Comprehensive Material Characterization and Simultaneous Model Calibration for Improved Computational Simulation Credibility

D. Thomas Seidl, Elizabeth M.C. Jones, and Brian T. Lester

September 13, 2021

SAND21XX-XXXXR

LDRD Project Number: 222471

Abstract

Computational simulation is increasingly relied upon for high-consequence engineering decisions, and a foundational element to solid mechanics simulations is a credible material model. Our ultimate vision is to interlace material characterization and model calibration in a real-time feedback loop, where the current model calibration results will drive the experiment to load regimes that add the most useful information to reduce parameter uncertainty. The current work investigated one key step to this Interlaced Characterization and Calibration (ICC) paradigm, using a finite load-path tree to incorporate history/path dependency of nonlinear material models into a network of surrogate models that replace computationally-expensive finite-element analyses. Our reference simulation was an elastoplastic material point subject to biaxial deformation with a Hill anisotropic yield criterion. Training data was generated using either a space-filling or adaptive sampling method, and surrogates were built using either Gaussian process or polynomial chaos expansion methods. Surrogate error was evaluated to be on the order of 10^{-5} and 10^{-3} percent for the space-filling and adaptive sampling training data, respectively. Direct Bayesian inference was performed with the surrogate network and with the reference material point simulator, and results agreed to within 3 significant figures for the mean parameter values, with a reduction in computational cost over 5 orders of magnitude. These results bought down risk regarding the surrogate network and facilitated a successful FY22-24 full LDRD proposal to research and develop the complete ICC paradigm.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Introduction and Executive Summary

Material (constitutive) models are mathematical representations of the complex responses of materials that predict stress in terms of loading (tension, compression, shear) and environmental histories (temperature and strain rate). As such, they are critical to any solid mechanics simulation like Finite Element Analysis (FEA). Often, there are multiple models that are applicable to a class of materials (metals), varying in the phenomena they represent (isotropic versus kinematic hardening) or assumed functions (different (an)isotropic yield surfaces). Each model has a unique set of material-dependent parameters that must be determined. Collecting the experimental data needed to identify these parameters is called material characterization, and the process of determining the parameter values is called model calibration. Selecting an appropriate material model is usually a balance between the desired fidelity, and the time and cost of characterization, calibration, and simulation.

Historically, characterization and calibration procedures rely predominately on tensile dog bone samples, which develop a homogeneous state of stress under uniaxial loading, allowing simple models to be calibrated analytically with global data (load-displacement curves). However, this simplicity comes at a cost: characterizing materials under different environmental conditions individually (strain rate, temperature, orientation) can quickly lead to a costly and time consuming test matrix. Further, models calibrated from simple monotonically-increasing, homogenous stress data may inhibit predictivity of FEA of components under complex, real-world loading conditions. Additionally, an appropriate material model and corresponding text matrix are typically selected a priori based on subject matter expertise, and only after a separate validation effort is performed can calibration results be evaluated. If they are unsatisfactory, a new model must be selected, potentially necessitating another experimental campaign [1–7]; overall, months if not years may be required to obtain a satisfactory calibration.

Improving upon tensile dog bones, more complicated specimens and test methods, such as notched tension, shear, torsion, interrupted or reversed loading, etc. can be employed to access other stress states. Analytical solutions for these experiments are usually unavailable, necessitating advanced calibration methods, such as finite element model updating (FEMU). To provide richer experimental data for calibrations, the community has begun employing advanced characterization methods like Digital Image Correlation (DIC), which delivers reliable, full-field, kinematic data (strains) [8]. Correspondingly, FEMU has been extended and other calibration methods like the Virtual Fields Method (VFM) have been developed [9–12], which take advantage of heterogeneous field data in the calibration cost function [13–16]. Because calibrations with high-dimensional field

data remain challenging, approaches such as wavelet decomposition, where the salient features of the data are preserved in a truncated series of modes, are currently being investigated [17, 18]. At Sandia, many of these advanced techniques are included or being added to the MatCal calibration software [19].

Despite these advances, applying these techniques for calibration of complicated, high-fidelity material models remains challenging. As a result, analysts may be compelled to revert to simpler characterization and calibration methods, as well as lower-fidelity models, which inevitably affects the predictivity of computational simulations and design or qualification decisions based thereon. Addressing these issues requires a departure from any existing schemes and the development of novel and transformative techniques to improve both the outcome and the process of material characterization and model calibration.

We propose to transform the paradigm by interlacing material characterization and model calibration to create a real-time feedback loop where in-situ calibration results actively drive the experiment through loadings that best improve the calibration and reduce parameter uncertainty. The "Interlaced Characterization and Calibration " (ICC) paradigm is summarized in Fig. 1 and has five key components: First, a multi-axial load frame subjects the specimen to complex stress states instead of uniaxial loading (Fig. 1b). Next, advanced, full-field diagnostics deliver rich characterization data compared to traditional load-displacement data (Fig. 1c). Third, a suite of models is calibrated with quantified parameter uncertainties using Bayesian inference, as opposed to a deterministic calibration of a single model (Fig. 1d). Fourth, the load frame is actively driven to loading regimes that will provide the most useful information to reduce parameter uncertainties (Fig. 1e). Finally, because Bayesian inference is prohibitively expensive with FEA, surrogate models will be pre-built offline prior to entering the feedback loop (Fig. 1a).

The complete ICC paradigm is the subject of a full LDRD project for FY22-24 entitled "Interlaced Characterization and Calibration for Improved Computational Simulation Credibility". The primary focus of this FY21 exploratory project was to investigate the feasibility of building a surrogate network based on the finite load-path tree (Fig. 1a). Existing versatile calibration approaches such as FEMU rely upon extensive use of FEA, requiring hundreds to thousands of simulations, with each simulation taking minutes to hours to run. Using these methods in real-time is infeasible. To overcome this issue, we will build surrogate models, which take only a fraction of a second to evaluate, in an offline, pre-computation stage. The surrogates directly replace FEA, taking material model parameters as inputs, and producing observable QoIs (e.g. DIC and FEA strain data) as outputs. Such surrogate models are normally state-dependent, meaning they produce QoIs for a single state of the specimen or instant in time within the loading. Therefore, a challenging aspect

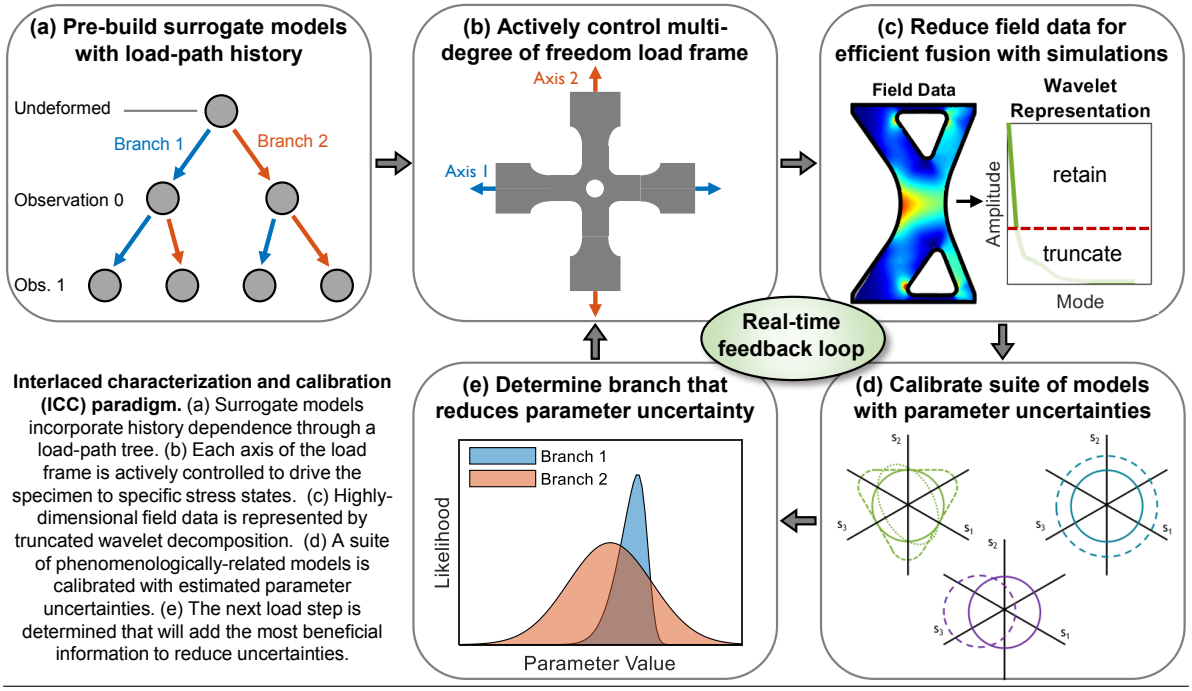

Figure 1: Overview of the complete ICC calibration workflow. Our work in this exploratory project concerns blocks (a) and (d).

of this approach is the incorporation of the history-dependence of plastically-deforming materials into the surrogates. To address path-dependency, we will reduce the infinite load-path space to a load-path tree with a finite number of load steps and branches at each step (Fig. 1a). Each branch represents a different potential load path, such as increasing/decreasing force along each axis of a quad-axis load frame (Fig. 1b).

Our approach was to start with a simple material point simulator (essentially a single-element FEA), where the two branches of the load-path tree correspond to two in-plane strains. We selected the Hill48 anisotropic yield criterion [20] as our material model, and focused on two of the six model parameters for demonstration purposes. The surrogate model takes the two model parameters as inputs, and produces the two in-plane stresses and one out-of-plane strain as outputs. A separate surrogate was built for each node in the load-path tree (gray circles in Fig. 1a), using

either Gaussian process [21, 22] or polynomial chaos expansion techniques [23, 24], with adaptive sampling methods [25–27]. Because the material point simulator itself is computationally inexpensive, we were able to show that the surrogates had extremely high accuracy compared to the original material point simulator. Additionally, we performed Bayesian inference at each node in the load-path tree and showed that certain load paths led to reduced parameter uncertainty compared to others, as hypothesized. This work has bought down risk associated with the surrogate network and load-path tree, and led to a successful proposal for FY22-24 for a full Engineering Science Research Foundation LDRD project to research the complete ICC paradigm.

Detailed Description of Research and Development and Methodology

Surrogate Network

Our means of reducing the infinite-dimensional load-path space is through the introduction of a graph we call the *surrogate network*, a rooted tree where each node represents a specific point in load-path space. In this way, we discretize the unbounded, original load-path space into a finite description we can analyze. Every node (other than the root node which represents the underformed specimen) has a corresponding surrogate model that approximately maps constitutive model parameters to quantities of interest (QoIs) that can, in principle, be computed from experimental measurements.

A prototypical example of a surrogate network for a two-dimensional loading space is depicted in Figure 2. The graph structure in this network is that of a *binary tree*, as each parent node has two children. We imagine that this graph stands in for an experiment that involves biaxial loading of a specimen, such that each left/right, orange/blue arrow represents an increment of applied strain along an axis.

The surrogate-based, online optimal experimental design (OED) workflow begins at the root node, where an initial deformation direction is chosen. For the purposes of this exploratory project, we assume that we deform the sample by a set amount of strain along the left branch, such that we arrive at the state represented by Node 1. Then data is collected in the experiment to produce observation 0, and Bayesian calibration is performed to compute the posterior distribution of the model parameters. We then must solve an instance of the *step selection problem*, i.e. given our current information about that parameters, which step (Node 3 or 4) is the most desirable for calibration? While critically important, the step selection problem was deemed out of scope for this year's exploratory project, and we revisit it in the discussion.

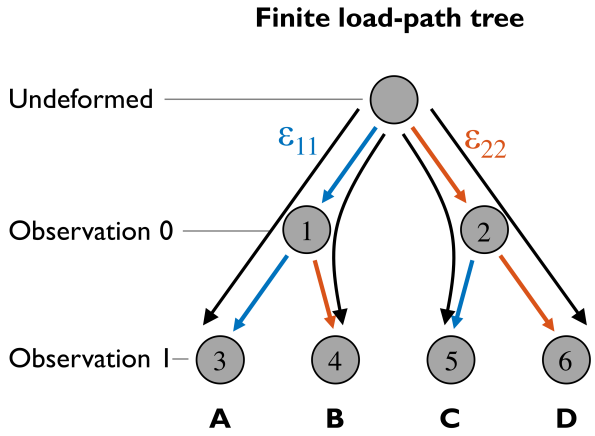

Figure 2: The exemplar material point simulator surrogate network considered in this study. It has four possible load-paths (A-D).

Now that a prototypical surrogate network has been described, we next introduce the material point simulator model for test specimen. Our goal is to replace the material point simulator with surrogate models that emulate it over some range of input parameters to a desired degree of accuracy. In the actual application, more costly finite element models will be utilized, but our use of the material point simulator as a reference simulation provides us with a simpler starting point to demonstrate the concept of the surrogate network.

Material Point Simulator

We consider as a representative model, an elastic-plastic material subject to rate- and temperature-independent deformation with isotropic hardening and anisotropic yield. A small-strain description is used although this and the previously stated assumptions may all be relaxed in subsequent investigations. For this particular effort, we utilize the Hill anisotropic yield model [20] for a material point subject to biaxial deformation under plane stress (i.e. $\sigma_{33} = 0$). The material is taken to be under purely biaxial loading such that shear stress and strain terms (i.e. σ_{ij} and ε_{ij} for $i \neq j$) are zero. Plastic flow generates a permanent deformation in the material that persists after all applied loads are removed. These effects are represented in the constitutive model through the linearized plastic strain tensor ε^p_{ij} and isotropic hardening variable κ.

Small-strain elastoplasticity assumes an additive decomposition of the total linearized strain tensor ε_{ij} into elastic ε^e_{ij} and plastic ε^p_{ij} components such that $\varepsilon_{ij} = \varepsilon^e_{ij} + \varepsilon^p_{ij}$. We note that in this model, plastic deformation is isochoric and consequently $\varepsilon^p_{kk} = 0$.

Elastoplastic behavior is characterized by a yield function $f := \phi - \bar{\sigma}$, where ϕ is the effective stress and $\bar{\sigma}$ is the current yield (flow) stress. If the yield function is less than zero (i.e. $f < 0$) the material behavior is linear elastic and there is no change in ε^p_{ij} and κ. However, if the yield function is zero, plastic deformation occurs, and these variables evolve according to the *flow rule*, *hardening law*, and *consistency constraint* ($f = 0$). The Cauchy stress is always determined by the relation

$$\begin{aligned} \sigma_{ij} &= \lambda \varepsilon^e_{kk} \delta_{ij} + 2\mu \varepsilon^e_{ij}, \\ &= \lambda \varepsilon_{kk} \delta_{ij} + 2\mu \left(\varepsilon_{ij} - \varepsilon^p_{ij} \right), \end{aligned} \tag{1}$$

where λ and μ are the Lamé constants. They are related to Young's modulus E and Poisson's ratio ν through $\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}$ and $\mu = \frac{E}{2(1+\nu)}$.

The yield function for this biaxial, plane stress model is

$$\phi = \left(F \left(\sigma_{22} \right)^2 + G \left(\sigma_{11} \right)^2 + H \left(\sigma_{11} - \sigma_{22} \right)^2 \right)^{\frac{1}{2}}, \tag{2}$$

where F, G, and H are the *Hill coefficients*. The yield stress is a monotonically increasing function of the isotropic hardening variable and is given by

$$\bar{\sigma}(\kappa) = \sigma_y + \bar{h}\kappa + A \left(1 - \exp\left(-n\kappa \right) \right), \tag{3}$$

with σ_y the initial yield stress, $\bar{h}$, the linear hardening modulus, A the Voce modulus, and n the Voce exponent. Lastly, the flow rule and hardening law may be combined to obtain the expression

$$\dot{\varepsilon}_{\gamma\gamma} = \dot{\kappa} \frac{\partial \phi}{\partial \sigma_{\gamma\gamma}}, \quad \gamma = 1, 2, 3 \text{ (no sum over } \gamma). \tag{4}$$

We collect the unknown state variables for this model in a vector $\boldsymbol{x} := \{\sigma_{11}, \sigma_{22}, \varepsilon_{33}, \kappa, \varepsilon^p_{11}, \varepsilon^p_{22}, \varepsilon^p_{33}\}$.

The model is path-independent and non-dissipative prior to the onset of plastic deformation (i.e. $\varepsilon_{ij}^{p} = 0$ and $\kappa = 0$), but it becomes path-dependent when plastic deformation occurs.

The material point simulator is discretized in pseudo-time using a backward Euler approach and is driven by the in-plane strains ε_{11} and ε_{22}. When a trial elastic step results in a yield function $f > 0$, the deformation is plastic and a return-mapping algorithm (which amounts to the solution of a nonlinear system of equations) is executed. Otherwise, the trial state is accepted, and the plastic state variables do not evolve [28].

Parameter	Value
E	70 GPa
ν	0.3
σ_y	200 MPa
h	200 MPa
A	200 MPa
n	20
F	0.55
G	0.4
H	0.5

Table 1: Parameter values for the material point simulator used to generate the stress curves in Figure 3.

Stress curves for load-paths A and B (c.f. Figure 2) with the parameters from Table 1 are shown in Figure 3. Points where experimental measurements are collected are represented by black diamonds. In the current study, experimental data corresponds to simulation results with known input parameters. The stress curves for the experiments represented by load-path A resembles traditional uniaxial stress-strain curves, except that here ε_{22} is fixed at zero, which causes the development of the σ_{22} stress. The reversal in the direction of applied strain for load-path B produces striking changes in the in-plane stresses, and the exponential nature of the hardening behavior is also easily observed.

While the material point simulator is certainly lower in fidelity than a FE model of the test specimen, we note that it is capable of producing the history-dependent behavior that is characteristic of plastic deformation. Furthermore, the anisotropic nature of the Hill model and nonlinear hardening behavior affect the in-plane stresses in complex ways when the direction of the applied strain is reversed. Another favorable property of the material point simulator is that is takes less than a second to run for a given load-path, which greatly reduces the computational burden of producing data for our initial forays into surrogate modeling. Lastly, the cost of the simulator is cheap enough

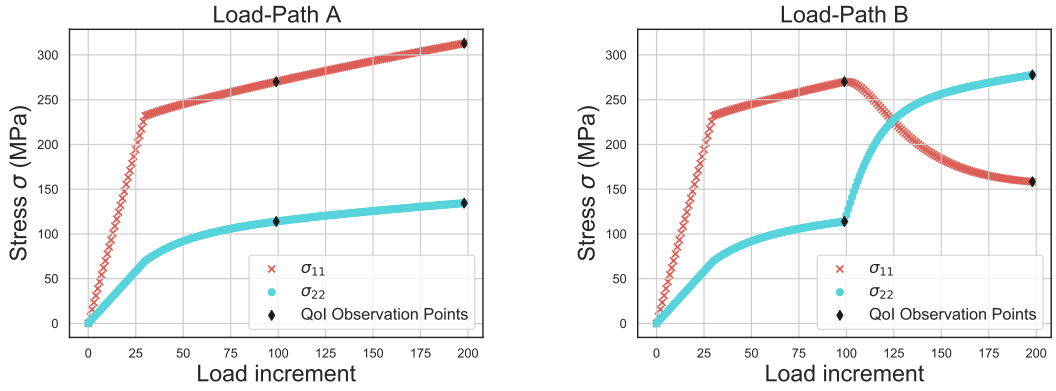

Figure 3: In-plane stresses for load-paths A and B.

that we can directly compare our Bayesian inference results to those obtained with the surrogates, which for most FE models is not a verification check that is practical to perform.

Surrogate Model Construction

Our means of dealing with the history dependence inherent in elastoplastic models is to remove it from consideration for all intents and purposes by building distinct surrogates for each node in the load-path graph. We utilize two approaches to generate build data (i.e. samples/points in the input space and corresponding function evaluations or QoIs in the output space) for the surrogates: a quasi-random sequence that produces a space-filling sample design and an adaptive sampling technique. For the remainder of the document, we limit the input space of the surrogates to the Hill coefficients F and G and the QoI to σ_{11} to facilitate visualization.

Halton samples [29] are generated using deterministic sequences of prime numbers and produce space-filling designs. They are a reasonable first strategy for the creation of surrogate build data. We created two distinct sets of samples for *training* and *test* data that each contain 200 points over the domain $[0.3, 0.7]^2$ using Halton sequences based on the primes $(2, 3)$ and $(5, 7)$, respectively.

Each F and G input pair was used to drive an execution of the material point simulator (other parameters specified according to the values in Table 1). There were a total of 100 increments between nodes with a strain increment $\Delta\epsilon = 0.005$, and the value of σ_{11} was stored at each observation layer. The resulting collections of input point-QoI pairs constitute the training and test

sets.

We considered Gaussian process (GP) and polynomial chaos expansion (PCE) surrogate models. GPs [21, 30, 31] use a kernel function to measure the similarity between points and have the notable benefit of providing both mean and variance predictions. In this study, we utilized an anisotropic squared exponential kernel. PCEs [23, 24, 32] are linear combinations of orthogonal polynomials that are chosen according to ther input random variables (RVs). In this case, the RVs correspond to the unknown Hill coefficients. We utilized an independent uniform RV for both parameters $\mathcal{U} \sim [0.3, 0.7]$ which results in Legendre polynomials for the PCE basis. The PCE basis and coefficients were determined according to the basis selection algorithm presented in [24]. All surrogates were built using PyApprox [33].

Adaptive sampling strategies are attractive because they tailor the design of the training set to reduce the approximation error in the surrogate so that in many cases they can produce surrogates with commensurate accuracy to space-filling designs with considerably fewer evaluations of the function the surrogate is meant to approximate. We will almost certainly need to employ an adaptive sampling approach when the function is a FE model to keep surrogate build costs manageable. In this work, we investigated the use of a recently-developed adaptive sampling strategy [34] based on the Cholesky decomposition of a GP's covariance matrix and compared the accuracy of surrogates built using this data to those built with a space-filling design.

Our primary motivation for developing surrogate models is to reduce the computational cost of Bayesian inference, which can easily require thousands to millions of function evaluations [30, 35]. We will now describe the inference problem and the names for the various probability distributions involved.

Bayesian Inference

Bayes' rule is derived from the definition of conditional probability:

$$\begin{aligned} P(\boldsymbol{z} \cap \boldsymbol{d}) &= P(\boldsymbol{z} \mid \boldsymbol{d})P(\boldsymbol{z}) = P(\boldsymbol{d} \mid \boldsymbol{z})P(\boldsymbol{d}), \\ P(\boldsymbol{z} \mid \boldsymbol{d}) &= \frac{P(\boldsymbol{d} \mid \boldsymbol{z})P(\boldsymbol{z})}{P(\boldsymbol{d})}. \end{aligned} \tag{5}$$

We want to infer the parameters $\boldsymbol{z}$ given experimental measurements of QoIs $\boldsymbol{d}$. The *prior distribution* $P(\boldsymbol{z})$ for parameters F and G is the product of two independent uniform random variables

$\mathcal{U}(0.3, 0.7)$. The probability distribution that quantifies how well our predictions match the data $P(\boldsymbol{d} \mid \boldsymbol{z})$ is called the *likelihood distribution*. The quantity in the denominator, $P(\boldsymbol{d})$, is called the *model evidence* or sometimes just the *evidence*. It is a scalar and normalizes the numerator so that the right hand side will integrate to 1 (a requirement of probability distributions). The distribution $P(\boldsymbol{z} \mid \boldsymbol{d})$ is known as the *posterior distribution*. It quantifies our knowledge about the model parameters given our prior beliefs and observed data.

When the noise in the data is assumed to be zero-mean, additive Gaussian noise, the likelihood distribution is

$$P(\boldsymbol{z} \mid \boldsymbol{d}) \propto \exp\left(-\frac{1}{2}(\boldsymbol{u}(\boldsymbol{z}) - \boldsymbol{d}) \cdot \Sigma^{-1}(\boldsymbol{u}(\boldsymbol{z}) - \boldsymbol{d})\right), \tag{6}$$

where $\boldsymbol{u}(\boldsymbol{z})$ is the *parameter-to-observable map* and Σ is the covariance matrix for the noise in the observations. When the noise is uncorrelated and homoscedastic, $\Sigma = \sigma_n^2 I$, where I denotes the identity matrix with dimension equal to the number of observations , and σ_n is the variance of the noise.

An approximation of the posterior can be directly computed from (5) and (6) when a cheap (and hopefully accurate) surrogate $\tilde{\boldsymbol{u}}(\boldsymbol{z})$ for $\boldsymbol{u}(\boldsymbol{z})$ is available and $P(\boldsymbol{z})$ is known. However, numerical integration of the denominator $P(\boldsymbol{d}) = \int_{\boldsymbol{z}} P(\boldsymbol{d}|\boldsymbol{z})P(\boldsymbol{z})\, d\boldsymbol{z}$ suffers from the *curse of dimensionality* with respect to the dimension of the parameter space. Consequently, direct evaluation is only practical when the parameter space is contains four or fewer dimensions [35].

Various techniques (e.g. Monte Carlo sampling, sparse grid quadrature) have been developed to make the process of Bayesian inference computationally tractable. Alternatively, Markov chain Monte Carlo (MCMC) methods can be used to produce samples from the posterior given the prior and likelihood. Density estimation techniques may be applied to those samples to produce a distribution for the posterior when one is desired, but they are also typically only feasible for low-dimensional problems. The reference [35] provides an introduction to all of these methods.

Traversing the load-path graph and performing calibration at each stage amounts to the need to solve a *sequential inference* problem. In sequential inference, the posterior for the previous inference problem becomes the prior for the current stage. We can write this problem for $i = 1, \ldots, N$ inference steps as

$$P^{i+1}(\boldsymbol{z}) = P(\boldsymbol{z} \mid \boldsymbol{d}^i) = P(\boldsymbol{d}^i \mid \boldsymbol{z}) P^i(\boldsymbol{z}), \tag{7}$$

where $\boldsymbol{d}^i$ is the calibration data and $P^i(\boldsymbol{z})$ is the prior for stage i, respectively.

We conclude this section with a discussion of the *Laplace approximation*. It is a Gaussian approximation of the posterior that can be obtained by solving a deterministic optimization problem and applying some post-processing. Assuming we have an *uninformative prior* and uncorrelated and homoscedastic noise, the posterior distribution π after N stages is

$$\pi(\boldsymbol{z}) \propto \exp\left(-\frac{1}{2\sigma_n^2} \sum_{i=1}^{N} \|\boldsymbol{u}^i(\boldsymbol{z}) - \boldsymbol{d}^i)\|^2 \right). \tag{8}$$

We now take the log of the posterior and multiply by -1 to obtain the *negative log-likelihood* function

$$J(\boldsymbol{z}) := -\log \pi(\boldsymbol{z}) = \frac{1}{2\sigma_n^2} \sum_{i=1}^{N} \|\boldsymbol{u}^i(\boldsymbol{z}) - \boldsymbol{d}^i)\|^2 + C. \tag{9}$$

The minimizer of (9) $\boldsymbol{z}^*$ is known as the *maximum likelihood estimate* (MLE). In the more general case when a prior is present, the analogous quantity is called the *maximum a posteriori* (MAP) point. Gradient-based optimization routines are often employed to minimize (9), which unsurprisingly require $\frac{dJ}{d\boldsymbol{z}}$ (or an approximation).

The post-processing step involves computing the Hessian of (9) at the MLE/MAP point and inverting it. One approach to computing the gradient and/or the Hessian is through a forward sensitivity analysis:

$$\begin{aligned} \frac{dJ}{dz_i} &= \frac{dJ}{du_k} \frac{du_k}{dz_i}, \\ \frac{d^2J}{dz_i dz_j} &= \frac{d^2J}{du_k du_l} \frac{du_k}{dz_i} \frac{du_l}{dz_j} + \frac{dJ}{du_k} \frac{d^2u_k}{dz_i dz_j}. \end{aligned} \tag{10}$$

The blue terms are known as the forward sensitivity matrices and some codes / surrogates can provide them. The red term (a rank-3 tensor) is less commonly available. The most straightforward approach to computing the Hessian H_{ij} is through a centered second difference approximation about $J(\boldsymbol{z}^*)$.

The Laplace approximation is the multivariate normal distribution $\mathcal{N}(\boldsymbol{z}^*, \boldsymbol{C})$, where the covariance is obtained by inverting the Hessian such that $C_{ij} = [\boldsymbol{H}^{-1}]_{ij}$.

Results and Discussion

We first present the errors on the test set for the two surrogates types and two sampling strategies we studied for a single QoI, namely σ_{11} for load-path B at observation point 1. Our presentation is limited to this single QoI for brevity. The results for the other QoIs displayed similar trends in accuracy.

The root-mean-squared error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE) are reported in Table 2. The definitions of these errors are given by the following expressions:

$$\begin{aligned} \text{MSE} &= \frac{1}{M}\sum_{i=1}^{N}\left(\tilde{\boldsymbol{u}}(\boldsymbol{z}_i) - \boldsymbol{u}(\boldsymbol{z}_i)\right)^2, \\ \text{MAE} &= \frac{1}{M}\sum_{i=1}^{N}|\tilde{\boldsymbol{u}}(\boldsymbol{z}_i) - \boldsymbol{u}(\boldsymbol{z}_i)|, \\ \text{MAPE} &= \frac{100}{M}\sum_{i=1}^{N}\left|\frac{\tilde{\boldsymbol{u}}(\boldsymbol{z}_i) - \boldsymbol{u}(\boldsymbol{z}_i)}{\boldsymbol{u}(\boldsymbol{z}_i)}\right|, \end{aligned} \tag{11}$$

where M is the number of samples in the test set.

The space-filling (Halton with 200 points) and adaptive (Cholesky with 20 points) training sets are shown in the top row of Figure 4. The number of points chosen was somewhat arbitrary for this simplified problem using the material point simulator. In the target application, the cost of the FE simulations will likely restrict the number of sampling points that can feasibly be accommodated; balancing the computational cost to generate build data and the surrogate error will be studied in the follow-on LDRD project. We note that the adaptive sampling approach tends to place points closer to the boundary. The bottom row of Figure 4 contains plots of the predictions from PCE

	Halton		Adaptive	
Error	GP	PCE	GP	PCE
RMSE (MPa)	1.5×10^{-4}	1.6×10^{-5}	1.0×10^{-2}	7.8×10^{-3}
MAE (MPa)	2.9×10^{-5}	6.8×10^{-6}	8.0×10^{-3}	6.3×10^{-3}
MAPE (%)	1.6×10^{-5}	4.4×10^{-6}	4.8×10^{-3}	3.8×10^{-3}

Table 2: Error metrics for the Halton test set from surrogates models for load-path B σ_{11} at Node 1 built using Halton (i.e. space-filling) and adaptive sampling schemes.

and GP surrogate models built using space-filling samples. The predictions from the adaptive sampling-based surrogates are not shown because they are visually indistinguishable from space-filling results.

We observe that all of the error metrics are lower for the surrogates built using the space-filling training set when compared to the surrogates built using the adaptive sampling scheme. The MAPE for the both of the adaptive sampling-based surrogates, however, is below 0.01 %, which in many scenarios could be an acceptable amount of approximation error. In addition, the adaptive approach required only a tenth of the number of function evaluations compared to the generation of the space-filling samples.

We next present inference results obtained from direct evaluation of Bayes' rule. As noted in the previous section, direct inference is only possible in low-dimensional parameter spaces due to the curse of dimensionality. However, the parameter space for the model problem presented in this work is two-dimensional, so direct inference is a viable approach. We discretized the parameter space into a regular grid with 101 points along each coordinate direction, evaluated the likelihood at each point, and used quadrature to compute the model evidence.

Figures 5 and 6 contain the results obtained from direct inference for synthetic calibration data from hypothetical experiments with differing amounts of noise σ_n^2 that utilize load-paths A and B, respectively, with true parameters $F = 0.65$ and $G = 0.44$ (all other parameters were set to the values in Table 1). The plots in the left column for inference from Node 1 are similar for both figures, as the only difference between them is the realization of random noise in the measurement at that stage. The differences in the right column (corresponding to Node 3 in Figure 5 and Node 4 in Figure 6), however, are easily noticed.

The inference made at observation layer 1 for load-path A (Node 3) looks much like that obtained from observation layer 0 (Node 1), which is not unexpected as the simulation has simply continued along the same strain path and thus has gained a minimal amount of new information

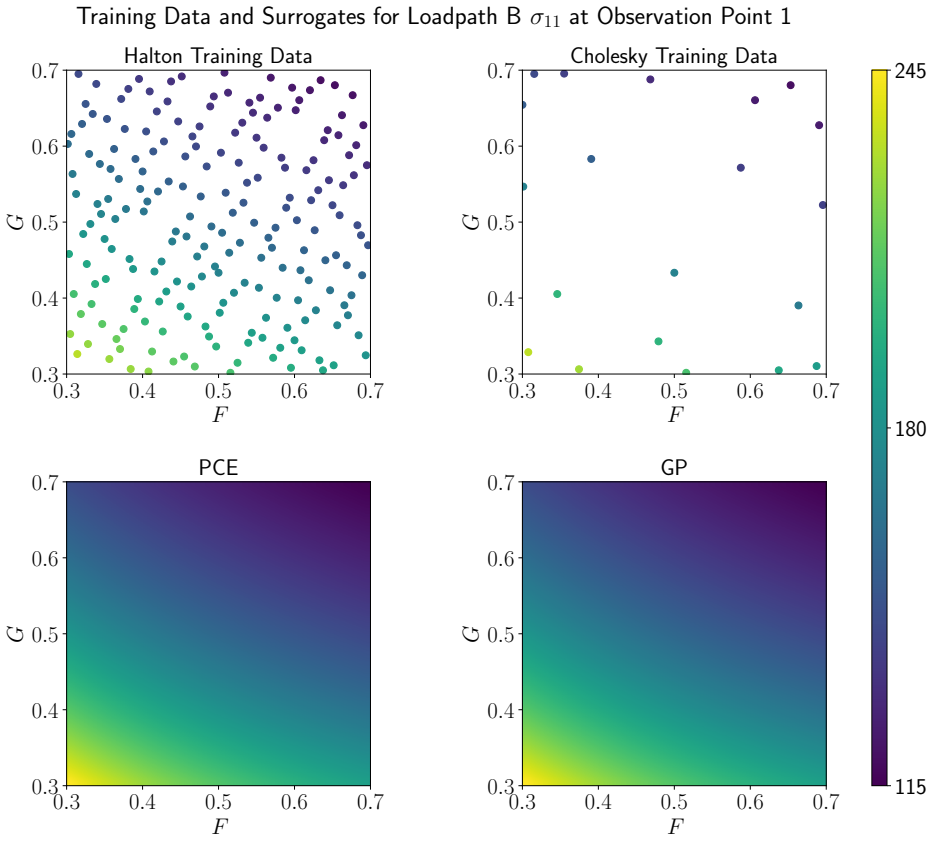

Figure 4: Top row: Training data sets for the surrogates studied in this paper. Bottom row: Predictions of the surrogates over the input space for the space-filling design. The QoI σ_{11} has units of MPa, and the input parameters F and G are unitless.

for calibration. On the other hand, load-path B imparts a strain in a different direction which causes the observation made at layer 1 (Node 4) to be more informative for calibration compared to observation layer 0 (Node 1), as evidenced by the "shrinking" of the posterior around a region in parameter space. Therefore, it would be advantageous to choose load-path B, as the parameter distributions it produces change the most between observation layers.

We also used quadrature to evaluate the posterior means obtained using the material point sim-

ulator and a PCE surrogate for the load-path B example with $\sigma_n^2 = 1.0$, results of which are shown in Table 3. The means for these posteriors agreed to at least three significant figures for the posteriors from observation layers 0 and 1, but the difference in computational cost was over five orders of magnitude in the surrogates' favor. The cost savings will be even greater when we replace the material point simulator with a FE model, but it may be a challenge to produce sufficiently accurate surrogates for our chosen QoIs given a higher-dimensional input space and more severely constrained budget for training data.

	Mean [F,G] Values	Computational Time (sec)
Material Point Simulator	[0.6324, 0.4445]	1.0×10^3
Surrogate Network	[0.6322, 0.4445]	6.5×10^{-3}

Table 3: Comparison of mean values of the posteriors and computational cost when using Bayesian inference with either the material point simulator or the surrogate network, for load-path B, observation layer 1 (Node 4), with QoI observation noise variance $\sigma_n^2 = 1.0$ MPa.

We note that direct integration quickly becomes intractable as the dimension of the parameter space grows. In higher dimensions, MCMC sampling is a means of producing samples from the posterior. We leave a deep investgation of such sampling algorithms and their use in step selection through Bayesian OED to future work.

Our final inference result concerns the Laplace approximation. Figure 7 shows results for load-path B inference when Laplace's approach (computed via finite differences) is used to generate an approximation of the posterior. Visual comparison with Figure 6 shows that while differences in the posteriors are apparent, the Laplace approximation does produce a close emulation of the true posterior for this problem. The cost of obtaining the Laplace approximation is considerably smaller than that of direct inference or MCMC sampling, and therefore these results suggest that we should keep this method in mind when we move to the use of costly FE models for the test specimen.

We have shown that we can produce reasonably accurate surrogate models for a simplified model of a characterization experiment as implemented through a material point simulator. Our work may be used as a template for constructing surrogates for QoIs extracted from FE models of the test specimen, which is our ultimate plan. Furthermore, our inference results demonstrate our ability to perform direct and approximate inference through evaluation of Bayes' rule and the Laplace approximation, a necessary step for obtaining the parameter distributions used to actively guide the experiment.

As a final note, we have deferred an in-depth exploration of the step selection problem to future work. However, we did begin to evaluate one possible strategy for choosing a branch given a current parameter distribution that involves quantifying the variance in QoIs for hypothetical observations obtained by "pushing" samples from the current posterior through the surrogates at the next observation layer. Our hypothesis is that the path with the largest variance would provide the most informative data for calibration and thus should be selected.

Figure 8 contains histograms generated by pushing samples from the posterior distribution obtained from inference at Node 1 in Figure 2 through the surrogates for Nodes 3 and 4. In this simplified example, we compute the standard unbiased sample variance. Load-path B appears to be more attractive according to our "greater variance" selection rule. We plan to investigate the utility of scalar measures of the sample covariance matrix for vectors of QoIs (e.g. trace, determinant) in the future.

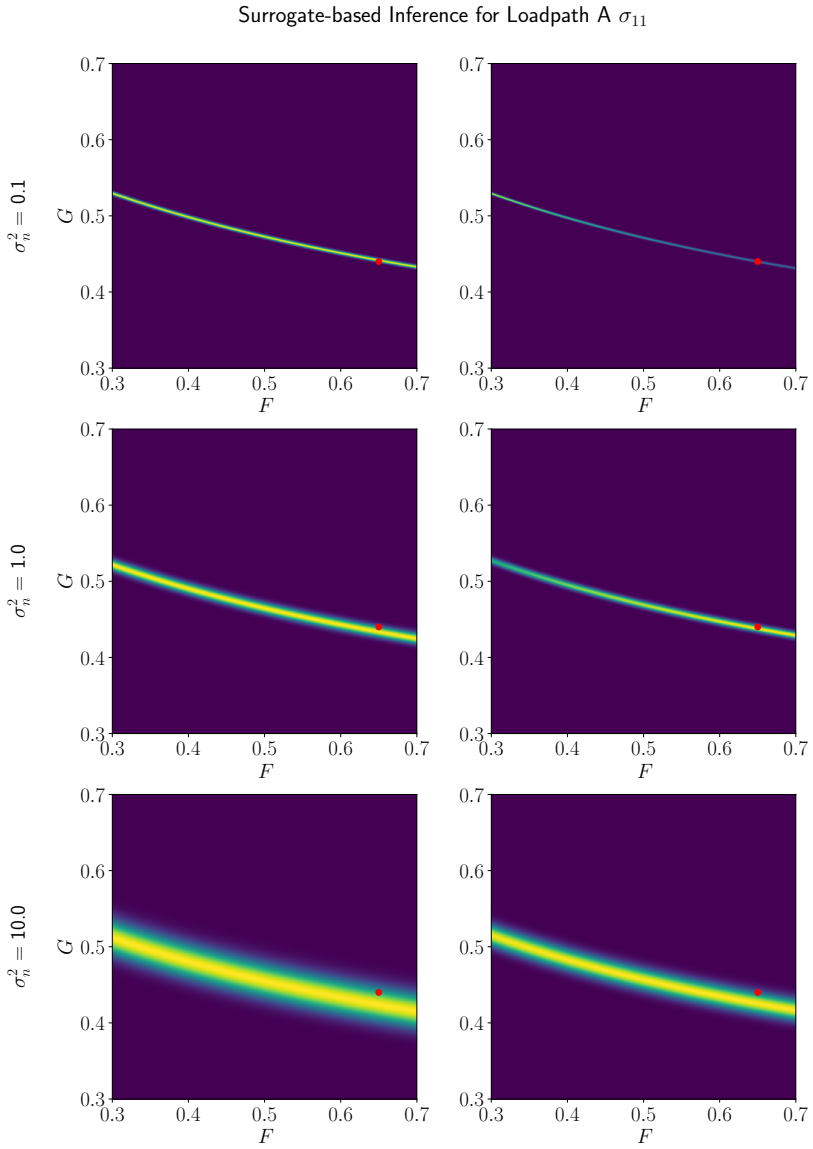

Figure 5: Posteriors obtained using direct inference for load-path A. Rows denote the amount of observation noise σ_n^2, and from left to right the columns represent inference at observation layers 0 and 1 (Nodes 1 and 3, respectively). The red dot is the true value of the parameters.

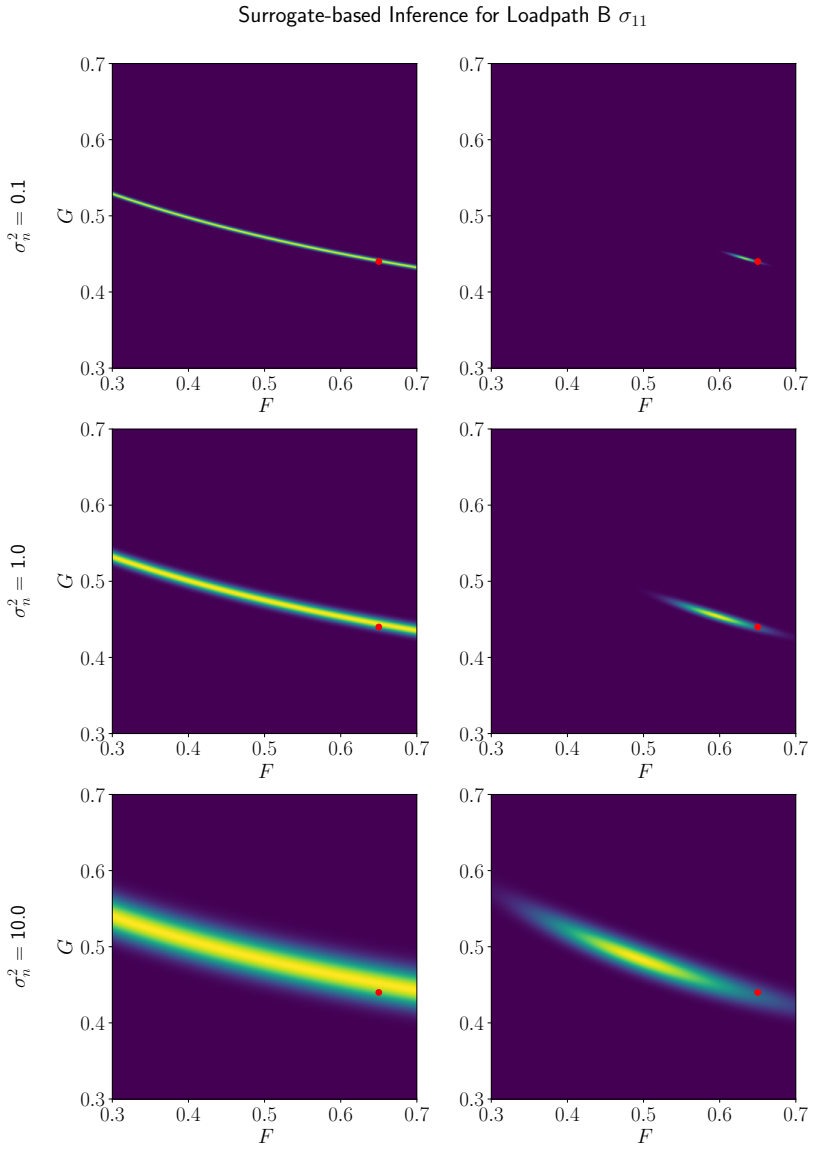

Figure 6: Posteriors obtained using direct inference for load-path B. Rows denote the amount of observation noise σ_n^2, and from left to right the columns represent inference at observation layers 0 and 1 (Nodes 1 and 4, respectively). The red dot is the true value of the parameters.

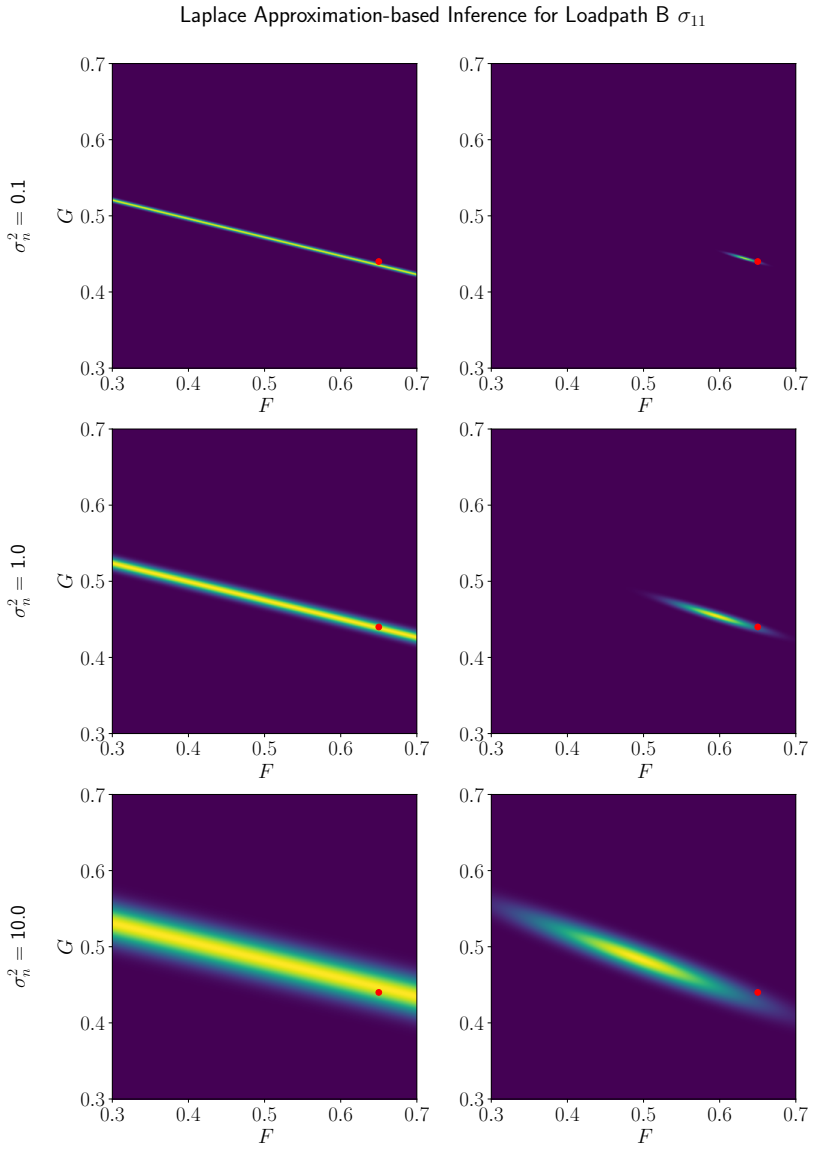

Figure 7: Laplace approximations of the posteriors for load-path B. Rows denote the amount of observation noise σ_n^2, and from left to right the columns represent inference at observation layers 0 and 1 (Nodes 1 and 4, respectively). The red dot is the true value of the parameters.

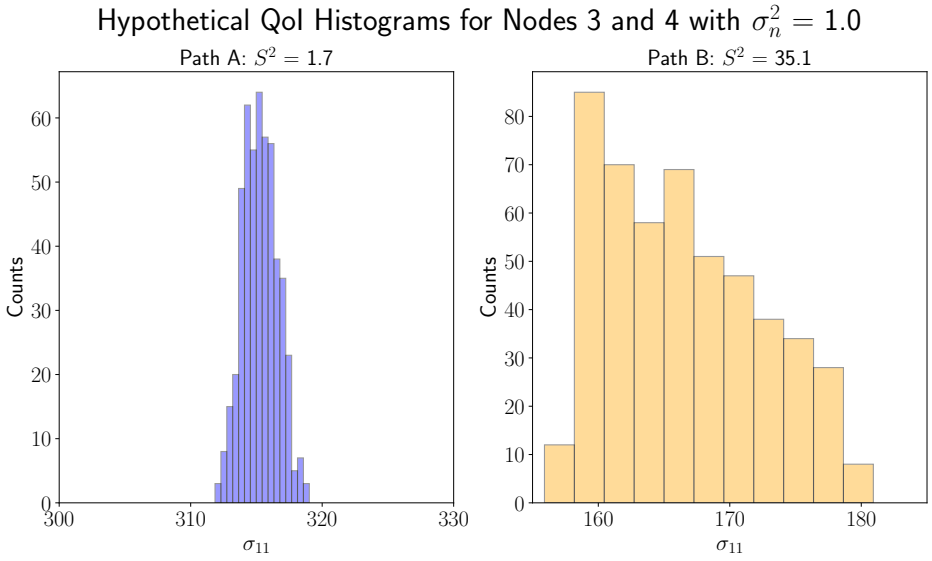

Figure 8: Histograms of hypothetical observations of σ_{11} obtained by sampling the posterior obtained from inference at observation layer 0 (Node 1) and evaluating those samples at the surrogates for the downstream nodes (Nodes 3 (left) and 4 (right)). The histogram that corresponds to choosing path B has a larger sample variance, which may mean that path is more desirable.

Anticipated Outcomes and Impacts

Bayesian inference provides uncertainty quantification on model parameters, as opposed to a deterministic calibration method such as FEMU, which only provides a single parameter set. However, FEA is too computationally expensive for Bayesian inference, which motivates the development of surrogate models. A challenge with surrogate models, though, is accounting for the history-dependence present in elastoplastic models. The key technical accomplishment of this exploratory LDRD was the successful incorporation of history/path dependency through a network of surrogate models via a finite load-path tree. Additional tools and capabilities developed include a framework for surrogate model construction based on the material point simulator and a workflow for Bayesian inference based on the surrogate network.

This exploratory project allowed a team of 3, who have been staff members for less than 5 years, to further develop their research careers at Sandia. Additionally, we have partnered with the University of Illinois at Urbana-Champaign, working with Professor John Lambros and graduate student Mr. Samuel Fayad, to investigate the sensitivity of the calibration results and load-step selection metrics to noise in DIC QoIs. This university work was funded through Advanced Scientific Computing (ASC) Physics and Engineering Models (P&EM), but complemented this exploratory LDRD project.

This exploratory LDRD project bought down risk regarding the surrogate network and facilitated a successful FY22-24 full ESRF LDRD proposal to research the complete ICC paradigm. The follow-on project will coordinate with and leverage an Integrated ASC project, “Calibrating Credible Material Models”. If the follow-on LDRD project is successful, the main immediate outcome will be a technical understanding on the major components of the ICC paradigm. Our long-term vision is that follow-on DE/ASC funding could transition the complete ICC technology to MatCal so it can be applied to multi-axial frames at both SNL-NM and SNL-CA. Such a tool could immediately be used for efficient and accurate calibrations of metals for the nuclear weapons community, since we are planning on using an anisotropic aluminum as an exemplar material and employing the modular plasticity framework developed at Sandia [36–38] to develop ICC paradigm. Looking further out, the results from the full LDRD project will serve as the basis for expanding ICC for other phenomenologies (e.g. rate and temperature dependence) and adapting ICC for other material classes.

Through the follow-on LDRD project, we seek to transform the state of practice of material characterization and model calibration through the development of the ICC paradigm. This project is well-aligned with the thrust areas of the Engineering Science Research Foundation. We will

develop new "Diagnostics and Experimental Capabilities for Engineering Applications" by interlacing experiments and calibration in a real-time feedback loop to improve both processes and their outcomes. Our proposed decomposition of field data will facilitate "Engineering Data Integration and Interrogation", by streamlining the fusion of experimental data and simulations. The ICC paradigm will create a new "Next Generation Engineering Simulation" framework, by producing credible material models with quantified parameter uncertainty.

The net result is being able to more expeditiously perform material model calibration, which is aligned with several of the Labs-Level Strategic Priorities. In addition to potential cost savings, the ICC paradigm will reduce lead times for calibration, allowing modeling and simulation to be employed earlier in the design cycle. This both delivers on current commitments (Priority 1) while creating a more flexible and responsive modeling capability (Priority 2). Through these processes, relevant subject matter expertise may be more readily deployed to a variety of projects (Priority 6). The combined experimental and modeling approaches proposed here seek to leverage the advanced capabilities of Sandia in both capacities, thereby "unleashing the power of Sandia" (Priority 7).

Conclusion

In this work, we have demonstrated our ability to build a surrogate network based on the finite load-path tree, and to use the network for Bayesian inference of our material model parameters. As an exemplar, our reference simulation was for an elastoplastic material point subject to biaxial deformation under plane stress conditions. The material point simulator employed the Hill anisotropic yield criterion with an isotropic Voce hardening function. The surrogate models took two parameters of the yield criterion (F and G) as inputs and produced two in-plane stresses and one out-of-plane strain as outputs.

Training data for the surrogates was generated using either a space-filling (Halton) or adaptive sampling (Cholesky) method, and surrogates were built using either Gaussian process (GP) or polynomial chaos expansion (PCE) methods. Using an independent set of test data, the error of the surrogates was evaluated compared to the reference material point simulator. With space-filling training data, the mean average percent error was on the order of 10^{-5} and 10^{-6} for GP and PCE surrogates, respectively, while the error was 10^{-3} for both surrogate methods with adaptive sampling. The acceptable balance of computational efficiency and error for the ICC paradigm is an open question, but adaptive sampling methods are a promising avenue to explore.

Given the surrogates, direct Bayesian inference was performed using both the material point simulator and the surrogate network. The means for the posteriors agreed to at least three sig-

nificant figures, but the difference in computational cost was over five orders of magnitude in the surrogates' favor. We also explored using the Laplace approximation to estimate the posterior and obtained a close emulation of the true posterior obtained with direct Bayesian inference with considerably smaller computational cost.

In summary, these results bought down risk regarding the surrogate network and facilitated a successful FY22-24 full LDRD proposal. In this follow-on work, we will research and develop the complete ICC paradigm. If successful, we will create a new calibration paradigm that enables next generation engineering simulation via rapid and simultaneous calibration of multiple material models with comprehensive material characterization that includes quantified uncertainty on model parameters. The methods and tools created will directly contribute to reducing design cycle time and lead to more credible simulations for use in high-consequence engineering decisions.

Acknowledgments

This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

[1] E. Corona, S. L. B. Kramer, W.M. Scherzinger, and A.R. Jones. Anisotropic plasticity model forms for extruded Al 7079: Part I, Calibration. *Int. J. Solids Struct.*, 213:135–147, 2021.

[2] K. N. Karlson, A.A. Hanson, and B.R. Antoun. Revisiting the pipe bomb V&V problem using the BCJ material model. *SAND Report*, SAND2018-13911, 2018.

[3] K. N. Karlson and G. Bergel. Hierarchical multi-scale modeling of laser weld failure: An initial investigation. *SAND Report*, SAND2019-15067R, 2019.

[4] E. Corona, S. L. B. Kramer, B.T. Lester, A.R. Jones, B. Sanborn, L. Shand, and C.J. Fietek. Thermal-mechanical elastic-plastic and ductile failure model calibrations for 304L stainless steel alloy. *SAND Report*, SAND2021-1752, 2021.

[5] E. Corona, S. L. B. Kramer, B.T. Lester, A.R. Jones, B. Sanborn, and C.J. Fietek. Thermal-mechanical elastic-plastic and ductile failure model calibrations for 6061-T651 aluminum alloy from plate. *SAND Report*, SAND2021-2983, 2021.

[6] E. Corona and G.E. Orient. An evaluation of the Johnson-Cook model to simulate puncture of 7075 aluminum plates. *SAND Report*, SAND2014-1550, 2014.

[7] E.M.C. Jones, E. Corona, A.R. Jones, W.M. Scherzinger, and S. L. B. Kramer. Anisotropic plasticity model forms for extruded Al 7079: Part II, Validation. *Int. J. Solids Struct.*, 213:148–166, 2021.

[8] M.A. Sutton, J.-J. Orteu, and H.W. Schreier. *Image Correlation for Shape, Motion, and Deformation Measurements*. Springer, New York, NY 10013, USA, 2009.

[9] E.M.C. Jones, J.D. Carroll, K.N. Karlson, S.L.B. Kramer, R.B. Lehoucq, P.L. Reu, D.T. Seidl, and D.Z. Turner. High-throughput material characterization using the virtual fields method. *SAND Report*, SAND2018-10635, 2018.

[10] E.M.C. Jones, K.N. Karlson, and P.L. Reu. Investigation of assumptions and approximations in the virtual fields method for a viscoplastic material model. *Strain*, 55(4), 2019.

[11] E. M. C. Jones, J. D. Carroll, K. N. Karlson, S. L. B. Kramer, R. B. Lehoucq, P. L. Reu, and D. Z. Turner. Parameter covariance and non-uniqueness in material model calibration using the virtual fields method. *Comp. Mater. Sci.*, 152:268–290, 2018.

[12] S. L. B. Kramer and W.M. Scherzinger. Implementation and evaluation of the virtual fields method: Determining constitutive model parameters from full-field deformation data. *SAND Report*, SAND2014-17871, 2014.

[13] S. Avril, M. Bonnet, A.-S. Bretelle, M. Grédiac, F. Hild, P. Ienny, F. Latourte, D. Lemosse, S. Pagano, E. Pagnacco, and F. Pierron. Overview of identification methods of mechanical parameters based on full-field measurements. *Exp. Mech.*, 48(4):381–402, 2008.

[14] F. Pierron and M. Grédiac. *The Virtual Fields Method: Extracting Constitutive Mechanical Parameters from Full-Filed Deformation Measurements*. Springer, 2012.

[15] M.I. Friswell and J.E. Mottershead. *Finite element model updating in structural dynamics*, volume 38 of *Solid Mechanics and Its Applications*. Kluwer Academic Publishers Group, 1995.

[16] M. Grédiac, F. Pierron, S. Avril, and E. Toussaint. The virtual fields method for extracting constitutive parameters from full-field measurements: A review. *Strain*, 42(4), 2008.

[17] Maher Salloum, Kyle N. Karlson, Helena Jin, Judith A. Brown, Dan S. Bolintineanu, and Kevin N. Long. Comparing field data using alpert multi-wavelets. *Computational Mechanics*, 66(4):893–910, 2020.

[18] M. Salloum, K.L. Johnson, J.E. Bishop, J.M. Aytac, D. Dagel, and B.G. van Bloemen Waanders. Adaptive wavelet compression of large additive manufacturing experimental and simulation datasets. *Comp. Mech.*, 63(3):491–510, 2018.

[19] M.W. Kury and K.N. Karlson. Matcal: A tool for improving the traceability and workflow for material calibration. *SAND Report*, SAND2019-11952C, 2019.

[20] R. Hill. A theory of the yielding and plastic flow of anisotropic metals. *P. Roy. Soc. Lond. A Mat.*, 193(1033):281–297, 1948.

[21] C.K.I. Williams and C.E. Rasmussen. *Gaussian Processes for Machine Learning*, volume 2. MIT Press Cambridge, MA, 2006.

[22] R.B. Gramacy. *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. CRC Press, Taylor and Francis Group, Boca Raton, Florida, USA, 2020.

[23] D. Xiu and G.E. Karniadakis. The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM J. Sci. Comput.*, 24(2):619–644, 2002.

[24] J.D. Jakeman, M.S. Eldred, and K. Sargsyan. Enhancing $\ell 1$-minimization estimates of polynomial chaos expansions using basis selection. *J. Comput. Phys.*, 289:18–34, 2015.

[25] H. Helmut, J.D. Jakeman, and P. Zaspel. Weighted greedy-optimal design of computer experiments for kernel-based and Gaussian Process model emulation and calibration. *SAND report*, SAND2020-3677R, 2020.

[26] J.D. Jakeman, F. Franzelin, A. Narayan, M. Eldred, and D. Plfuger. Polynomial chaos expansions for dependent random variables. *Comput. Method Appl. M.*, 351:643–666, 2019.

[27] A. Narayan and J.D. Jakeman. Adaptive Leja sparse grid contructions for stochastic collocation and high-dimensional approximation. *SIAM J. Sci. Comput.*, 36(6):A2952–A2983, 2014.

[28] J.C. Simo and T.J.R. Hughes. *Computational Inelasticity*, volume 7. Springer Science & Business Media, 2006.

[29] J.H Halton. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.*, 2(1):84–90, 1960.

[30] K.P. Murphy. *Machine learning: A probabilistic perspective*. MIT press, 2012.

[31] C.M. Bishop. *Pattern recognition and machine learning*. springer, 2006.

[32] R.G. Ghanem and P.D. Spanos. *Stochastic finite elements: A spectral approach*. Courier Corporation, 2003.

[33] J.D. Jakeman. *PyApprox: Probabilistic analysis and approximation of data and simulation*, 2021.

[34] H. Harbrecht, John D. Jakeman, and Peter Zaspel. Cholesky-based experimental design for gaussian process and kernel-based emulation and calibration. *Communications in Computational Physics*, 29(4):1152–1185, 2021.

[35] R.C. Smith. *Uncertainty quantification: Theory, implementation, and applications*, volume 12. SIAM, 2013.

[36] LAMÉ Team. Library of advanced materials for engineering (LAMÉ). *SAND Report*, SAND2020-3408, 2020.

[37] B.T. Lester and W.M. Scherzinger. Adiabatic heating in modular plasticity models. *SAND Report*, SAND2019-15003R, 2019.

[38] B.T. Lester and W.M. Scherzinger. Modular plane stress plasticity material model. *SAND Report*, SAND2019-3151R, 2019.