

Plasticity models of material variability based on uncertainty quantification techniques

F. Rizzi

*Scalable Modeling and Analysis Department, Sandia National Laboratories, P.O. Box 969,
Livermore, CA 94551, USA*

R.E. Jones*

*Mechanics of Materials Department, Sandia National Laboratories, P.O. Box 969,
Livermore, CA 94551, USA*

J.A. Templeton

*Thermal/Fluid Science and Engineering Department, Sandia National Laboratories, P.O.
Box 969, Livermore, CA 94551, USA*

J.T. Ostien

*Mechanics of Materials Department, Sandia National Laboratories, P.O. Box 969,
Livermore, CA 94551, USA*

B.L. Boyce

*Materials Mechanics and Tribology Department, Sandia National Laboratories, P.O. Box
5800, Albuquerque, NM 87185, USA*

*Corresponding author: rjones@sandia.gov

Plasticity models of material variability based on uncertainty quantification techniques

F. Rizzi

*Scalable Modeling and Analysis Department, Sandia National Laboratories, P.O. Box 969,
Livermore, CA 94551, USA*

R.E. Jones*

*Mechanics of Materials Department, Sandia National Laboratories, P.O. Box 969,
Livermore, CA 94551, USA*

J.A. Templeton

*Thermal/Fluid Science and Engineering Department, Sandia National Laboratories, P.O.
Box 969, Livermore, CA 94551, USA*

J.T. Ostien

*Mechanics of Materials Department, Sandia National Laboratories, P.O. Box 969,
Livermore, CA 94551, USA*

B.L. Boyce

*Materials Mechanics and Tribology Department, Sandia National Laboratories, P.O. Box
5800, Albuquerque, NM 87185, USA*

Abstract

The advent of fabrication techniques like additive manufacturing has focused attention on the considerable variability of material response due to defects and other micro-structural aspects. This variability motivates the development of an enhanced design methodology that incorporates inherent material variability to provide robust predictions of performance. In this work, we develop plasticity models capable of representing the distribution of mechanical responses observed in experiments using traditional plasticity models of the mean response and recently developed uncertainty quantification (UQ) techniques. We demonstrate that the new method provides predictive realizations that are superior to more

*Corresponding author: rjones@sandia.gov

traditional ones, and how these UQ techniques can be used in model selection and assessing the quality of calibrated physical parameters.

Keywords: Plasticity, additive manufacturing, material variability, uncertainty quantification.

1. Introduction

Variability of material response due to defects and other micro-structural aspects has been well-known for some time [1, 2, 3, 4]. In many engineering applications inherent material variability has been ignorable, and traditionally the design process is based on the mean or lower-bound response of the chosen materials. Material failure is a notable exception since it is particularly sensitive to outliers in the distributions of micro-structural features [5, 6, 7].

Currently, additive manufacturing (AM) is of particular technological interest and provides strong motivation to not only model the mean response of materials but also their intrinsic variability. Additive manufacturing has the distinct advantages of being able to fabricate complex geometries and accelerate the design-build-test cycle through rapid prototyping [8]; however, currently, fabrication with this technique suffers from variability in mechanical response due to various sources, including defects imbued by the process, the formation of residual stresses, and geometric variation in the printed parts. As an example, high throughput tensile data from Boyce *et al.* [9] clearly shows pronounced variability in the resultant yield and hardening.

Given this current state of the technology, the need to enhance design methodology to account for this variability in order to meet performance thresholds with high confidence is clear. In this work, we leverage tools from uncertainty quantification (UQ) [10, 11, 12] to provide material variability models, realizations, and, ultimately, robust performance predictions.

It is well-known that any model is an approximation of the physical response of a real system. Typically, models are characterized by many parameters, and thus appropriately tuning them becomes a key step toward reliable predictions.

The most common approach to model calibration is least-squares regression which yields a deterministic result appropriate for design to the mean. Bayesian inference methods provide a more general framework for model calibration and parameter estimation by providing a robust framework for handling multiple
30 sources of calibration information as well as a full joint probability density on the target parameters. Traditionally, Bayesian techniques have been applied in conjunction with additive noise models that are appropriate for modeling external, uncorrelated influences on observed responses. Recently, a technique to embed the modeled stochasticity in distributions on the physical parameters
35 of the model itself was developed by Sargsyan, Najm, and Ghanem [13], and in this work we adapt it to model the inherent variability of an AM metal [9]. This is not the only method available in this emerging field of probabilistic modeling of physical processes for engineering applications. There are commonalities between many of the methods. Notably, the work of Emery *et al.* [7] applied
40 the stochastic reduced order model (SROM) technique [14] to weld failure. The SROM technique has many of the basic components of embedded noise model: a surrogate model of the response to physical parameters, a means of propagating distributions of parameters with Monte Carlo (MC) sampling and computing realistic realizations of the predicted response.

45 In Sec. 2 of this work, we describe the selected experimental dataset [9] that motivates this effort and provides calibration data. This deep dataset provides real-world relevance that a synthetic dataset would not; however, we apply some pre-processing and simplifying assumptions to facilitate the task of developing the methodology. In Sec. 3, we review the basic plasticity theory that provides
50 the basis for the material variability models developed in Sec. 4. In Sec. 4, we develop the methods necessary to perform Bayesian calibration of the material parameters: selection of prior distributions to represent the state of knowledge prior to calibration, design of the likelihood function that determines how close the model response is to the calibration data, and the Markov chain Monte
55 Carlo sampling needed to evaluate the posterior distribution of the parameters that quantifies their means and uncertainties. In particular, we adapt both the

traditional additive error [15] and the newer embedded error [13] UQ methods to the representation of the observed mechanical response; and we develop surrogate models of the full finite element simulation tailored to the elastic-plastic response of interest to facilitate efficient Monte Carlo sampling. In Sec. 5, we provide the results of the surrogate response building and calibration processes in order to compare the two methods in light of the selected data. We also employ sensitivities provided by the surrogate in order to discuss model selection, and make assessments about the importance of the various parameters. In Sec. 6, we discuss the results in light of a simple analytic version of the representation problem that serves to illustrate the flow of the calibration process and emphasizes the attributes that make the embedded noise model particularly suitable to representing inherent material variability. We also describe how the variability models can be used in an enhanced design process. In Sec. 7, we emphasize the innovations of the proposed approach to modeling the mechanical response to microstructural material variability.

2. Experimental Data

We focus this work on the analysis of high-throughput, micro-tension experimental measurements of additively manufactured stainless steel. From the experiments of Boyce *et al.* [9], we have six experimental datasets, each consisting of 120 stress-strain curves from the array of nominally identical dogbone-shaped specimens shown in Fig. 1(a). (The data from distinct builds of the array are referred to as *batches* throughout the remainder of the manuscript.) Each stress-strain curve Fig. 1(b) is qualitatively similar and behaves in a classically elastic-plastic fashion; however, the material displays a range of yield strengths, hardening and failure strengths and some variability in its elastic properties.

To simplify the data and remove some of the uncertainties associated more with the loading apparatus than the material, we omit the pre-load cycle to approximately 0.2% strain. The remainder of the mechanical response is monotonic tensile loading at a constant strain rate, see Fig. 1(c). We associate zero

strain reference configuration with the zero-stress, mildly worked material resulting from the pre-load cycle. The resulting stress, σ , and strain, ε , values are derived from the customary engineering stress and strain formulas. We assume the measurement noise to be Gaussian with $\pm 0.009\%$ standard deviation in the strain measurement and ± 20.0 MPa in the stress measurement based on the analysis of the random variations for individual curves and the noise in their zero-stress/zero-strain intercepts. Since we do not try to model failure in this effort, we discard tests that do not reach at least 3% strain. This threshold was chosen to be sufficiently large that each sample curve is well within the plastic regime (and near peak stress), and yet retain sufficient data to enable calibration. This preprocessing yielded $N_b = 6$ batches of stress data with $N_i = \{64, 77, 91, 79, 64, 46\}$ curves, respectively. To make the data suitable for the inverse problem of parameter calibration, we interpolate each curve and extract $n_\varepsilon = 151$ points over the interval (0,3)% to finally arrive at

$$\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^{N_b}, \quad \text{with} \quad \mathcal{D}_i = \{\mathcal{D}_i^{(k)}\}_{k=1}^{N_i} \quad \text{and} \quad \mathcal{D}_i^{(k)} = \{\sigma_j^{(i,k)}\}_{j=0}^{n_\varepsilon-1}, \quad (1)$$

where i enumerates the batches, k enumerates the N_i stress curves within the i -th batch, and $\sigma_j^{(i,k)} = \sigma^{(i,k)}(\varepsilon_j)$ represents the stress measured at the j -th strain value, $\varepsilon_j = 0.03j/(n_\varepsilon - 1)$, for the k -th curve of the i -th batch. The resulting dataset is shown in Fig. 1(c).
85

To expedite development of the appropriate analysis and modeling of materials with significant intrinsic variations, we assume all variability beyond the nearly negligible measurement noise stems from the underlying material response. This will lead to conservative estimates of material variability; however, 90 given relevant data, variations in the as-built geometry could be included in the variability analysis or corrected for in pre-processing of the stress-strain data.

3. Plasticity Theory

To model the observed behavior which resembles standard von Mises plastic response, we adopt a standard finite deformation framework [16] with a multiplicative decomposition of the deformation gradient into elastic and plastic

parts

$$\mathbf{F} = \mathbf{F}_e \mathbf{F}_p , \quad (2)$$

where \mathbf{F}_e is associated with lattice stretching and rotation, and \mathbf{F}_p is associated with plastic flow. Following Ref. [17], we assume an additive stored energy potential written in terms of the elastic deformation

$$W = \frac{\kappa}{2} \left(\frac{1}{2} (J_e^2 - 1) - \log(J_e) \right) + \frac{\mu}{2} (\text{tr}[\bar{\mathbf{b}}_e] - 3) . \quad (3)$$

Here, the elastic volumetric deformation is given by $J_e = \det(\mathbf{F}_e) = \det(\mathbf{F})$ since plastic flow is assumed to be isochoric, and the deviatoric elastic deformation is measured by $\bar{\mathbf{b}}_e = J_e^{-2/3} \mathbf{F}_e \mathbf{F}_e^T$. We associate the elastic constants κ and μ with the bulk modulus and shear modulus, respectively, and relate them to Young's modulus, E , and Poisson's ratio, ν , via the linear elastic relations $\kappa = \frac{E}{3(1-2\nu)}$ and $\mu = \frac{E}{2(1+\nu)}$. The Kirchhoff stress resulting from the derivative of the stored energy potential, W , is

$$\boldsymbol{\tau} = \frac{\kappa}{2} (J_e^2 - 1) \mathbf{I} + \mathbf{s} \quad \text{with} \quad \mathbf{s} = \mu \text{dev}[\bar{\mathbf{b}}_e] . \quad (4)$$

For the inelastic response, we employ a J_2 (von Mises) yield condition between an effective stress derived from $\mathbf{s} = \text{dev}[\boldsymbol{\tau}]$ and an associated flow stress, Υ , as

$$f := \sqrt{\frac{2}{3}} \|\mathbf{s}\| - \Upsilon \leq 0 . \quad (5)$$

The rate independent, associative flow rule is written in the current configuration as the Lie derivative of the elastic left Cauchy-Green tensor (*cf.* Ref. [17])

$$L_v \mathbf{b}_e = -\frac{2}{3} \gamma \text{tr}[\mathbf{b}_e] \frac{\mathbf{s}}{\|\mathbf{s}\|} . \quad (6)$$

The Lagrange multiplier γ enforces consistency of the plastic flow with the yield surface, obeys the usual Kuhn-Tucker conditions, and can be interpreted as the rate of plastic slip. Finally, we make the flow stress

$$\Upsilon(\bar{\epsilon}_p) = Y + H\bar{\epsilon}_p + K(1 - \exp(-B\bar{\epsilon}_p)) , \quad (7)$$

a function of the equivalent plastic strain

$$\bar{\epsilon}_p = \sqrt{\frac{2}{3}} \int_0^t \gamma dt , \quad (8)$$

and the following parameters: initial yield, Y ; linear hardening coefficient, H , and nonlinear exponential saturation modulus K and exponent B . In tension, 95 the yield strength, Y , determines the onset of plasticity; the hardening coefficient H determines the linear trend of the post-yield behavior; and K, B superpose a more gradual transition in stress-strain from the trend determined by Young's modulus E in the elastic regime to H in the plastic regime. These material parameters form the basis of our analysis of material variability. To 100 be clear, this standard J_2 plasticity model is a coarse-grained representation of the microstructural variations that engender the variability in the mechanical response, with the plastic strain representing a wide variety of underlying inelastic mechanisms and the physical definitions of the material parameters shaping our interpretation of the underlying causes of the variable response.

105 We approximate the tensile test with a boundary value problem on a rectangular parallelepiped of the nominal gauge section with prescribed displacements on two opposing faces and traction free conditions on the remaining faces to effect pure tension. Finite element simulations are performed in ALBANY [18] using the constitutive model described in this section. The engineering stress σ 110 and strain ε corresponding to that measured in the experiments are recovered from the reaction forces, prescribed displacements, original cross-sectional area and gauge length.

4. Calibration formulation

In general, a calibration problem involves searching for the parameters θ 115 of a given model that minimize the difference between model predictions and observed data. In this work, we adopt a Bayesian approach to the calibration problem [15, 19, 20, 21, 13, 22]. In contrast to least-squares fitting resulting in a single set of parameter values, in a Bayesian perspective the parameters are considered random variables with associated probability density functions 120 (PDFs) that incorporate both prior knowledge and measured data. The choice of Bayesian methods is well motivated by the data which agree with the chosen

model to a high degree, but uncertainty is present in the model parameters both within and across all batches. Bayesian calibration results in a joint distribution of the parameters $p(\boldsymbol{\theta}|\mathcal{D}, M)$ that best fits the available observations \mathcal{D} given the model choice M . The width of the distribution depends on the consistency of the model with the data and the amount of data. By using this probabilistic framework and physical interpretations of the parameters, we aim to quantify the material variability.

4.1. Bayesian inference for parameter calibration

Consider our model M for stress $\sigma = M(\varepsilon; \boldsymbol{\theta})$ comprised of Eqs. (4–8), where ε is the independent variable and $\boldsymbol{\theta} = \{E, Y, H, K, B\}$ are the parameters of interest. By setting to $\{H, K, B\}$ or $\{K, B\}$ zero we can form a nested sequence of models with 2, 3, or 5 parameters with perfect plastic, linear hardening, or saturation hardening phenomenology, respectively. Given that we only have one dimensional tension data, we fix the Poisson's ratio $\nu = 0.3$; however, we allow the Young's modulus, E , to vary, so that the locus of yield points is not constrained to a line.

Bayes rule relates the data and prior assumptions on the parameters into the posterior density of the target parameters as

$$p(\boldsymbol{\theta}|\mathcal{D}, M) = \frac{p(\mathcal{D}|\boldsymbol{\theta}, M) p(\boldsymbol{\theta}|M)}{p(\mathcal{D}|M)} , \quad (9)$$

where $p(\mathcal{D}|\boldsymbol{\theta}, M)$ is the likelihood of observing the data \mathcal{D} given the parameters $\boldsymbol{\theta}$ and model M , $p(\boldsymbol{\theta}|M)$ is the prior density on the parameters reflecting our knowledge *before* incorporating the observations, and $p(\mathcal{D}|M)$ is the evidence. It is important to note that the denominator is typically ignored when sampling from the posterior since it is a normalizing factor, independent of $\boldsymbol{\theta}$, that ensures the posterior PDF to integrate to unity. Here, we will employ relatively uninformative uniform prior densities based on plausible ranges for each of the parameters (more details will be given in Sec. 5). Experimental data influences the resulting posterior probability only through the likelihood $p(\mathcal{D}|\boldsymbol{\theta}, M)$, which is based on the difference between the data \mathcal{D} and the model predictions

$M(\varepsilon; \boldsymbol{\theta})$. The likelihood plays an analogous role to the cost/objective function in traditional fitting/optimization in the sense that it defines what model predictions are close to the data. Specific forms of the likelihood will be discussed in Sec. 4.3. As Eq. (9) suggests, the outcome is conditioned on the model chosen which leads questions of model discrepancy, comparison and selection which will be briefly discussed in Sec. 4.3 and Sec. 5. In general, given the complexities of the model M , the posterior density $p(\boldsymbol{\theta}|\mathcal{D}, M)$ is not known in closed form and one has to resort to numerical methods to evaluate it. Markov Chain Monte Carlo (MCMC) methods [23, 24] provide a suitable way to sample from the posterior density, and to estimate it using, *e.g.*, kernel density estimation.

4.2. Surrogate Model

MCMC sampling of the posterior density involves many sequential evaluations of the model $M(\varepsilon; \boldsymbol{\theta})$. Since the finite-element based forward model $M(\varepsilon; \boldsymbol{\theta})$ is relatively expensive to query (each tension simulation takes approximately 1 cpu-hour), the inverse problem of determining the parameters $\boldsymbol{\theta}$ becomes infeasible. We overcome this by building an efficient, sufficiently accurate surrogate model $\tilde{M}(\varepsilon; \boldsymbol{\theta})$ of the physical response $M(\varepsilon, \boldsymbol{\theta})$ over the region of interest with a polynomial chaos expansion (PCe, see Appendix A for a brief review of polynomial chaos expansions). The fidelity of the surrogate with the full model will be discussed in detail in Sec. 5.1.

Since rough bounds of each of the parameters can be estimated from the data and knowledge of similar materials, we assume that the model parameters $\boldsymbol{\theta}$ follow a uniform prior density $p(\boldsymbol{\theta})$. We construct a corresponding PC expansion of the random vector $\boldsymbol{\theta}$ via

$$\boldsymbol{\theta}(\boldsymbol{\xi}) = \sum_{I=0}^{P_\theta} \boldsymbol{\theta}_I \Psi_I(\boldsymbol{\xi}), \quad (10)$$

where $\boldsymbol{\xi}$ is a vector of standard uniform random variables and P_θ defines the number of terms in the expansion. A corresponding expansion of the model

response can be written as

$$\tilde{M}(\varepsilon; \boldsymbol{\xi}) = \sum_{I=0}^{P_M} \sigma_I(\varepsilon) \Psi_I(\boldsymbol{\xi}), \quad (11)$$

and serves as a suitable surrogate model which, given $\{\sigma_I\}$, can be evaluated by drawing samples from the distribution of $\boldsymbol{\xi}$ and then evaluating the polynomial expansion, Eq. (11). Methods to obtain the coefficients $\{\sigma_I\}$ are discussed in ¹⁷⁰ Appendix A. A by-product of this expansion is the direct access to parametric sensitivities over the range of interest. The sensitivities of the response to the input parameters can be used to identify the most influential parameters. This analysis is particularly useful for parameter elimination in problems with large ¹⁷⁵ dimensionality, as described in more detail in Sec. 5.

4.3. Likelihood Formulation

As mentioned before, the likelihood is the term in Eq. (9) that accounts for the data. To formulate the likelihood, one needs to reason about what data is available and its relationship with the model predictions. From the discussion ¹⁸⁰ in Sec. 2, one can argue that the batches are independent. Furthermore, within a given batch, we assume all the N_i stress-strain curves are independent since each experiment is a self-contained test, performed on separate specimens, *i.e.* the variability of each specimen is the result of its specific microstructure.

We consider two different formulations of the likelihood, which lead to different ¹⁸⁵ formulations of the inverse problem, and hence models of the material variability. The formulations differ by how they account for measurement noise and other variability, and how they are affected by (systematic) model discrepancies. Since each formulation leads to qualitatively different predictions, interpretations, and realizations, we are interested in how each is able to discriminate material variability from other sources of randomness. In this section ¹⁹⁰ and in the Results section we will discuss how, given that plastic strain is a coarse metric of the inelastic deformation in additively manufactured materials, discrepancies between the observed data and the model predictions can be

interpreted physically. The results in Sec. 5 will illustrate how the posterior
195 responds to the quantity of data and its variability.

4.3.1. Additive error formulation

Consider the k -th stress-strain curve from the i -th batch which consists of a sequence of stress observations $\{\sigma_j^{(i,k)}\}_{j=0}^{n_\varepsilon-1}$ obtained at the strain locations $\{\varepsilon_j\}_{j=0}^{n_\varepsilon-1}$. A widely-adopted approach is to express the discrepancy between an observation and surrogate model prediction using an additive noise model as

$$\sigma_j^{(i,k)} = \tilde{M}(\varepsilon_j; \boldsymbol{\theta}) + \eta_j^{(i,k)}, \quad (12)$$

where $\{\eta_j^{(i,k)}\}_{j=0}^{n_\varepsilon-1}$ is the (i, k) sample from the set of random variables $\{\eta_j\}_{j=0}^{n_\varepsilon-1}$ capturing the discrepancy between observations and model predictions at a given ε_j . This formulation is predicated on the assumption that the model $\tilde{M}(\varepsilon; \boldsymbol{\theta})$
200 accurately represents the true, physical process occurring with fixed, but unknown, parameters. This a strong assumption (and one of the main deficiencies of this approach) since models are, in general, only approximations of observed behavior. Nevertheless, this is a commonly used method due to its simplicity.

In lieu of a completely characterized measurement error model (which is rarely obtained in practice), it is reasonable to assume the errors to be independent and identically distributed (i.i.d.) Gaussian random variables with zero mean, *i.e.* $\eta_j \sim \mathcal{N}(0, \varsigma^2)$, where ς^2 is the variance. This yields the following likelihood

$$p(\mathcal{D}_i^{(k)} | \boldsymbol{\theta}, \tilde{M}) = \prod_{j=0}^{n_\varepsilon-1} (2\pi\varsigma^2)^{-1/2} \exp\left(-\frac{(\sigma_j^{(i,k)} - \tilde{M}(\varepsilon_j; \boldsymbol{\theta}))^2}{2\varsigma^2}\right), \quad (13)$$

where we recall that $\mathcal{D}_i^{(k)}$ represents the stress observations collected from the k -th stress-strain curve of the i -th batch. By assuming that each curve is independent from another, we can write

$$p(\mathcal{D}_i | \boldsymbol{\theta}, \tilde{M}) = \prod_{j=0}^{n_\varepsilon-1} \prod_{k=1}^{N_i} (2\pi\varsigma^2)^{-1/2} \exp\left(-\frac{(\sigma_j^{(i,k)} - \tilde{M}(\varepsilon_j; \boldsymbol{\theta}))^2}{2\varsigma^2}\right). \quad (14)$$

for the full dataset of the i -th batch. The standard deviation ς can either be
205 fixed in advance, if some knowledge about the experimental process is available,

or it can be inferred along with the target parameters $\boldsymbol{\theta}$. Moreover, it can be assumed to be either constant or varying with the data points.

The basic additive error formulation can be enriched by adding a term capturing the discrepancy between the model prediction and truth represented by the physical data, leading to

$$\sigma_j^{(i,k)} = \tilde{M}(\varepsilon_j; \boldsymbol{\theta}) + \hat{\eta}_j^{(i,k)} + \eta_j^{(i,k)}, \quad (15)$$

where $\hat{\eta}$ represents the discrepancy between the model prediction and truth. A structure for the model error is more difficult to prescribe than that for the data error. In fact, the calibrated model now effectively becomes $\tilde{M}(\varepsilon, \boldsymbol{\theta}) + \hat{\eta}(\varepsilon)$, and not simply the original $\tilde{M}(\varepsilon, \boldsymbol{\theta})$. Given that this additional term is not physically associated with the presumed sources of non-measurement variability its applicability outside the training regime is delicate. Lastly, this additive term can yield difficulties because it can lead to violations of physical laws [25, 26].

²¹⁵ 4.3.2. *Embedded model discrepancy*

A more suitable approach to representing variability embedded in the physical model involves adding the model discrepancy error [27, 28] to the parameters

$$\sigma_j^{(i,k)} = \tilde{M}(\varepsilon_j; \boldsymbol{\theta} + \hat{\boldsymbol{\eta}}) + \eta_j^{(i,k)}, \quad (16)$$

where $\eta_j^{(i,k)}$ is an additive noise term akin to that in Eq. (13). In this case, $\boldsymbol{\theta} + \hat{\boldsymbol{\eta}}$ is a random vector with density and moments to be estimated, whereas $\eta_j^{(i,k)}$ is determined by *a priori* estimates of measurement noise. The random vector $\boldsymbol{\theta} + \hat{\boldsymbol{\eta}}$ can be represented with a PCe. For instance, for a single parameter θ we can write

$$\theta + \hat{\eta} = \sum_I \alpha_I \Psi(\xi). \quad (17)$$

The problem is thus transformed into a density estimation problem, where our objective is now to estimate $\boldsymbol{\alpha} = \{\alpha_0, \alpha_1, \dots\}$ that parametrize and define the density of $\theta + \hat{\eta}$. This is in contrast to the conventional use of Bayesian inference for parameter estimation, *i.e.* additive error formulations, in which one strictly infers the *parameter* and not its density. Also, the data for our present

calibration problem motivates the embedded approach since it suggests the uncertainties are aleatoric/irreducible rather than epistemic/reducible.

In the conventional case, as more data is taken into account, the width of the posterior density shrinks, tending to a Dirac delta function at the true parameter value assuming informative data and negligible model discrepancy. On the other hand, in the present context of *density estimation*, the objects of inference are the parameters α , and the posterior density is thus on α . Thus, the more data is taken into account, provided the data is sufficiently informative, the distribution on α narrows while the width of the distribution of the parameters θ remains finite and conforms to the data. For this embedded technique, the model calibration problem thus involves finding the posterior distribution on α via Bayes theorem Eq. (9)

$$p(\alpha|\mathcal{D}, \tilde{M}) \sim p(\mathcal{D}|\alpha, \tilde{M}) p(\alpha|\tilde{M}), \quad (18)$$

where α has been substituted for θ and, again, $p(\alpha|\mathcal{D}, \tilde{M})$ is the posterior, $p(\mathcal{D}|\alpha, \tilde{M})$ is the likelihood, and $p(\alpha|\tilde{M})$ is the prior. Once the posterior for the parameters is characterized, it can be propagated through the model to obtain the posterior predictive distributions for quantities of interest (namely stress in this case). The key feature of these predictions is that their uncertainty is affected by both parameter and model uncertainties. For brevity, we leave the full mathematical details of this embedded approach, including the likelihood formulation, to Appendix B.

5. Results

In this section, we present the details of the construction of the particular surrogate models from the full finite element plasticity model, their calibration to the experimental data, and the physical interpretation of the resulting predictions and parameter estimates. Most of the numerical results presented below are obtained using the UQ Toolkit [29] package.

5.1. Surrogate Model

To describe the material stress-strain behavior we analyze, calibrate and compare three nested plasticity models of increasingly complex phenomenology, namely perfect plasticity, linear hardening and saturation hardening. As mentioned in Sec. 3, we focus on five parameters: Young’s modulus, E ; yield strength, Y ; hardening modulus, H ; saturation modulus, K ; and saturation exponent, B , which control the elastic-plastic stress response. We build Legendre-Uniform PC expansions of these parameters by assuming that they are uniformly distributed over a chosen range

$$\begin{aligned} E &= 200 + 80 \xi_1 \text{ [GPa]}, \\ Y &= 1.2 + 0.5 \xi_2 \text{ [GPa]}, \\ H &= 3.005 + 2.995 \xi_3 \text{ [GPa]}, \\ K &= 0.2 + 0.2 \xi_4 \text{ [GPa]}, \\ B &= 1750 + 1250 \xi_5 \end{aligned} \tag{19}$$

where $\{\xi_1, \xi_2, \xi_3, \xi_4, \xi_5\} \sim \mathcal{U}(-1, 1)$ are independent identically distributed (i.i.d.) standard uniform random variables. We chose these parameters ranges to be
²⁴⁰ large enough that the corresponding predictions can capture the variability of the experimental data shown in Fig. 1b. Also, we remark that the expansion with i.i.d. random variables is a common step to build the surrogate model. Any correlations between the parameters will then be discovered through the inverse problem, see *e.g.* Refs. [20, 21, 22]. The priors for α are constructed such that
²⁴⁵ the target physical parameters have their target priors, Eq. (19).

5.1.1. Two parameter perfect plasticity model

For the two parameter model, the stress is expressed as a function of strain and Young’s modulus, E , and yield strength, Y according to

$$\sigma(\varepsilon) = M^{(2)}(\varepsilon; \{E, Y\}), \tag{20}$$

where we use the superscript “(2)” to identify this as the two parameter model. As mentioned before, the full tension simulation is expensive to evaluate, so we

leverage the PC expansions of the inputs, Eq. (19), to create PCes of the stress at each strain $\{\varepsilon_j\}$ value

$$\sigma(\varepsilon_j) = \tilde{M}_j^{(2)}(\xi_1, \xi_2) \approx M^{(2)}(\varepsilon_j, E(\xi_1), Y(\xi_2)), \quad j = 0, \dots, n_\varepsilon - 1. \quad (21)$$

To build this sequence of PCes, $\{\tilde{M}_j^{(2)}, j = 0, n_\varepsilon - 1\}$, we employ regression using uniform random samples in the two-dimensional space (ξ_1, ξ_2) . In particular, we generate 1004 training (1000 samples in the inner domain and four additional ones for the corners) to build the surrogate model, and 250 validation samples to assess its accuracy and check for over-fitting. (We chose regression over the more computationally efficient stochastic collocation via sparse quadrature grids given the superior results in constructing the piece-wise smooth surrogates to be introduced forthwith.) Fig. 2 shows the training samples mapped back to the physical domain (E, Y) using Eq. (19), and the corresponding plasticity simulations. Given the number of samples, the regression approach is suitable for constructing PC expansions of up to ninth order. Computing a higher-order expansion would lead to an under-determined problem. The ensemble of responses, shown in Fig. 2b, are clearly piece-wise linear, as expected, with a slight downward slope in the post-yield response due to finite deformation effects.

For the two parameter model, a key feature is that the stress-strain behavior has a discontinuity of the first derivative when the behavior switches from elastic to plastic as the stress, which depends on E , exceeds the particular yield Y value. This dependence on both parameters can be observed in Fig. 3, which shows the stress plotted as a function of E and Y at three different strain locations, $\varepsilon = 0.12, 0.39, 1.95\%$. These points have been chosen such that the first, $\varepsilon = 0.12\%$, is within a regime for which all points in the (E, Y) space are fully elastic; the second value, $\varepsilon = 0.39\%$, is in a mixed elastic-plastic regime, and the third point, $\varepsilon = 1.95\%$, identifies a fully plastic regime. Fig. 3 shows that within the elastic and plastic regimes, the behavior is linear, and the separator is a line. We will exploit this observation, and use the plastic strain (which is zero in the elastic regime) as a classifier to sub-divide the training data.

To compare the accuracy of global PC expansions up to ninth order and
275 a piecewise linear surrogate built over the elastic and plastic sub-domains, in
Fig. 4 we show the relative error based on the ℓ_2 -norm and ℓ_∞ -norm for the
various PCes as a function of the strain. From Fig. 4 we observe that a *global*
linear PCe is accurate where the regime is either fully elastic or plastic, but
inaccurate in the mixed region where the discontinuity in the response makes
280 a global representation sub-optimal. Also, as we increase the order of the PCe
from first to fifth order, the results do not change within the elastic and plastic
regions, but improve in the mixed region. However, when the order of the
expansion is at least fifth order, the errors do not decrease as rapidly which
suggests over-fitting. Lastly, the low-order piece-wise linear surrogate has the
285 lowest error in both norms across the strain range and, hence, it is more suitable
than a high order global surrogate for this model response.

5.1.2. Three parameter linear hardening model

Augmenting the two parameter model with the post-yield phenomenology
controlled by the hardening modulus (H) results in the three parameter model,
290 $\sigma(\varepsilon) = M^{(3)}(\varepsilon; \{E, Y, H\})$. In this case, we use a total of 2008 training samples
(2000 in the inner domain and the 8 additional ones for the corners), and 500
validation points.

Again, we build global polynomial surrogates of increasing order and a piece-
wise low order polynomial surrogate to represent the response surface. To cap-
295 ture the additional complexity in the post-yield response we used a quadratic
PC in the plastic regime of the mixed, piece-wise surrogate which is linear in
the elastic regime. The resulting ℓ_2 -norm and ℓ_∞ -norm errors (not shown for
brevity) have the same trends as those for the two parameter model shown in
Fig. 4 due to similarity in the slope discontinuity of the response and, likewise,
300 the piece-wise surrogate is the best representation of the full simulation response
for linear hardening.

5.1.3. Five parameter saturation hardening model

The five parameter model adds the saturation modulus (K) and saturation exponent (B), yielding $\sigma(\varepsilon) = M^{(5)}(\varepsilon; \{E, Y, H, K, B\})$. To build and check this surrogate, we collected 5032 training and 500 validation samples. As shown in Fig. 5, the five parameter model adds a smoother elastic-plastic transition through the presence of the saturation modulus and exponent.

Fig. 6 shows the relative error based on the ℓ_2 -norm and ℓ_∞ -norm for the various PCes as a function of the strain. As with the piece-wise surrogate for the three-parameter model, here we use a linear PCe for the elastic and a quadratic PCe for the plastic. Unlike the results for the two simpler models, the piece-wise surrogate does not outperform the global surrogates. In this case, a global PCe of order ≥ 6 gives the lowest errors. Since the errors are comparable for these polynomials, orders > 6 are likely over-fitting the full simulation data.

5.2. Sensitivity Analysis

As mentioned, one advantage of building a PCe surrogate is that one can obtain global Sobol sensitivities of a target quantity of interest with respect to the input parameters [30]. Here we compute the total sensitivities [30] of the stress using the surrogate built at each strain point. Fig. 7 shows the sensitivities obtained over the range of the surrogate for the: (a) two, (b) three, and (c) five parameter models. The sensitivities are influenced by the range chosen to build the surrogate model. The experimental stress-strain curves show little hardening and so the surrogate was constructed with over a narrow range of hardening parameters, Eq. (19), which appropriately minimizes their importance. Also, it is apparent that the relative importance of the parameters evolves with strain. As expected, Fig. 7 shows that the Young's modulus E is the dominant parameter within the elastic regime. At larger strains, in the transition between the elastic and plastic regimes, E becomes gradually less important and the yield Y starts to dominate. In the more complex models, the hardening H and the saturation parameters K and B play a relative minor role due to the fact that the data displays little hardening and the two parameter model is a good rep-

resentation of the majority of the stress behavior. In both the three and five parameter models the sensitivity to H is slim and almost negligible, whereas the added K in the five parameter model is clearly not negligible apparently for its role in determining the “knee” in the stress-strain at the elastic-plastic transition. More discussion of this behavior will be given in Sec. 5.3.3.

5.3. Calibration

In this section, we discuss and contrast the results obtained from the inverse problem formulated the additive error and those obtained using the embedded formulation. We use the three parameter, linear hardening model as a reference case, which we discuss in detail, and then show the main results for the other models.

5.3.1. Inversion with the additive error model

We assume the measurement noise to be constant, ς^2 , along the strain axis, *i.e.* the measurement error does not depend on the strain ε . Hence, the parameters to be inferred are $\boldsymbol{\theta} = \{E, Y, H, \varsigma^2\}$. As priors, we choose uniform densities with ranges coinciding with those chosen to build the surrogate model in Eq. (19). For the variance, we choose a uniform prior over the positive axis. This is typically appropriate because the surrogate might not be reliable outside the range where it was computed on. The prior plays a minor role if a substantial amount of data is available, making the problem likelihood-informed rather than prior-informed in the large data limit.

We leverage this case to highlight some key features of Bayesian calibration applied to the present problem of representing material variability: first, the dependence on the type of surrogate model used, second, the batch-to-batch differences in the resulting parameters, third, the correlations between the target parameters, and finally, convergence of the results with the amount of data used in the inverse problem.

Fig. 8 shows the joint posteriors between the physical parameters, Fig. 8(a,b,c), and the marginalized posterior for the standard deviation of the measurement

error, Fig. 8(d), obtained using 10 stress-strain curves randomly chosen from \mathcal{D}_3 . Results are shown for each of the surrogates we constructed. The convergence of the resulting PDFs with polynomial order gives us confidence that the higher order global surrogates and the piece-wise surrogate lead to sufficiently accurate 365 posterior parameter distributions. Since the distributions are not skewed, it is apparent that the extracted elastic modulus E is not correlated with the yield Y , or the hardening H , whereas there is a weak negative correlation between yield and hardening. We conjecture that the pre-yield data informs E independent 370 of the other parameters and, likewise, the yield point informs Y ; however, there is a trend in the experimental data, see Fig. 1c, for high yield points to lead to subsequent low post yield slopes and vice versa. The maximum *a posteriori* (MAP) value of the inferred standard deviation is around 0.027 GPa. This value is larger than the one estimated directly from the experimental data 375 (0.020 GPa). This is expected because we are not accounting for model error and, therefore, the model discrepancy is lumped into the measurement error. Note that the range used for the plots are much smaller than those originally chosen for the surrogate construction, Eq. (19).

Fig. 9 shows the joint posteriors among the physical parameter for the separate batches obtained using 40 randomly selected curves from each batch. 380 Clearly, the mean parameters of the batches are quite variable and the distributions are, for the most part, distinct and well-separated. However, the correlation structures are similar, suggesting that the batches behave qualitatively in the same way.

Fig. 10 illustrates the convergence trend in the posterior distributions as a 385 function of the number N of curves used in the calibration. The panels of Fig. 10 show samples of the joint posteriors $p(E, Y)$, $p(E, H)$ and $p(Y, H)$ obtained from the third batch, as a function of the number N of curves used to run the problem. The densities shift and narrow as more data is taken into account, and, given the range of the graphs, it appears that the densities have more-or-less converged 390 which suggests sufficient data has been obtained to characterize the parameters.

Fig. 11 shows the predictions using the posterior distribution along with an

ensemble of curves from the third batch used to run the inference. The error bars represent the posterior predictive uncertainty stemming from the data noise, while the black solid line represents the mean prediction and the gray band represents the ± 2 standard deviations due to posterior uncertainty. In
395 this case, the data set leads to a narrow posterior distribution around the mean curve, which is reflected in a tight gray band. Overall, we see that the model prediction and the superimposed error bars cover the experimental data and, hence, provide a good representation of it. However, the inferred value of the
400 measurement noise is not representative of the actual noise that was estimated from the experimental data since the lack of model discrepancy term amplifies the inferred data noise. Hence, despite the predictive analysis showing that we are able to recover the variability of the data, our conclusion is this is an unsuitable inversion model since the full variability is represented as measurement error. Consequently, predictions made using this model would underestimate
405 and inadequately account for the actual variability in the material.

5.3.2. *Inversion with embedded error model*

In this section, we discuss some key calibration results obtained using the embedded error model.

410 Fig. 12 shows the posterior predictive plots for the 2, 3, and 5 parameter models along with an ensemble of stress-strain curves from \mathcal{D}_3 . (For brevity we only show the results for one representative batch, since the others yield similar results.) For each plot, the blue bars represent the posterior uncertainty from the data noise, the black solid line represents the mean prediction, the dark gray band represents the ± 2 standard deviations due to posterior uncertainty, and the light gray band represents the ± 2 standard deviations due to model error. We observe that in all cases the experimental data is captured and described well
415 by the model predictions. Here, the variability of the data is mostly described by the prediction uncertainty due to parameter variability, while the data noise
420 is small and comparable to the estimate of the measurement error obtained from the data itself. This is the key difference with respect to the additive results

shown in Fig. 11. By accounting for model discrepancy through the embedded error terms, we are able to characterize the variability more properly because it is not artificially lumped in the data noise. The contribution stemming from the 425 posterior uncertainty is again quite small, suggesting that we have accounted for sufficient amount of data. The results for three models show that overall they have similar predictive capabilities.

5.3.3. Comparison of additive and embedded error models

To illustrate the differences between the additive and embedded error-based 430 inference results, Fig. 13 shows the posterior PDFs for each parameter using the data from the third batch obtained from the additive inference (left column) and embedded error approach (right column). We note in both cases, the contribution stemming from the posterior uncertainty of the parameter estimates is again quite small, as shown in Fig. 11 and Fig. 12, suggesting that we have 435 accounted for a sufficiently large amount of data. The key difference in the parameter posterior distributions is the embedded term which explicitly enables the calibrated models to reflect the inherent variability in the material parameters. In contrast, calibration with the additive error term tends to attribute variability to that component of the model which cannot be interpreted as a 440 specific level of variability in any particular parameter.

Posterior PDFs of the two parameters common to all the models, the Young's modulus E and yield strength Y , highlight this difference in uncertainty quantification strategies. Fig. 13 shows that the embedded approach attributes a significant amount of variability in E based on the data while the additive 445 approach has a small uncertainty. In effect, the distribution in the additive approach converges to the best average value for E , since E is treated as a constant rather than a distribution. All variability arising from data variation, as opposed to measurement error and finite sample size, is attributed to the additive term where it is confused with all other uncertainties. With the 450 embedded approach, the distributions of the elastic modulus E for the three models are broad, flat, and fairly consistent. Likewise, the difference in the two

approaches is readily apparent in the distributions of Y shown in Fig. 13. The distributions of the yield strength Y for the two simpler plasticity models are essentially the same, as they treat yield as a well-defined point unlike the more complex 5-parameter, saturation hardening model. When interpreted in light of the model sensitivities shown in Fig. 7 and discussed in Sec. 5.2, Y is well informed by the data, resulting in the narrow distribution and small uncertainties in the additive approach and consistent broad distributions using the embedded approach. The low sensitivity to H in the 3 and 5-parameter models gives rise to broad Gaussian distributions with the additive formulation indicating more informative data is needed. This deficiency also contributes to the qualitative differences in the densities for H and Y with the additive formulation. In contrast to H , the 5 parameter model is sensitive to both the Y and K , as shown in Fig. 7, and yet the data is not sufficient to fully inform each independently. Instead, the calibration results in a broad joint PDF of the two parameters since they effect similar changes in the model's behavior near the elastic-plastic transition. This confounding effect also likely gives rise to the bimodal posterior distributions of these parameters. Lastly, given that the sensitivity to B is essentially negligible, it is not surprising that the prior, restricted by the range of the surrogate, exerts significant influence on the posterior distribution of this parameter in both formulations.

Fig. 14 shows 100 posterior predictive realizations obtained using the 5 parameter model for the additive approach Fig. 14(a,c) and embedded approach Fig. 14(b,d) calibrated using \mathcal{D}_3 . The top row shows the results obtained by sampling the joint posterior density of the parameters $\boldsymbol{\theta}$, and pushing these samples through the model $\tilde{M}(\varepsilon_j; \boldsymbol{\theta})$ only; while the bottom row shows the results with the contribution of the measurement noise $\eta_j^{(3,k)}$. It is evident from comparing Fig. 14 to the data shown in Fig. 1b that the embedded error approach yields a suitable representation where the variability of the response is determined by the variability of the material parameters. On the contrary, the additive approach yields a tight envelope of predictions, and the full variability is only captured by the added (and over-estimated) measurement noise. Com-

paring individual realizations to the curves obtained experimentally, it is clear that the embedded approach with noise yields curves that match quite well the 485 trends observed in the experiment, the classical method deviates considerably and is highly dominated by the high frequency, uncorrelated noise.

6. Discussion

Following the discussion begun in Sec. 4.3.2, we will use a simplification of the model $M(\varepsilon, \theta)$ to summarize the key concepts in this work and help generalize the intuition needed to model variability. Let us consider only the elastic response so that the nominal model is:

$$\sigma(\varepsilon) = M(\varepsilon; E) = E\varepsilon, \quad (22)$$

and limit our attention to data for a single batch in the elastic regime. The embedded model of the data is

$$\sigma_j = M(\varepsilon_j; E + \alpha \hat{\eta}) + \eta_j = (E + \alpha \hat{\eta})\varepsilon_j + \eta_j, \quad (23)$$

where $\hat{\eta}$ and η_j are mean zero random variables. The additive model omits $\alpha \hat{\eta}$ which varies the slope of the stress-strain curve. In both cases,

$$\mathbb{E}[\sigma_j] = (E + \alpha \hat{\eta})\varepsilon_j + \eta_j = E\mathbb{E}[\varepsilon_j], \quad (24)$$

so, in the limit of infinite informative data, both formulations recover the correct mean $E = \frac{\mathbb{E}[\sigma_j]}{\mathbb{E}[\varepsilon_j]}$. This is illustrated in the comparison of Fig. 14a with 490 the average of Fig. 14b in the elastic regime. The difference between the two formulations becomes clear when examining the variance at a given ε_i and covariance across all samples ε_i of the calibrated representations. Recall that the stress-strain data for a single specimen has virtually no measurement noise and yet the stress-strain curves are essentially lines with slopes that vary across a 495 batch. Without the $\alpha \hat{\eta}$ term, the simplest additive model, where the sequence of random variables $\{\eta_j\}$ are assumed independent and identically distributed, obtains a variance around the mean that is determined by the total variance of

the dataset. Furthermore, the realizations are not smooth. They have a wide and constant variation around their mean trends and, consequently, have distinct offsets in stress at zero strain, as can be seen in Fig. 14c. In contrast to this model employing only uncorrelated noise η_j , the embedded model accounts for most of the variation in the data with a distribution of slopes effected by the $\alpha \hat{\eta}$ term. This is consistent with the variations in the dataset which is composed of highly correlated data for each test, *i.e.* each test gives essentially the same linear relationship between σ_j and ε_j at every j . The fan-like ensemble of realizations shown in Fig. 14d clearly represents the continuity and the particular type of variation seen in the data.

This basic illustration was motivated by our data where each individual experiment is well-described by the hypothesized model. When that is the case, the example shows that the embedded model better represents the intrinsic material variability and, by extension, the underlying physics. If, however, external measurement noise were the primary source of uncertainty, the additive model would provide a good representation. As in traditional constitutive modeling, a rational choice of how to formulate the representation of observed variability can only be assessed by examination of the data, and then confirming the validity of that choice by comparing synthetically replicated experiments to the observed behavior, as in Fig. 14. This choice can be guided by examining whether or not the apparent noise is correlated with the mean behavior and the model prediction, as was done in this work.

In general, there are three kinds of uncertainty that should be considered during calibration to experimental data: (A) external measurement error, (B) intrinsic variability, and (C) model form error. Given that measurement noise is typically uncorrelated with the underlying physical response it is typically modeled with white noise. Moreover, it is reducible by replicating the experiment and collecting more data in the sense that the posterior distributions of model parameters converge and narrow. In contrast, variability in the material properties cannot be reduced by increased data gathering, although more data will typically better inform the estimated joint distribution of material parameters.

The hallmark of inherent variability is individual experimental curves which are
530 well explained by a model with appropriate physical parameters, but have sys-
tematic parametric discrepancies across the set of curves. The embedded error
formulation is well-suited to represent this source of variability. Finally, model
form error refers to relevant physics which are unincorporated in the model, and
manifests itself through discrepancies between model predictions and the actual
535 data. When present, a model of a single realization will display a systematic
discrepancy from the data it is trying to emulate. Since this error can confound
the determination of the other errors it is crucial to perform model selection, as
was done in this work albeit for a dataset that was generally well-represented by
all members of the model family. (It should be noted the embedded formulation
540 [13] was originally developed to mitigate this type of error.)

7. Conclusions

We have presented a method which can model the variability of a mate-
rial which is well-described by existing plasticity models of mean response, but
contains microstructural variability leading to different macroscopic material
545 properties. By leveraging the embedded error method of Sargsyan *et al.* [13]
(which, as mentioned, was originally developed to address model discrepancy),
we can mathematically represent the material variability as material parameters
drawn from a well-calibrated joint distribution. This Bayesian approach is con-
sistent with our understanding of microstructural variability and appropriate for
550 UQ studies requiring forward propagation of this variability. In particular, we
developed a constitutive model of variability that is amenable to non-intrusive
sampling and adaptable to direct evaluation in simulation codes that handle
fields of distributions. This enables a robust design methodology that can pre-
dict performance margins with high confidence.

555 Another important contribution of this work is to contrast the proposed ap-
proach with commonly used uncertainty formulations. The standard, additive
error formulation appropriately accounts for the uncertainty in the experiments

arising from measurement error. Yet, in the case of inherent variability, it characterizes all the uncertainty as measurement error which results in unwarranted confidence in the material properties and an inability to correctly understand how that variability would manifest in applications. We have demonstrated that the embedded method accurately characterizes the aleatoric uncertainty present in the experimental observations and enables “black-box” engineering UQ analysis. It gives insight into what aspects of a homogeneous, macroscale constitutive model are most strongly affected by microstructural variability and enables quantitative model selection. It is able to distinguish the variability of different batches and therefore assess their relative performance. It demonstrates convergence with increasing data and the convergence of common parameters in a nested hierarchy of models. Moreover, it is capable of representing both significant external noise and inherent variability in a unified formulation. In future work, we will extend the methodology to the post-necking failure behavior of similar materials.

Acknowledgments

This work was supported by the LDRD program at Sandia National Laboratories, and its support is gratefully acknowledged. B.L. Boyce would like to acknowledge the support of the Center for Integrated Nanotechnologies. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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.

- [1] R. Hill, Elastic properties of reinforced solids: some theoretical principles, *Journal of the Mechanics and Physics of Solids* 11 (5) (1963) 357–372.
- [2] S. Nemat-Nasser, Averaging theorems in finite deformation plasticity, *Mechanics of Materials* 31 (8) (1999) 493–523.

585 [3] D. McDowell, S. Ghosh, S. Kalidindi, Representation and computational structure-property relations of random media, *JOM Journal of the Minerals, Metals and Materials Society* 63 (3) (2011) 45–51.

590 [4] K. K. Mandadapu, A. Sengupta, P. Papadopoulos, A homogenization method for thermomechanical continua using extensive physical quantities, *Proc. R. Soc. A* 468 (2142) (2012) 1696–1715.

[5] R. Dingreville, C. C. Battaile, L. N. Brewer, E. A. Holm, B. L. Boyce, The effect of microstructural representation on simulations of microplastic ratcheting, *International Journal of Plasticity* 26 (5) (2010) 617–633.

595 [6] C. C. Battaile, J. M. Emery, L. N. Brewer, B. L. Boyce, Crystal plasticity simulations of microstructure-induced uncertainty in strain concentration near voids in brass, *Philosophical Magazine* 95 (10) (2015) 1069–1079.

600 [7] J. M. Emery, R. V. Field, J. W. Foulk, K. N. Karlson, M. D. Grigoriu, Predicting laser weld reliability with stochastic reduced-order models, *International Journal for Numerical Methods in Engineering* 103 (12) (2015) 914–936.

[8] W. E. Frazier, Metal additive manufacturing: a review, *Journal of Materials Engineering and Performance* 23 (6) (2014) 1917–1928.

605 [9] B. L. Boyce, B. C. Salzbrenner, J. M. Rodelas, L. P. Swiler, J. D. Madison, B. H. Jared, Y.-L. Shen, Extreme-value statistics reveal rare failure-critical defects in additive manufacturing, *Advanced Engineering Materials*.

[10] O. Le Maître, O. M. Knio, *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*, Springer Science & Business Media, 2010.

610 [11] D. Xiu, *Numerical methods for stochastic computations: a spectral method approach*, Princeton university press, 2010.

[12] R. C. Smith, Uncertainty quantification: theory, implementation, and applications, Vol. 12, Siam, 2013.

[13] K. Sargsyan, H. Najm, R. Ghanem, On the statistical calibration of physical models, *International Journal of Chemical Kinetics* 47 (4) (2015) 246–276.

615 [14] R. Field, M. Grigoriu, J. Emery, On the efficacy of stochastic collocation, stochastic Galerkin, and stochastic reduced order models for solving stochastic problems, *Probabilistic Engineering Mechanics* 41 (2015) 60–72.

[15] M. C. Kennedy, A. O'Hagan, Bayesian calibration of computer models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63 (3) (2001) 425–464.

620 [16] J. C. Simo, A framework for finite strain elastoplasticity based on maximum plastic dissipation and the multiplicative decomposition: Part i. continuum formulation, *Computer methods in applied mechanics and engineering* 66 (2) (1988) 199–219.

625 [17] J. Simo, T. Hughes, Computational Inelasticity, Springer New York, New York, NY, 1998.

[18] A. G. Salinger, R. A. Bartlett, A. M. Bradley, Q. Chen, I. P. Demeshko, X. Gao, G. A. Hansen, A. Mota, R. P. Muller, E. Nielsen, et al., Albany: Using component-based design to develop a flexible, generic multiphysics analysis code, *International Journal for Multiscale Computational Engineering* 14 (4).

630 [19] D. Sivia, Data Analysis: A Bayesian Tutorial, Oxford Science, 1996.

[20] F. Rizzi, O. Knio, H. Najm, B. Debusschere, K. Sargsyan, M. Salloum, H. Adalsteinsson, Uncertainty Quantification in MD Simulations. Part II: Inference of force-field parameters, *SIAM J. Multiscale Model. Simul.* 10 (4) (2012) 1460–1492.

[21] F. Rizzi, R. E. Jones, B. J. Debusschere, O. M. Knio, Uncertainty quantification in md simulations of concentration driven ionic flow through a silica nanopore. ii. uncertain potential parameters, *The Journal of Chemical Physics* 138 (19) (2013) 194105.
640

[22] Y. M. Marzouk, H. N. Najm, L. A. Rahn, Stochastic spectral methods for efficient bayesian solution of inverse problems, *Journal of Computational Physics* 224 (2) (2007) 560 – 586.

[23] D. Gamerman, H. F. Lopes, *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*, CRC Press, 2006.
645

[24] B. A. Berg, A. Billoire, *Markov chain Monte Carlo simulations*, Wiley Online Library, 2008.

[25] M. Salloum, J. A. Templeton, Inference and uncertainty propagation of atomistically-informed continuum constitutive laws, part 1: Bayesian inference of fixed model forms, *International Journal for Uncertainty Quantification* 4 (2).
650

[26] M. Salloum, J. A. Templeton, Inference and uncertainty propagation of atomistically informed continuum constitutive laws, part 2: Generalized continuum models based on gaussian processes, *International Journal for Uncertainty Quantification* 4 (2).
655

[27] K. Sargsyan, H. N. Najm, R. Ghanem, On the statistical calibration of physical models, *International Journal of Chemical Kinetics* 47 (4) (2015) 246–276.

[28] C. Safta, M. Blaylock, J. Templeton, S. Domino, K. Sargsyan, H. Najm, Uncertainty quantification in LES of channel flow, *International Journal for Numerical Methods in Fluids* 83 (4) (2017) 376–401.
660

[29] B. Debusschere, K. Sargsyan, C. Safta, K. Chowdhary, UQ Toolkit, <http://www.sandia.gov/UQToolkit> (2017).

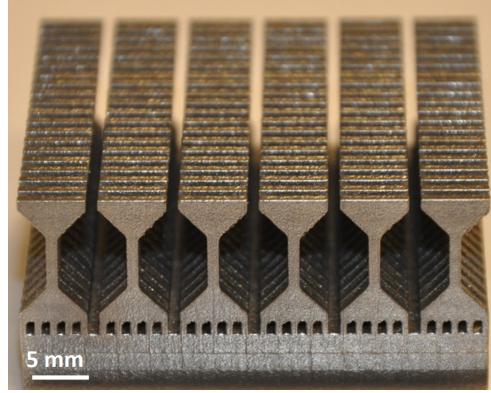
[30] B. Sudret.

⁶⁶⁵ [31] R. Ghanem, P. Spanos, Stochastic Finite Elements: A Spectral Approach, Springer Verlag, New York, 1991.

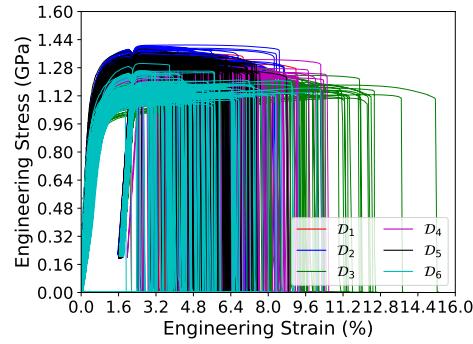
[32] D. Xiu, G. Karniadakis, The Wiener-Askey polynomial chaos for stochastic differential equations, *SIAM Journal on Scientific Computing* 24 (2) (2002) 619–644.

⁶⁷⁰ [33] O. Le Maître, O. Knio, *Spectral Methods for Uncertainty Quantification*, Springer, New York, NY, 2010.

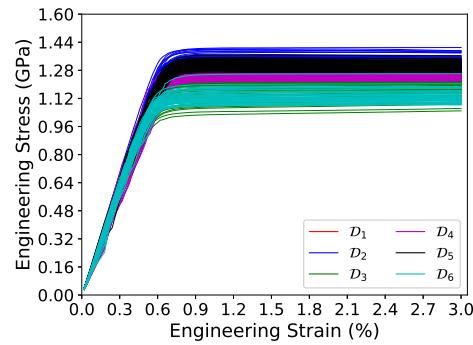
[34] R. Cameron, W. Martin, The orthogonal development of nonlinear functionals in series of Fourier-Hermite functionals, *Annals of Mathematics* 48 (1947) 385–392.



(a)

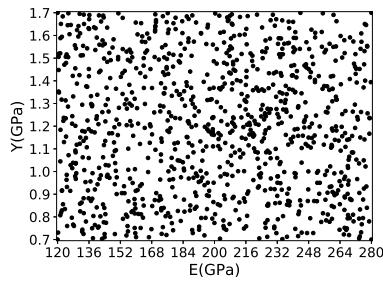


(b)

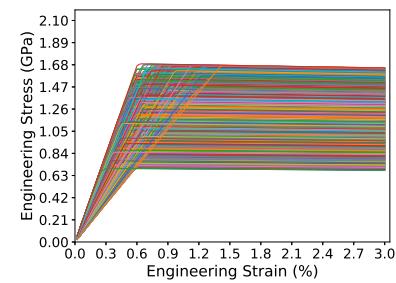


(c)

Figure 1: (a) An array of nominally identical micro-tension ‘dogbone’ specimens, (b) experimental data from Boyce *et al.* [9] color-coded by batch, and (c) the reduced data set used in this work.

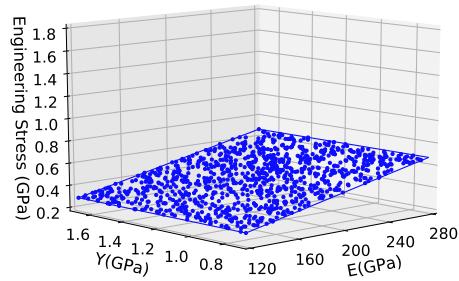


(a)

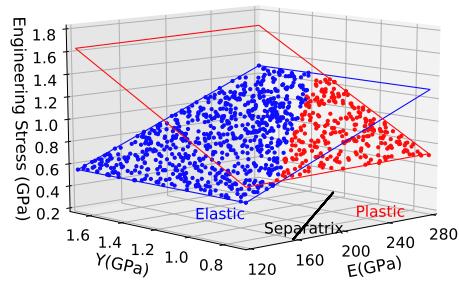


(b)

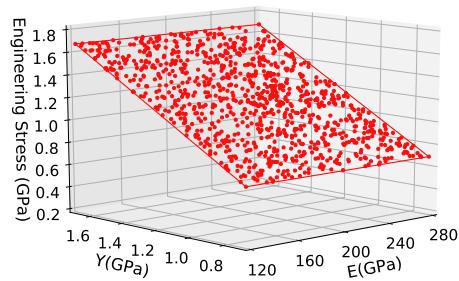
Figure 2: Two parameter surrogate data: (a) training samples in the space (E, Y) and (b) corresponding stress-strain curves used to build the surrogate for the two parameter model.



(a)

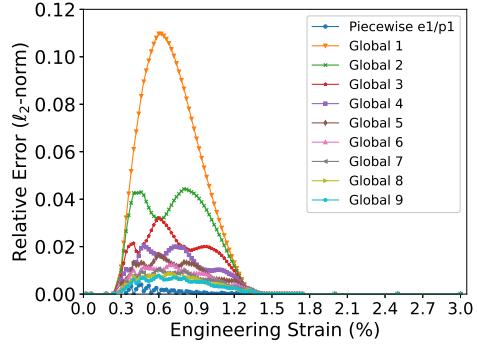


(b)

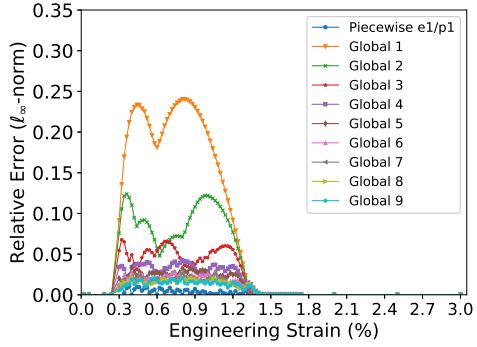


(c)

Figure 3: Stress plotted as a function of elastic modulus (E) and yield (Y) at three different strains: (a) fully elastic, $\varepsilon = 0.12\%$, (b) elastic-plastic, 0.39% , and (c) fully plastic 1.95% , over the domain of the two parameter model.



(a)



(b)

Figure 4: Surrogate error for the two parameter model at each target strain location based on (a) the ℓ_2 -norm (a) and (b) the ℓ_∞ -norm. Results are shown for global PC expansions of orders up to nine, as well as the piece-wise surrogate based on linear polynomials for both the elastic and plastic response.

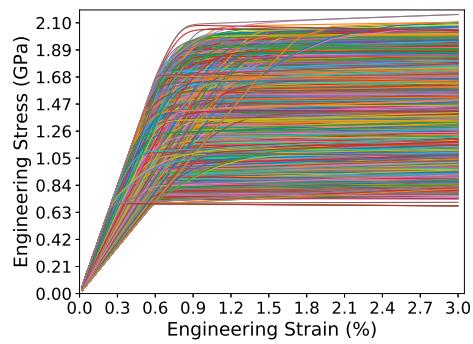
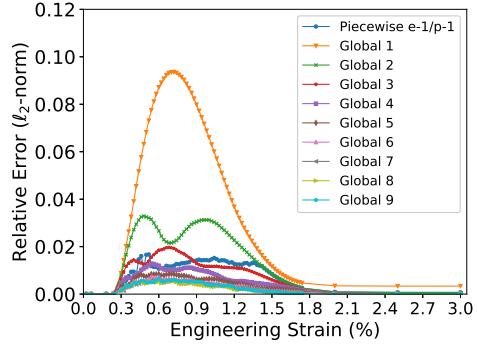
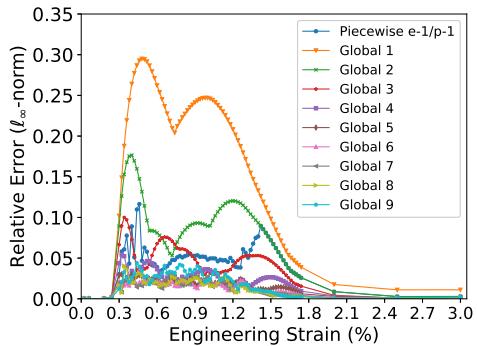


Figure 5: Stress-strain curves samples used to build the surrogate for the five parameter model.

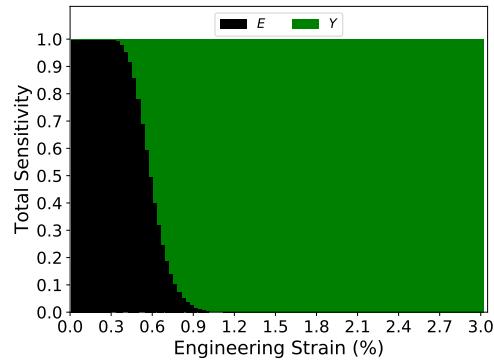


(a)

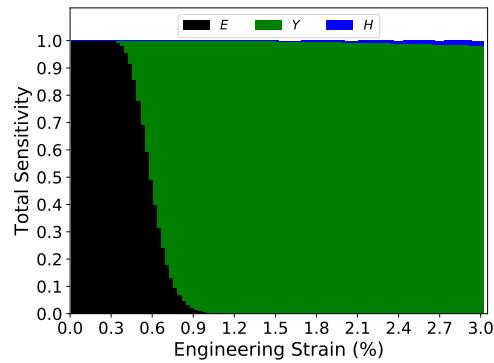


(b)

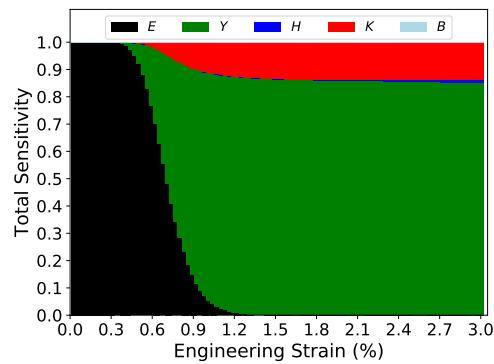
Figure 6: Surrogate error for the five parameter model at each target strain location based on (a) the ℓ_2 -norm and (b) the ℓ_∞ -norm. Results are shown for global PC expansions of orders up to nine, as well as the piece-wise surrogate with a linear polynomial over the elastic regime, and a quadratic within the plastic regime.



(a)



(b)



(c)

Figure 7: Total sensitivities as a function of the strain obtained for the (a) two, (b) three, and (c) five parameter model.

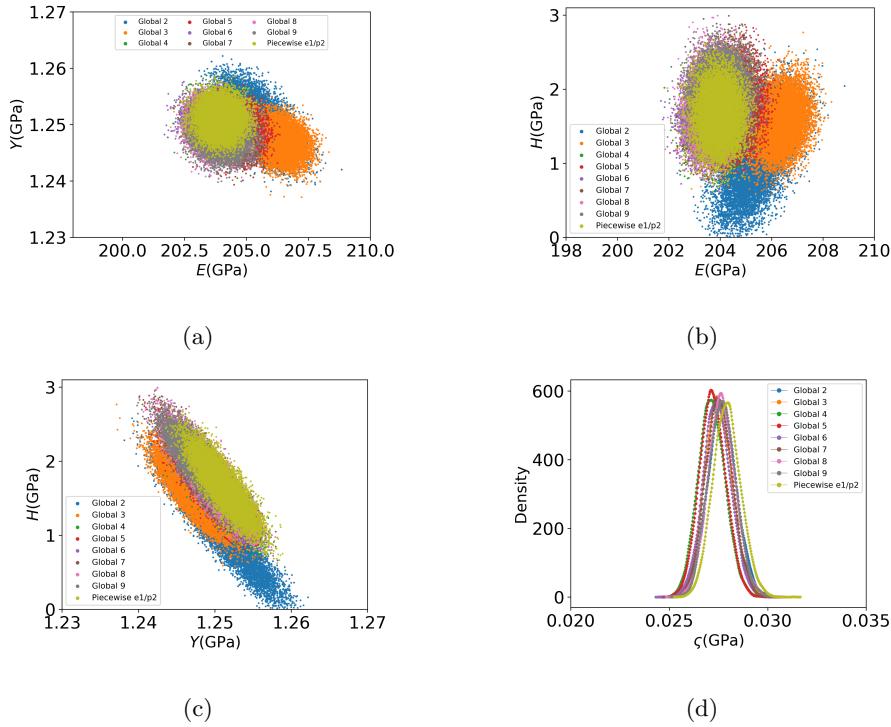


Figure 8: Results for the additive-model-based inversion run with the three parameter model using 10 stress-strain curves from \mathcal{D}_3 showing the effect of the surrogate. We plot samples of the joint posteriors: (a) $p(E, Y)$, (b) $p(E, H)$, and (c) $p(Y, H)$, as well as (d) the marginalized posterior for the standard deviation of the measurement error, $p(\zeta)$.

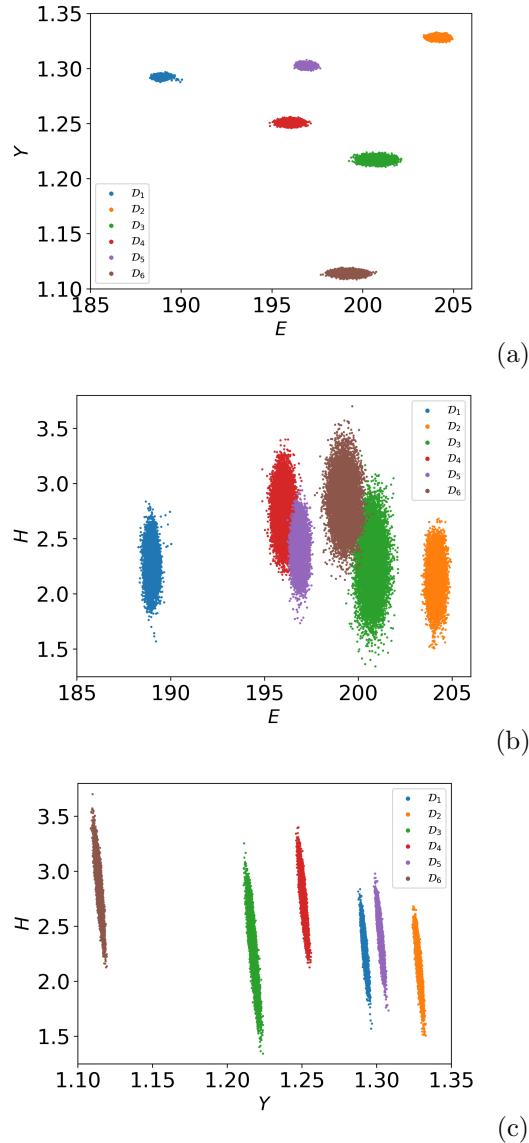


Figure 9: Batch-to-batch comparison of the joint posteriors: (a) $p(E, Y)$, (b) $p(E, H)$, and (c) $p(Y, H)$ (c) resulting from additive inversion with the piecewise elastic/plastic surrogate for the three parameter model using 40 stress-strain curves for each batch.

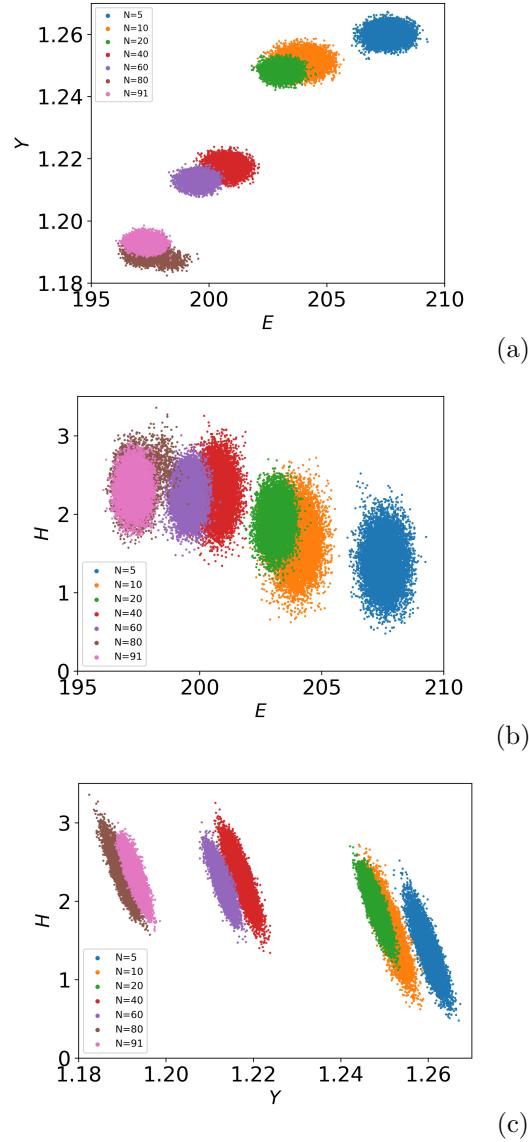


Figure 10: Convergence of posterior distributions with data: (a) $p(E, Y)$, (b) $p(E, H)$, and (c) $p(Y, H)$. Results are obtained using the additive inversion with the piece-wise elastic/plastic surrogate for the three parameter model $\tilde{M}^{(3)}$ and different number N of curves from batch \mathcal{D}_3 .

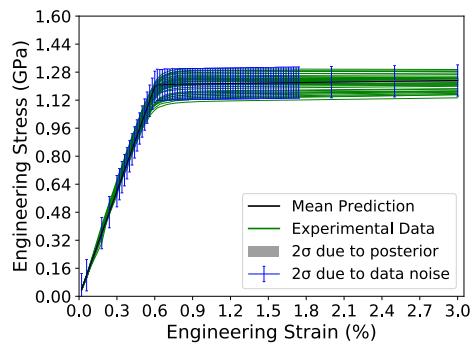


Figure 11: Posterior predictive results obtained for the additive inversion using the three parameter model along with 40 curves (green) from \mathcal{D}_3 . The blue bars represent the posterior uncertainty from the data noise, while the black solid line represents the mean prediction and the gray band represents the ± 2 standard deviations due to posterior uncertainty.

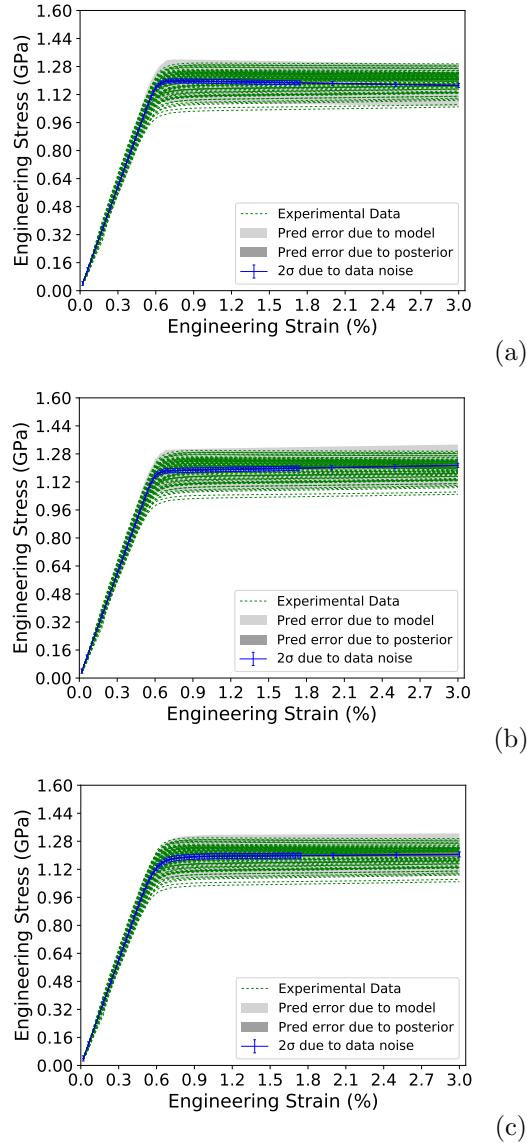


Figure 12: Posterior predictive results obtained with the embedded model-based inversion using the (a) 2, (b) 3, and (c) 5-parameter models using 91 curves (green) from \mathcal{D}_3 . The blue bars represent the posterior uncertainty from data noise, the black solid line represents the mean prediction, the light gray band represents the ± 2 standard deviations due to model inadequacy and the dark gray band is the posterior uncertainty.

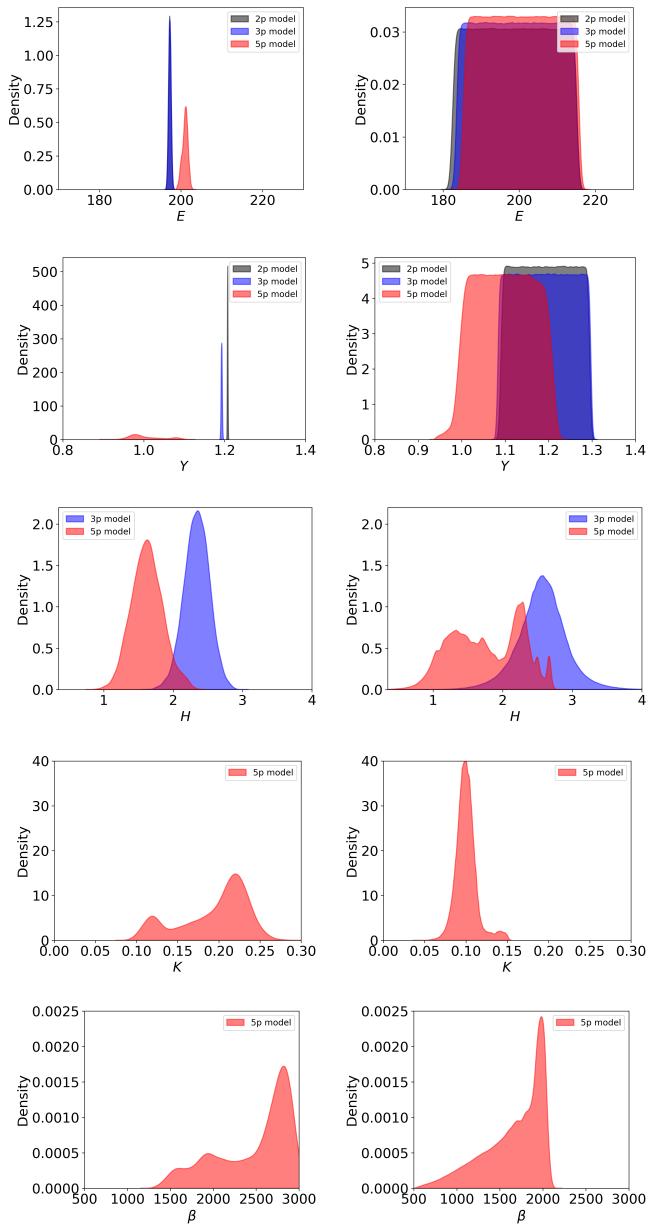


Figure 13: Comparison of posterior PDFs of the material parameters obtained using the additive (left) and embedded (right) model inference using the data from \mathcal{D}_3 .

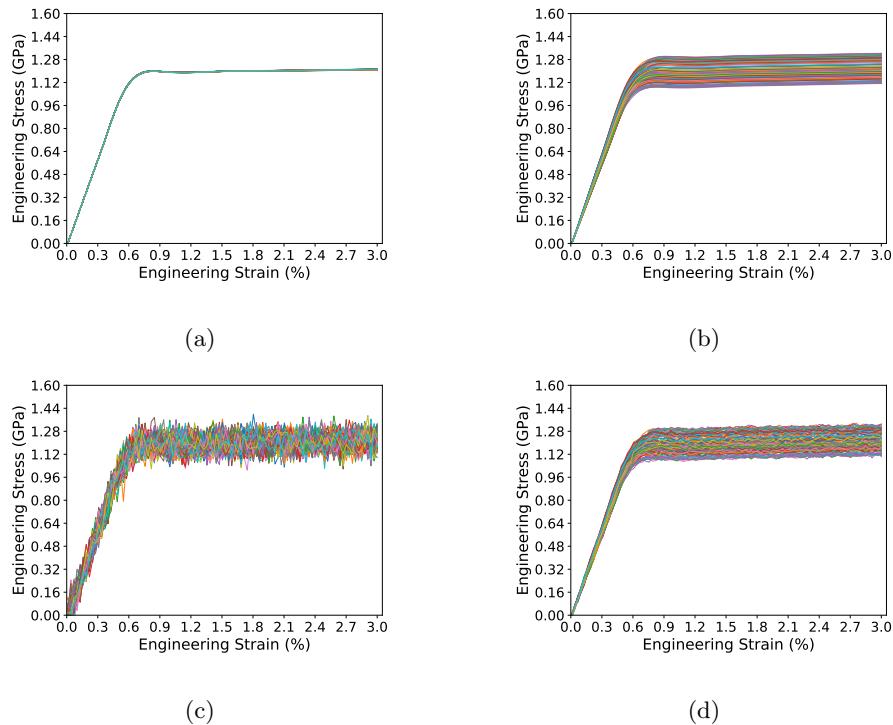


Figure 14: Sample posterior predictive realizations obtained using the 5 parameter model for the additive (left) and embedded (right) approach calibrated using \mathcal{D}_3 . The top row shows the results obtained by sampling the joint posterior density and pushing these samples through the model only, while the bottom row shows the results with the contribution of measurement error.

675 **Appendix A. Polynomial chaos expansion**

A polynomial chaos expansion (PCe) is a spectral representation of a random variable. Here we provide a brief description of the PCe construction and refer readers to Refs. [31, 32, 33] for more details. A PCe representation of any real-valued random variable λ with finite variance is an expansion of the form

$$\lambda = \sum_{|I|=0}^{\infty} \lambda_I \Psi_I(\xi_1, \xi_2, \dots), \quad (\text{A.1})$$

where ξ_I are independent identically distributed (i.i.d.) standard random variables, λ_I are the coefficients, $I = (I_1, I_2, \dots) \forall I_j \in \mathbb{N}_0$ is an infinite-dimensional multi-index, $|I| = I_1 + I_2 + \dots$ is the ℓ_1 norm, and Ψ_I are multivariate normalized orthogonal polynomials written as products of univariate orthonormal polynomials:

$$\Psi_I(\xi_1, \xi_2, \dots) = \prod_{j=1}^{\infty} \psi_{I_j}(\xi_j). \quad (\text{A.2})$$

The basis functions ψ_{I_j} are polynomials of order I_j in the independent variable ξ_j orthonormal with respect to the probability density $p(\xi_j)$. For instance, if the germ ξ is a standard Gaussian random variable, then the PCe is based on Hermite polynomials. Different choices of ξ_j and ψ_m are available via the 680 generalized Askey family [32]. The PCe (A.1) converges to the true random variable λ in the mean-square sense [34].

For computational purposes, the infinite dimensional expansion (A.1) must be truncated:

$$\lambda = \sum_{I \in \mathcal{I}} \lambda_I \Psi_I(\xi_1, \xi_2, \dots, \xi_{n_s}), \quad (\text{A.3})$$

where \mathcal{I} is some index set, and n_s is some finite stochastic dimension that typically corresponds to the number of stochastic degrees of freedom in the system. For example, one popular choice for \mathcal{I} is the *total-order* expansion of 685 degree p , where $\mathcal{I} = \{I : |I| \leq p\}$, see *e.g.* Ref. [33].

Given the expansion for the input $\lambda(\xi)$, the PCe for a target quantity of interest Q produced by the model evaluation $Q = M(\lambda)$ can be written in a

similar form

$$Q(\xi) = M(\lambda(\xi)) = \sum_{I \in \mathcal{I}} q_I \Psi_i(\xi_1, \xi_2, \dots, \xi_{n_s}). \quad (\text{A.4})$$

Methods to compute PC coefficients are broadly divided into two groups, namely intrusive and non-intrusive [33]. The former involves substituting the expansions into the governing equations, and applying orthogonal projection to the resulting equations, resulting in a larger and modified system for the PCe
690 coefficients. This approach is applicable when one has access to the full forward model and can readily modify the governing equations in the simulator. The other, non-intrusive approach is more generally applicable and involves finding an approximation in the subspace spanned by the basis functions by evaluating the original model many times.

One such non-intrusive method relies on orthogonal projection of the solution

$$q_I = \mathbb{E}[M(\lambda)\Psi_I] = \int_{\Xi} M(\lambda(\xi))\Psi_I(\xi)p(\xi) d\xi. \quad (\text{A.5})$$

and is known as non-intrusive spectral projection (NISP). In general, this integral must be estimated numerically. An alternative method of non-intrusively obtaining PCe coefficients is regression, which involves solving the linear system:

$$\underbrace{\begin{bmatrix} \Psi_{I^1}(\xi^{(1)}) & \dots & \Psi_{I^K}(\xi^{(1)}) \\ \vdots & & \vdots \\ \Psi_{I^1}(\xi^{(K)}) & \dots & \Psi_{I^K}(\xi^{(K)}) \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} q_{I^1} \\ \vdots \\ q_{I^K} \end{bmatrix}}_c = \underbrace{\begin{bmatrix} M(\lambda(\xi^{(1)})) \\ \vdots \\ M(\lambda(\xi^{(K)})) \end{bmatrix}}_M, \quad (\text{A.6})$$

695 where Ψ_{I^n} is the n th basis function, q_{I^n} is the coefficient corresponding to that basis, and $\xi^{(m)}$ is the m th regression point. In the regression matrix \mathbf{A} each column corresponds to a basis element and each row corresponds to a regression point from the training set.

Appendix B. Embedded discrepancy

As discussed in Ref. [27], the embedded discrepancy likelihood often involves highly nonlinear and near-degenerate features, thus forcing one to find an alternative way to approximate it in a computationally feasible manner. Sargsyan *et*

al. [27] suggest several options based on the assumption of conditional independence between the data points. In this work, we rely on the Gaussian approximation to the marginalized likelihood, which for the i -th batch \mathcal{D}_i , can be written as

$$p(\mathcal{D}_i | \boldsymbol{\alpha}, \tilde{M}) = \frac{1}{(2\pi)^{\frac{N_\varepsilon n_\varepsilon}{2}}} \prod_{j=0}^{n_\varepsilon-1} \prod_{k=1}^{N_i} \frac{1}{\varsigma_j(\boldsymbol{\alpha})} \exp \left[-\frac{(\mu_j(\boldsymbol{\alpha}) - \sigma_j^{(i,k)})^2}{2\varsigma_j^2(\boldsymbol{\alpha})} \right], \quad (\text{B.1})$$

where

$$\mu_j(\boldsymbol{\alpha}) \equiv \mu(\varepsilon_j; \boldsymbol{\alpha}) = \mathbb{E}_{\boldsymbol{\xi}}[\tilde{M}(\varepsilon_j; (\boldsymbol{\theta} + \hat{\boldsymbol{\eta}})(\boldsymbol{\xi}))] \quad (\text{B.2})$$

and

$$\varsigma_j^2(\boldsymbol{\alpha}) \equiv \varsigma^2(\varepsilon_j; \boldsymbol{\alpha}) = \mathbb{V}_{\boldsymbol{\xi}}[\tilde{M}(\varepsilon_j; (\boldsymbol{\theta} + \hat{\boldsymbol{\eta}})(\boldsymbol{\xi}))] \quad (\text{B.3})$$

are the mean and variance of the model at fixed $\boldsymbol{\alpha}$ and strain point. These moments are computed by constructing a PCe for the outputs by propagating the PCe of the input argument in Eq. (17):

$$\tilde{M}(\varepsilon; \boldsymbol{\theta} + \hat{\boldsymbol{\eta}}) = \tilde{M} \left(\varepsilon; \sum_I \alpha_I \Psi_I(\boldsymbol{\xi}) \right) \approx \sum_I \tilde{M}_I(\varepsilon; \boldsymbol{\alpha}) \Psi_I(\boldsymbol{\xi}). \quad (\text{B.4})$$

This can be done using NISP mentioned in Appendix A together with quadrature, and the moments can be computed from the expansion coefficients as

$$\mu(\varepsilon; \boldsymbol{\alpha}) \approx \tilde{M}_0(\varepsilon; \boldsymbol{\alpha}) \quad \text{and} \quad \varsigma^2(\varepsilon; \boldsymbol{\alpha}) \approx \sum_{I \neq 0} \tilde{M}_I(\varepsilon; \boldsymbol{\alpha}). \quad (\text{B.5})$$

After obtaining $\boldsymbol{\alpha}$ using Bayesian calibration and the likelihood just discussed the model can be used in a predictive manner. Let $\phi(\varepsilon; \boldsymbol{\alpha}) = \tilde{M}(\varepsilon; (\boldsymbol{\theta} + \hat{\boldsymbol{\eta}})(\boldsymbol{\xi}))$ be the probabilistic prediction of the model for a fixed $\boldsymbol{\alpha}$. We remark that even if $\boldsymbol{\alpha}$ is fixed, the prediction $\tilde{M}(\varepsilon; (\boldsymbol{\theta} + \hat{\boldsymbol{\eta}})(\boldsymbol{\xi}))$ remains probabilistic because of the additional variability of the random variable $\boldsymbol{\theta}$ stemming from its PCe with $\boldsymbol{\xi}$. We can then inspect the *posterior predictive* random variable $\phi(\varepsilon; \boldsymbol{\alpha})$. The posterior predictive random variable has the following mean

$$\phi_{\text{mean}}(\varepsilon) = \mathbb{E}_{\boldsymbol{\alpha}}[\mu(\varepsilon; \boldsymbol{\alpha})] \quad (\text{B.6})$$

and variance

$$\phi_{\text{var}}(\varepsilon) = \underbrace{\mathbb{E}_{\boldsymbol{\alpha}}[\varsigma^2(\varepsilon; \boldsymbol{\alpha})]}_{\text{model error}} + \underbrace{\mathbb{V}_{\boldsymbol{\alpha}}[\mu(\varepsilon; \boldsymbol{\alpha})]}_{\text{posterior uncertainty}}, \quad (\text{B.7})$$

700 where $\mathbb{V}_{\boldsymbol{\alpha}}$ denotes variance with respect to the posterior distribution of $\boldsymbol{\alpha}$. Here $\varsigma^2(\boldsymbol{\alpha})$ is the forward-propagated variance of the function $\tilde{M}(\varepsilon; (\boldsymbol{\theta} + \hat{\boldsymbol{\eta}})(\boldsymbol{\xi}))$ at location ε for given $\boldsymbol{\alpha}$. The contribution of model error and posterior uncertainty are identified and separated leveraging the law of total variance. The uncertainty due to model error is independent of how much data we use during the inference process and, thus, it can be only improved by refining the model and/or its accuracy. On the other hand, the posterior uncertainty tends to shrink with more data. Note that, in practice, the posterior distribution is described via samples, therefore the expectation ($\mathbb{E}_{\boldsymbol{\alpha}}$) and the variance ($\mathbb{V}_{\boldsymbol{\alpha}}$) are computed via Monte-Carlo integration using MCMC samples.