop measurement system, and the results plotting the rate of normalized resistance (related to crack growth through a simple polynomial) as a function of fatigue cycle are shown in figure (7, right), which also shows the inverse rate of normalized resistance [34]. The measurement in this experiment contains observations throughout the entire crack propagation event, potentially including crack growth stages 1-3, which is shown to correspond to a transient, unknown α [22, 24]. The failure mechanism of the experiment was ductile fracture, making the selection of the exact t_f cycle non-trivial. The potential drop system was removed at cycle 545,455, when the coupon began to exhibit a narrowing of cross-section (necking), and the specimen was removed from the fatigue testing machine at cycle 555,000 before complete fracture. We chose to express the actual t_f as a range between cycle 545,455 and 565,455.

Table 2: Fatigue test specimen dimensions

Parameter	**Value**
W (mm)	50
B (mm)	25
a (mm)	15.5
Maximum Load, P_{max} (kN)	11
Load Ratio, R (-)	0.1

The non-stationary nature of α (nonlinearity measure) in the real, unaltered fatigue test caused an influential difference compared to the synthesized data which was generated with a constant α value.

Table 3: Quantified uncertainties of each input parameter of the empirical crack growth law.

Parameter	**Mean Value**	**Standard Error**
Measured Crack Length, a0 (mm)	Updates with each inspection	1
Critical Crack Length, af (mm)	38	Not considered
Paris Constant, ln(C)	-25.5	0.264
Paris Exponent, m	2.88	Not considered
Maximum Load, Pmax (kN)	35	3.5
Load Ratio, R	0.1	Not considered
Geometry, Y(a)	Calculated from standards	Not considered

Performing Bayesian inference on the entire series of data (which may include multiple stages of crack growth) can introduce biased estimation due to the non-stationarity of parameters. Standard practice of implementing the FFM often includes truncating the visibly nonlinear inverse feature rate data, which can introduce human error influenced by where the user chooses to truncate. To mitigate the effects of this potential bias and subjectivity, we designed a method to truncate data estimated to be recorded during different-than-current crack growth stages. We utilized PYMC3 to develop a Bayesian model that incorporates two distinct, uncorrelated sets of Θ (all of the sampled parameters) determined by a switch point selected by the maximum posterior belief of a sampled α switch point parameter (noted as α_{sp}), i.e. completely separating the sampled parameters after a discrete cycle where α has is estimated to change, signifying the change in crack growth stage. Though crack growth has been observed to have 3 stages, our model allowed for only 2 sets of Θ because the model provided poor results when a second switch point (3 stage model) was introduced; we hypothesize that this is due to a relatively low amount of data in the third crack growth stage. This α-switching model used the same likelihood function (Eq. (10)) as the previously proposed model, with slightly altered Θ priors, summarized in Table (4). Due to previous literature showing that crack growth after stage 1 exhibits $\alpha \geq 2$ behavior [22, 24], the prior selected for the second α parameter was constrained to a uniform distribution from 2.0 to 2.2. The prior selected for the α_{sp} parameter was a uniform distribution from 0 to current time t.

Figure (8) shows the results of the α-switching model compared to t_f (regression) estimations using full and truncated data sets. For current cycle=$\{0.55t_f, 0.65t_f, 0.75t_f, 0.98t_f\}$, the Bayesian model consistently estimated α_{sp} to have a posterior distribution with a mean cycle of approximately $112,300$, around 20% through the experiment, shown on each plot. The dashed line shows the t_f (regression) using the full test data including stage 1 crack growth, which greatly underestimates the actual t_f for every estimation cycle. We also chose to show t_f (regression) using only data after the Bayesian-estimated α_{sp} cycle, which improved estimation, but still yielded negatively biased results for cycles $0.55t_f$-$0.57t_f$, and overestimated the actual t_f for cycle $0.98t_f$. We attribute the underestimation to the selection of the α_{sp}; the α-switching model selected what appears to be the inflection point of the data, leaving some nonlinear data remaining after truncation. The overestimation of the truncated t_f (regression) at $0.98t_f$ may be due to the inverse feature rate's nonlinear behavior at later fatigue cycles. In implementation, we hypothesize that t_f (regression) would benefit from truncating data from a later cycle. The Bayesian α-switching model was able to provide more accurate t_f estimations which were able to include the actual t_f range for all cycle predictions, and we saw our model converging towards the normal distribution as more data was made available. The interval of the actual t_f was

represented in high density regions of each (Bayesian) t_f PDF, with high accuracy in cycles past $0.65t_f$. The α-switching Bayesian model was able to objectively provide a distribution for the crack stage growth transition cycle, improving the accuracy of t_f (regression) when using truncated data, while also providing highly accurate (Bayesian) t_f estimation. During future use of the Bayesian α-switching model, feature data may be provided to the model within only a single stage of crack growth. To prevent the model from falsely selecting an α switch-point and separating model parameters, the uniformity and convergence of α_{sp} can be evaluated to motivate the selection of the number of switch-points (if any). Beyond the FFM, the identification of switch-points may be useful in analyzing data through conventional analysis such as Paris Law; segmenting the data allows more accurate characterization of the different stages.

Table 4: Parameters of the α-switching Bayesian model (real data)

Parameter	**Prior**
t_f	Improper uniform distribution with lower bound of t
k_1	Improper uniform distribution with lower bound of 0
k_2	Improper uniform distribution with lower bound of 0
α_1	Uniform distribution between 1.5 and 2.5
α_2	Uniform distribution between 2.0 and 2.2
α_{sp}	Uniform distribution between 0 and t
σ_l	Half normal distribution, standard deviation of 0.8

5 Summary and Conclusions

When limited data are available for analysis, future events can be difficult to accurately predict. For failure events (t_f) which exhibit positive feedback failure mechanisms, the failure forecast method (FFM) allows for t_f predictions made using a single measured feature, unaffected by application specific experimental parameters and their uncertainty. Classically, the FFM implementation assumes linearity of the inverse feature rate, and creates a single estimation of t_f from a single realization of data from linear regression. Assuming the inverse feature rate's linearity causes a rigidity in the FFM because the governing phenomena often exhibit nonlinear inverse feature rate behavior, which has been observed in early and late stage crack growth [15, 16]. In this paper, we developed a Bayesian statistical model which relaxes this linearity assumption, and samples the posterior distribution of t_f, allowing for the generation of probability distributions for t_f.

For the specific application of fatigue crack growth, t_f (Bayesian) was compared against a statistical model of the "classic" FFM implementation, t_f (regression). We first compared the two models on synthetic data based on a real accelerated fatigue test published in [24], and then compared both models' t_f estimation made by the models on an unaltered accelerated fatigue test from [34].

In the study performed with synthetic data, t_f (Bayesian) was able to accurately converge to t_f estimation to the true t_f values for $\alpha = \{1.75, 2.0, 2.25\}$ across all noise levels, while the linear regression model was only able to converge to the correct t_f for data generated with $\alpha = 2.0$ across all noise levels. t_f estimation results of this study are shown in figures (5) and (6). The Bayesian model behaved mostly independent of the true α value of each synthetic data simulation, save for slight

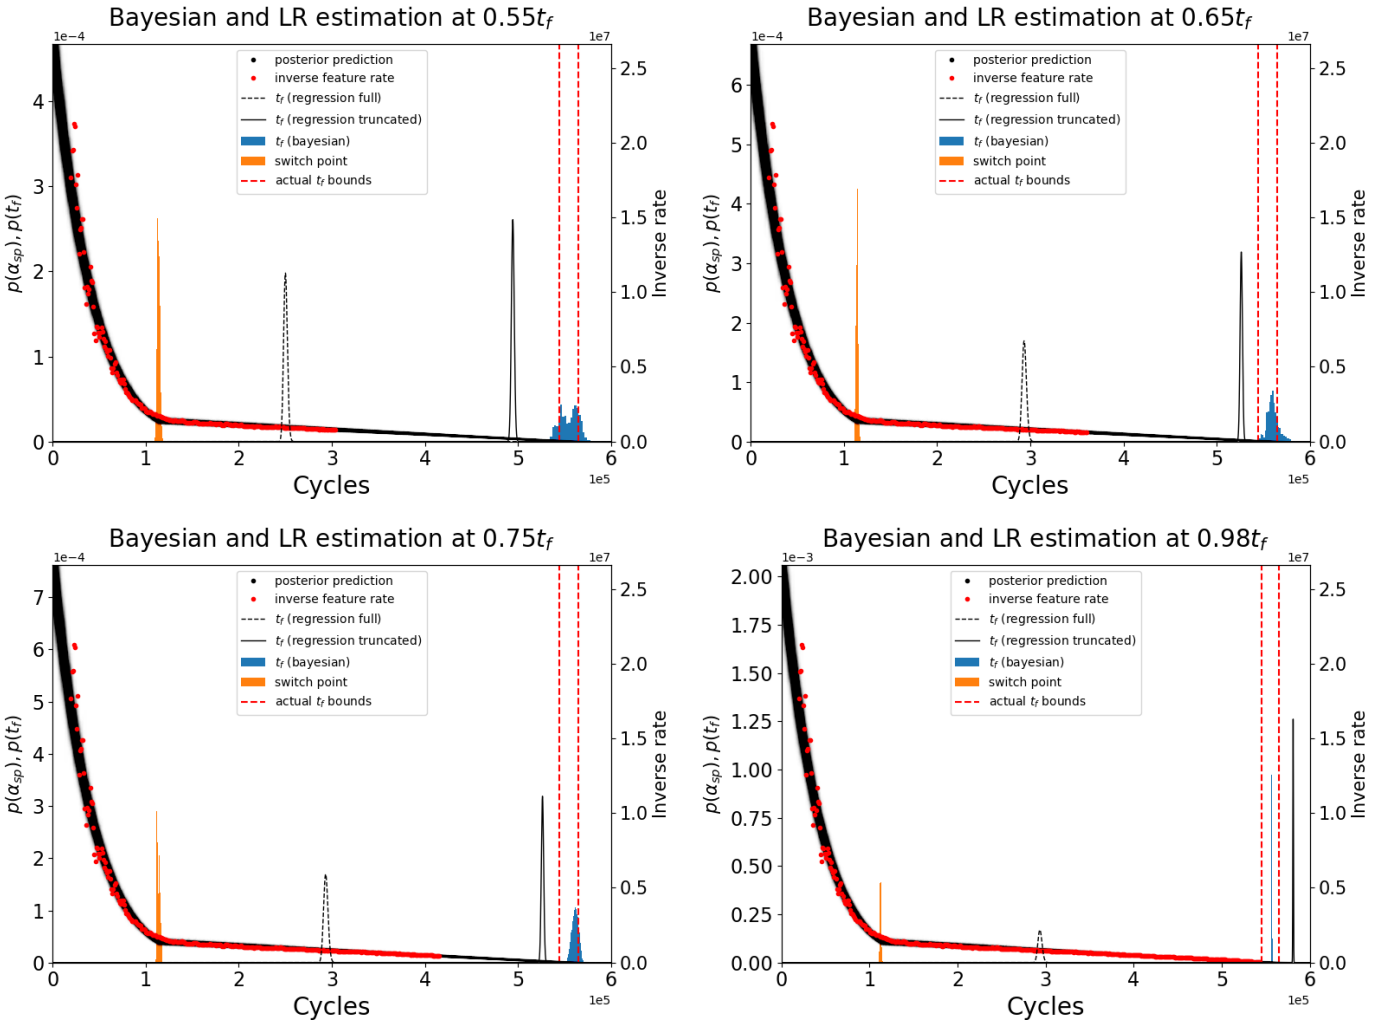

Figure 8: Results from using the α-switching Bayesian model and linear regression model on data from a real un-altered fatigue test for cycle predictions $0.55t_f$-$0.98t_f$. Black shaded regions show the posterior prediction, orange regions show the posterior distribution of α_{sp}, blue regions show the t_f (Bayesian) posterior distribution, dashed and solid black lines show the t_f (regression) distribution for non-truncated data and truncated data, respectively. Inverse feature rate data is shown in red dots, and the range of actual t_f is shown in dashed red lines. Primary y-axis has units of probability density for distributions $p(\alpha_{sp})$ and $p(t_f)$, and secondary axis has units of inverse feature rate for experimental data and Bayesian posterior predictive.

increase in t_f distribution spread as the α value increased. The t_f (Bayesian) estimations across all α levels developed a negative bias at cycles $0.4t_f - 0.6t_f$, provided a positive bias at cycle $0.7t_f$, before converging to the true t_f value at cycles $0.8t_f$ and $0.9t_f$. The amount of bias was proportional to the level of noise. Joint distributions of sampled variables t_f-α in the Bayesian model also exhibited the same trend for all α values; these variables were observed to be highly negatively correlated.

The t_f estimations from both models were also compared on data from an un-altered fatigue test from [34]. To account for the non-stationarity often observed in the inverse feature rate data, the Bayesian model was adapted to incorporate an α switch point, allowing for two sets of Θ (Bayesian model parameters) to be sampled before and after the α switch point. The α-switching Bayesian model was able to accurately estimate t_f for cycles $0.55t_f$ onward, by effectively sampling the transition point α_{sp} and truncating the data after stage 1 crack growth, which was predicted to occur approximately 20% through the fatigue experiment. The α-switching Bayesian model was also used to estimate a truncation point for the linear regression model, which performed much more accurately when using data only after α_{sp}.

The benefits of truncating the early stage crack growth motivates future work which explores the benefits of windowing data, impacting the t_f for both models. Some discussion is presented in [23], regarding windowing feature data for the linear regression model, but there is no work exploring the effects on the Bayesian model of windowing data.

As seen in figure (6), the Bayesian model showed a very high negative correlation between t_f and α. This result also motivates future iterations of the Bayesian model which accounts for this correlation by establishing jointly distributed t_f and α parameters, or simplifying the model based on principal component analysis.

Acknowledgements

The U.S.A. authors acknowledge support for this work was provided by the University of Dayton on subcontract FA8650-16-D-0311 under master contract with Air Force Research Laboratory/Eglin Air Force Base. The U.K. authors acknowledge support from the UK Engineering and Physical Sciences Research Council via the UK Research Centre in NDE, EP/L022125/1.

References

[1] C. R. Farrar and K. Worden, "An introduction to structural health monitoring.," Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences, 2007.

[2] N. Chen and K. L. Tsui, "Condition monitoring and remaining useful life prediction using degradation signals: revisited," IIE Transactions, vol. 45, pp. 939–952, sep 2013.

[3] X.-S. Si, W. Wang, C.-H. Hu, D.-H. Zhou, and M. G. Pecht, "Remaining useful life estimation based on a nonlinear diffusion degradation process," IEEE Transactions on Reliability, vol. 61, no. 1, pp. 50–67, 2012.

[4] S. K. Everton, M. Hirsch, P. Stravroulakis, R. K. Leach, and A. T. Clare, "Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing," Materials & Design, vol. 95, pp. 431–445, 2016.

[5] P. Paris and F. Erdogan, "A critical analysis of crack propagation laws," ASME Journal of Basic Engineering, vol. 85, no. 1, pp. 528–534, 1963.

[6] X.-S. Si, W. Wang, C.-H. Hu, and D.-H. Zhou, "Remaining useful life estimation–a review on the statistical data driven approaches," European Journal of Operational Research, vol. 213, no. 1, pp. 1–14, 2011.

[7] J. Z. Sikorska, M. Hodkiewicz, and L. Ma, "Prognostic modelling options for remaining useful life estimation by industry," Mechanical Systems and Signal Processing, vol. 25, no. 5, pp. 1803–1836, 2011.

[8] S. Beretta and M. Carboni, "Experiments and stochastic model for propagation lifetime of railway axles," Engineering fracture mechanics, vol. 73, no. 17, pp. 2627–2641, 2006.

[9] S. Beretta and A. Villa, "A RV approach for the analysis of fatigue crack growth with NASGRO equation," in 4th International ASRANAT Colloquium, pp. 1–7, 2008.

[10] C. Mallor, S. Calvo, J. L. Núñez, R. Rodríguez-Barrachina, and A. Landaberea, "Full second-order approach for expected value and variance prediction of probabilistic fatigue crack growth life," International Journal of Fatigue, vol. 133, p. 105454, 2020.

[11] M. Jimenez-Martinez and M. Alfaro-Ponce, "Fatigue damage effect approach by artificial neural network," International Journal of Fatigue, vol. 124, pp. 42–47, 2019.

[12] J. F. Barbosa, J. A. F. O. Correia, R. C. S. F. Júnior, and A. M. P. De Jesus, "Fatigue life prediction of metallic materials considering mean stress effects by means of an artificial neural network," International Journal of Fatigue, vol. 135, p. 105527, 2020.

[13] T. FUKUZONO, "A new method for predicting the failure time of a slope," in Proceedings of 4th International Conference and Field Workshop on Landslide., 1985, pp. 145–150, 1985.

[14] B. Voight, "A method for prediction of volcanic eruptions," Nature, 1988.

[15] B. Voight, "A relation to describe rate-dependent material failure," Science, 1989.

[16] R. R. Cornelius and P. A. Scott, "A materials failure relation of accelerating creep as empirical description of damage accumulation," Rock Mechanics and Rock Engineering, vol. 26, no. 3, pp. 233–252, 1993.

[17] R. R. Cornelius and B. Voight, "Seismological aspects of the 1989–1990 eruption at Redoubt Volcano, Alaska: The Materials Failure Forecast Method (FFM) with RSAM and SSAM seismic data," Journal of Volcanology and Geothermal Research, vol. 62, no. 1-4, pp. 469–498, 1994.

[18] R. R. Cornelius and B. Voight, "Graphical and PC-software analysis of volcano eruption precursors according to the Materials Failure Forecast Method (FFM)," Journal of Volcanology and Geothermal Research, vol. 64, no. 3-4, pp. 295–320, 1995.

[19] G. B. Crosta and F. Agliardi, "Failure forecast for large rock slides by surface displacement measurements," Canadian Geotechnical Journal, vol. 40, no. 1, pp. 176–191, 2003.

[20] A. Boue, P. Lesage, G. Cortés, B. Valette, and G. Reyes-Dávila, "Real-time eruption forecasting using the material Failure Forecast Method with a Bayesian approach," Journal of Geophysical Research: Solid Earth, vol. 120, no. 4, pp. 2143–2161, 2015.

[21] Y. Lavallée, P. G. Meredith, D. B. Dingwell, K.-U. Hess, J. Wassermann, B. Cordonnier, A. Gerik, and J. H. Kruhl, "Seismogenic lavas and explosive eruption forecasting," Nature, vol. 453, no. 7194, p. 507, 2008.

[22] J. Corcoran, "Rate-based structural health monitoring using permanently installed sensors," Proc. R. Soc. A, vol. 473, no. 2205, p. 20170270, 2017.

[23] M. D. Todd, M. Leung, and J. Corcoran, "A probability density function for uncertainty quantification in the failure forecast method," in Proceedings of the 9th European Workshop on Structural Health Monitoring, Manchester, UK July, pp. 9–11, 2018.

[24] M. S. H. Leung, J. Corcoran, P. Cawley, and M. D. Todd, "Evaluating the use of rate-based monitoring for improved fatigue remnant life predictions," International Journal of Fatigue, vol. 120, pp. 162–174, 2019.

[25] A. F. Bell, M. Naylor, M. J. Heap, and I. G. Main, "Forecasting volcanic eruptions and other material failure phenomena: An evaluation of the failure forecast method," Geophysical Research Letters, 2011.

[26] A. Bevilacqua, E. B. Pitman, A. Patra, A. Neri, and M. Bursik, "Probabilistic Enhancement of the Failure Forecast Method Using a Stochastic Differential Equation and Application to Volcanic Eruption Forecasts," Frontiers in Earth Science, 2019.

[27] A. Boue, P. Lesage, G. Cortés, B. Valette, G. Reyes-Dávila, R. Arámbula-Mendoza, and A. Budi-Santoso, "Performance of the 'material Failure Forecast Method'in real-time situations: A Bayesian approach applied on effusive and explosive eruptions," Journal of Volcanology and Geothermal Research, vol. 327, pp. 622–633, 2016.

[28] A. F. Bell, M. Naylor, and I. G. Main, "The limits of predictability of volcanic eruptions from accelerating rates of earthquakes," Geophysical Journal International, vol. 194, no. 3, pp. 1541–1553, 2013.

[29] D. V. Hinkley, "On the ratio of two correlated normal random variables," Biometrika, vol. 56, no. 3, pp. 635–639, 1969.

[30] A. Gelman, H. S. Stern, J. B. Carlin, D. B. Dunson, A. Vehtari, and D. B. Rubin, Bayesian Data Analysis. Chapman and Hall/CRC, 2013.

[31] A. Gelman, D. Simpson, and M. Betancourt, "The prior can often only be understood in the context of the likelihood," Entropy, vol. 19, no. 10, p. 555, 2017.

[32] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, "Probabilistic programming in Python using PyMC3," PeerJ Computer Science, vol. 2, p. e55, 2016.

[33] M. D. Hoffman and A. Gelman, "The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.," Journal of Machine Learning Research, vol. 15, no. 1, pp. 1593–1623, 2014.

[34] J. Corcoran, C. M. Davies, P. Cawley, and P. B. Nagy, "A Quasi-DC Potential Drop Measurement System for Material Testing," IEEE Transactions on Instrumentation and Measurement, vol. 69, no. 4, pp. 1313–1326, 2019.