



AFRL-RX-WP-TR-2021-0134

ENSEMBLE PREDICTIONS OF MATERIAL BEHAVIOR FOR ICMSE FOR ADDITIVE STRUCTURES

**Stephen R. Niezgoda and Denielle Ricciardi,
The Ohio State University
Department of Materials Science and Engineering
Department of Mechanical and Aerospace Engineering**

**24 AUGUST 2021
Final Report**

Distribution Statement A. Approved for public release: distribution is unlimited.

**AIR FORCE RESEARCH LABORATORY
MATERIALS AND MANUFACTURING DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7750
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

Qualified requestors may obtain copies of this report from the Defense Technical Information Center (DTIC) (<http://discover.dtic.mil>).

AFRL-RX-WP-TR-2021-0134 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

SPARKMAN.DANIEL.M
URRAY.1384350808

Digitally signed by
SPARKMAN.DANIEL.MURRAY.13843
50808
Date: 2022.03.07 10:14:38 -05'00'

DANIEL SPARKMAN
Work Unit Manager
Materials State Wareness Branch
Structural Materials Division

JONES.ERIC.L.1
299413990

Digitally signed by
JONES.ERIC.L.1299413990
Date: 2022.03.07 11:31:35
-05'00'

ERIC JONES
Branch Chief
Materials State Awareness Branch
Structural Materials Division

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 24 August 2021		2. REPORT TYPE Final Report		3. DATES COVERED (From – To) 24 March 2017 – 24 June 2021	
4. TITLE AND SUBTITLE ENSEMBLE PREDICTIONS OF MATERIAL BEHAVIOR FOR ICMSE FOR ADDITIVE STRUCTURES				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA8650-17-1-5277	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Stephen R. Niezgoda Denielle Ricciardi				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER X1CK	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Ohio State University Office of Sponsored Programs 1960 Kenny Road Columbus OH 43210-1016				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory Materials and Manufacturing Directorate Wright-Patterson Air Force Base, OH 45433-7750 Air Force Materiel Command United States Air Force				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RXCA	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RX-WP-TR-2021-0134	
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release: distribution is unlimited.					
13. SUPPLEMENTARY NOTES Report contains color.					
14. ABSTRACT Integrated Computational Materials Engineering (ICME) is an engineering approach whereby model linkages as well as experimental and computational integration are exploited in order to efficiently explore materials processing-performance relationships. Many materials models and simulations are deterministic in nature. Achieving a statistical confidence in a simulation output requires, first, the identification of the various sources of error and uncertainty affecting the simulation results. Statistical inference can then be used to recover information about unknown model parameters by conditioning on available data while taking into account the various sources of uncertainty. A statistical random effects hierarchical framework was developed and demonstrated. First for parameter estimation and response prediction in a phenomenological viscoplastic self-consistent (VPSC) crystal plasticity model [1]. Inference is performed under two different scenarios: 1) with the consideration of model discrepancy, modeled through a Gaussian process and 2) without the consideration of model discrepancy. Second, uncertainty quantification and propagation is demonstrated for calculations of phase equilibrium for the binary AgCu system. A Metropolis-Hastings Markov chain Monte Carlo algorithm is used to estimate model parameters with a quantified level of uncertainty, which induces a probability distribution in the simulation output, and then is quantified using posterior summaries.					
15. SUBJECT TERMS ICMSE, material models, model aggregation, model uncertainty					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 147	19a. NAME OF RESPONSIBLE PERSON (monitor) Daniel Sparkman
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUBER (include area code) (937) 255-1340

Table of Contents

LIST OF FIGURES	ii
LIST OF TABLES	v
1: SUMMARY	1
2: INTRODUCTION	4
3: METHODS, ASSUMPTIONS, AND PROCEDURES.....	25
4: RESULTS AND DISCUSSION.....	73
5: CONCLUSIONS	97
6.0 REFERENCES.....	99
APPENDIX A.....	115
APPENDIX B	121
APPENDIX C	126
LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS	139

LIST OF FIGURES

<u>Figure</u>	<u>Page</u>
Figure 2 1: Reliability-based design vs deterministic design [47]	9
Figure 2 2: Aleatoric vs epistemic uncertainty [2]	11
Figure 2 3: Frequentist vs Bayesian inference. θ is the true value and $\hat{\theta}$ is the 'best guess' provided by the estimator [60].	14
Figure 2 4: Schematic of a Metropolis-Hastings Markov chain Monte Carlo simulation from [66].....	16
Figure 2 5: Schematics comparing the (a) fixed effects (left) and (b) random effects (right) models [71].....	18
Figure 2 6: Various sources of uncertainty present in computational models that must be accounted for to perform reliable inference and prediction.	23
Figure 3 1: Visualization of a Bayesian random effects model. Arrows represent the direction of conditional dependence. Clear nodes represent unobserved/unknown quantities while shaded nodes indicate observations.....	31
Figure 3 2: Calibration data generated via dislocation-based EVP model for the stress- strain response up to 9% strain of miniature copper samples under tension. Six experimental data sets were created to simulate sample-to-sample and experimental variability	36
Figure 3 3: (a) Diagram illustrates the data generation and computation of posterior distribution for the conventional Bayesian problems. The dashed directed lines indicates the model executions. (b) Visualization of data generation and derivation of posterior distribution for expensive simulators. The grey elements indicate difference from the conventional Bayesian inference problems.	43
Figure 3 4: The upper left plot shows the model outputs from eight training calibration parameter values (solid curves) and the true process (blue dashed curve) generated in the case study. The upper right plot is the basis functions constructing the basis spline space. The bottom plots show five basis coefficient of eight training data (circles) and the true process (star).	47
Figure 3 5: Plots of Matern covariance function with $\nu = 3/2$, squared exponential covariance function and exponential covariance function where the x-axis is the Euclidean distance of two input points.	47
Figure 3 6: Bayesian hierarchical model of calibration for expensive simulators. Directed connection between nodes represent the direction of conditional dependence. The shaded nodes are the observed data and the clear nodes are the unknown quantities.....	48
Figure 3 7: Plots illustrate the full conditional distributions of basis coefficients. The blue curves are the mean of full conditional distribution and the red and yellow curves represents the 95% credible bounds.	50
Figure 3 8: Visualization of the Bayesian random effects model with discrepancy. Arrows represent the direction of conditional dependence. Clear nodes represent unobserved/unknown quantities while shaded nodes indicate observations.	56
Figure 3 9: Multiple sample functions drawn from a GP prior with a smooth covariance structure. The cases shown are a) a stationary covariance function and b) a non-stationary covariance function.....	58
Figure 3 10: Multiple sample functions drawn at random from the prior specified by a GP which favors smooth functions and a stationary covariance function with parameters a) $w=.025, \sigma^2 = 10$ b) $w=.25, \sigma^2 = 10$, c) $w = .025, \sigma^2 = 100$, d) $w = .25, \sigma^2 = 100$	59
Figure 3 11: Schematic of the toy problem showing the true behavior ζ (blue line), the model of the behavior $m(t,\theta)$ (magenta line) and the observations y_n (open circles).....	60

Figure 3 12: A comparison of inference on the toy example without (a & b) and with (c & d) the contribution of model misfit.....	62
Figure 3 13: Analysis for VPSC conditioned on the elastic + plastic regimes of deformation not accounting for model discrepancy. The red points are observations and the black lines show approximately 200 samples from the posterior distribution.	64
Figure 3 14: The black dots are the experimental data. The red line is the basis spline and the blue line is the VPSC output where the corresponding iteration generates the mode posterior. The subplot is the a part of the plot zoomed in.....	68
Figure 3 15: Visualization of the Bayesian random effects model with discrepancy. Arrows represent the direction of conditional dependence. Clear nodes represent unobserved/unknown quantities while shaded nodes indicate observations.	71
Figure 3 16: Schematic of Metropolis-Hastings algorithm	72
Figure 4 1: Correlation plots for a) a random effect and b) the overall effect. Marginal posterior (black) and prior (red) density estimates for each parameter are shown on the diagonal; bi-variate posterior (colors by intensity) and prior (red) contours	74
Figure 4 2: Posterior marginal contours for a) (τ_1, ξ_1) and b) (τ_0, ξ_0). Solid contours are used for random effects; dashed contours are used for overall effects.	75
Figure 4 3: 95% HPI (outer dashed lines) and posterior mean (central dashed line) for the posterior predictive distribution. 200 sampled predictive stress-strain curves (grey lines) are shown, and the MAP VPSC evaluation (solid red line).....	76
Figure 4 4: Correlation plots for a) a random effect and b) the overall effect. The diagonal subplots contain the marginal histogram for each parameter with the overall parameter histograms being plotted with each respective prior density (red line).	79
Figure 4 5: Contour Plot of estimated bi-variate distribution between L0 and L1. The random effect contours are shown with solid lines and the overall effect countours are shown with dashed lines. Note that the contours of the random effect in the bottom left corner overlap.....	80
Figure 4 6: Forward model evaluations based on 530 samples from the posterior distribution for the Gibbs free energy as a function of atomic fraction copper at 1000K: a) the entire range of 0-1 atomic fraction copper b) a magnified range of 0-.1 atomic fraction copper is shown. This distribution can be used to propagate parameter uncertainty to phase diagrams.	81
Figure 4 7: Random draws from the prior of the discrepancy in a) the elastic region and b) the plastic region c) over the full domain. Note the difference in the scale of the vertical axes. .	82
Figure 4 8: Marginal posterior histograms and parameter correlations for a) a representative random effect and b) the overall effect,with prior draws included in red for the overall effect.	83
Figure 4 9: Marginal posterior contours for a) (τ_1, ξ_1) and b) (τ_0, ξ_0). Random effect marginal contours are indicated by solid lines and the overall effect contours are indicated by a dashed line.	84
Figure 4 10: The mean posterior draw (blue), the VPSC model evaluated at the mean parameter values (red) and the mean discrepancy (green) are shown in a) with an inset showing the elastic region. 200 draws from the a) posterior and b) posterior predictive distributions are shown with the mean and 95% highest posterior intervals shown.	86
Figure 4 11: Stress-strain response predictions (a) and posterior draws (b) up to 50% strain under ideal rolling boundary conditions and a ‘copper’ initial texture. Samples of size 200 are shown.	90

Figure 4 12: Histograms showing the propagated uncertainty in atomic fraction Cu (xcu) for phase boundaries a) α , $\alpha + \beta$ and b) $\alpha + \beta$, β . The outer and middle dashed lines (red) bound the 95% HPI and mean posterior value, respectively. As comparison, boundaries reported by other sources [171–173, 169, 174] are plotted as well.	92
Figure 4 13: The posterior histogram and correlation contour plot of the posterior distributions.	95
Figure 4 14: The upper plots show the basis coefficients of the training data (dots) and of the true process (blue star). The bottom plot is the true projected spline (red) and 95% point-wise credible bound of basis splines (blue).	96
Figure A 1: Scaled marginal posterior densities from case study I for random effects (solid lines) and the overall effect (dashed lines) for each Voce hardening parameter.....	119
Figure A 2: Scaled marginal histograms for Redlich-Kister parameters at 1000K. The parameters are L0 in (a) and L1 in (b) and include all random effects (with 'o' markers) as well as the overall effect parameters (dashed lines).	120
Figure A 3: Marginal histogram for the error precision plotted with its prior density distribution (red line) for case study I.....	120
Figure B 1: Scaled marginal posterior densities from case study I for random effects (solid lines) and the overall effect (dashed lines) for each Voce hardening parameter.....	125

LIST OF TABLES

<u>Table</u>	<u>Page</u>
Table 2 1: Multi-scale modeling in material mechanics [2].....	6
Table 3 1: Prior hyperparameters for case study I with corresponding prior means and variances.	38
Table 3 2: Prior hyperparameters for case study II with corresponding prior means and variances.	41
Table 3 3: Prior hyperparameters for the case study with corresponding prior means and variances.	66
Table 4 1: Posterior summaries for the overall effect parameters and error precision.	76
Table 4 2: Posterior summaries for the overall effect parameters and error precision at 1000K....	80
Table 4 3: Posterior summaries for the overall effect parameters and error precision.	84
Table 4 4: Expected values, variances, and 95% highest posterior density intervals of the phase boundaries in terms of atomic fraction of Cu in the Binary AgCu system.....	91
Table 4 5: The summary table of posterior distributions.....	95
Table A 1:Element-wise posterior correlation for a random effect from case study I	117
Table A 2: Element-wise posterior correlation for the overall effect from case study I.....	117
Table A 3: Element-wise posterior correlation for the random effect from case study II	118
Table A 4: Element-wise posterior correlation for the overall effect from case study II	118
Table B 1: Element-wise posterior correlation for a representative random effect from case study I.....	124
Table B 2: Element-wise posterior correlation for the overall effect from case study I.....	124

1: SUMMARY

In the past several decades, there has been an unprecedented demand for the discovery and design of new materials to support rapidly advancing technology. This demand has fueled a push for Integrated Computational Materials Engineering (ICME), an engineering approach whereby model linkages as well as experimental and computational integration are exploited in order to efficiently explore materials processing-performance relationships. Tailored simulations allow for the reduction of expensive and lengthy experiments, emphasizing the need to establish a statistical confidence in component designs and manufacturing processes from the simulations, rather than experiments, in a principled way. Since many materials models and simulations are deterministic in nature, the use of sophisticated tools and techniques are required.

Achieving a statistical confidence in a simulation output requires, first, the identification of the various sources of error and uncertainty affecting the simulation results. These sources include machine and user error in collecting calibration data, uncertain model parameters, random error from natural processes, and model inadequacy in capturing the true material property or behavior. Statistical inference can then be used to recover information about unknown model parameters by conditioning on available data while taking into account the various sources of uncertainty.

In this work, Bayesian inference is used to quantify and propagate uncertainty in simulations of material behavior. More specifically, a random effects hierarchical framework is used since it provides a way to account for uncertainty stemming from random natural processes or conditions. This is especially important in many materials modeling applications where the random microstructure plays an important role in dictating material behavior. In addition to this, in many cases experiments are quite costly, so in order to obtain sufficient data for calibration, a compilation of data from multiple sources (i.e., labs, publications, experimental techniques) may be necessary. As a result, there is not only error within each experiment due to noise, resolution or machine/user error, but also variations in the data from sample-to-sample (or experiment-to-experiment). The random effects model allows for the consideration of this variability that occurs between samples to robustly calibrate model parameters and predict material behavior.

The statistical framework is first demonstrated for parameter estimation and response prediction in the phenomenological viscoplastic self-consistent (VPSC) crystal plasticity model [1]. Inference is performed under two different scenarios: 1) with the consideration of model discrepancy, modeled through a Gaussian process and 2) without the consideration of model discrepancy. Second, uncertainty quantification and propagation is demonstrated for calculations of phase equilibrium for the binary AgCu system. In all cases, uncertainty due to uncertain model inputs, observation error, as well as sample- to-sample variability are considered. A Metropolis-Hastings Markov chain Monte Carlo algorithm is used to estimate model parameters with a quantified level of uncertainty. This uncertainty is then propagated to induce a probability distribution in the simulation output and quantified using posterior summaries such as the mean, mode, and highest posterior intervals (HPI).

Uncertainty Quantification for Parameter Estimation and Response Prediction: Generalizing the Random Effects Bayesian Inferential Framework to Account for Material and Experimental Variability

Integrated Computational Materials Engineering (ICME) is an engineering approach where the materials, manufacturing process, and component designs are optimized concurrently before an actual physical component is realized. This requires the integration of models across vast length and timescales. A key benefit of ICME is the ability to reduce the bulk of expensive and lengthy experiments via tailored simulation. However, ICME introduces new challenges and limitations as the statistical confidence in the final design and manufacturing process must be established from simulation rather than experimental observation and testing. The computational materials science community has not formally adopted tools for verification and validation or UQ for materials simulations. In this study, a Bayesian hierarchical model is considered which accounts for parameter uncertainty, the inherent variability in the properties of material samples tested, and measurement noise. The Bayesian inferential framework is used to calibrate model parameters given calibration data and also to make forward predictions with a confidence level established through highest posterior density intervals. The generality of the framework is demonstrated through two case studies: 1) parameter estimation for a crystal plasticity (CP) model which provides key microstructural and grain-level deformation information to use within the ICME chain 2) the estimation in uncertainty in thermodynamic phase stability from multiple databases for phase stability.

Uncertainty Quantification Accounting for Model Discrepancy Within a Random Effects Bayesian Framework

In materials modeling, there is a great variety of models with an impressive breadth of fidelity available by which to predict material behaviors. With the growing emphasis on establishing the predictive confidence of these models in a principled way, there is a benefit in using phenomenological or empirical models which are computationally well-suited for statistical simulations. Since these models are often missing physics or have incomplete knowledge built into their mathematical construct, they may not capture the physical system or process adequately over the entire domain. This can lead to a systematic discrepancy, otherwise known as model misfit, between the model output and the system it represents over at least part of the domain. However, by accounting for this discrepancy in uncertainty analyses, reliable inference and prediction on the model parameters and material behavior can be made, despite the missing physics. The previous statement is contingent on two conditions: 1) The discrepancy is systematic, and 2) There is a strong prior knowledge on the structure of the discrepancy, which is required to minimize issues with non-identifiability among unknown model components. This work uses a random effects Bayesian hierarchical framework to perform inference and prediction in a phenomenological crystal plasticity model. Uncertainty information about unknown model parameters is recovered by conditioning on data and by taking into account uncertainty stemming from aleatoric sample-to-sample variability, observation error as well as model form error. The uncertainty quantification framework is demonstrated on inference for the phenomenological meso-scale VPSC crystal plasticity model, which does not contain physics describing the elastic regime of deformation. Inference is performed via a Bayesian approach enabled by posterior simulation. Posterior uncertainty in the unknown Voce hardening parameters is propagated to generate a distribution of the stress-strain response in both the elastic and plastic regimes. Additionally, posterior predictive distributions are simulated to establish uncertainty bounds for future unobserved data.

Bayesian Computer Model Emulation Using a Gaussian Process

Computer models are widely used to represent materials phenomena and analyze underlying structures of real-world systems. These models are often computationally expensive or inaccessible, making direct

uncertainty quantification impractical. This paper presents a Bayesian inference method which performs calibration and uncertainty quantification of complex materials models at a low computation cost. The method calibrates a materials model while treating the code as a ‘black box’ through a response surface model. The input-output relationship of the deterministic function is identified through a small number of physical model evaluations. Within a Bayesian hierarchical framework, we calibrate and quantify the uncertainty of input parameters. The Bayesian calibration method is demonstrated through the VPSC crystal plasticity model.

2: INTRODUCTION

The reliable use and effectiveness of these simulations depends on the accurate assessment of the sources of error, or uncertainties, which affect their outcome. It is a seemingly obvious but often ill-addressed fact that no model, and consequently simulation, is perfect. There are gaps between the model and the physical system it aims to imitate, gaps in the understanding of the actual physical system or process, and an inability to solve some models in a closed form. These factors all contribute to the various sources of uncertainty within the simulation which influence the prediction of the system. An understanding of how these sources affect the prediction can be obtained through mature tools developed through mathematical probability theory and statistical analyses which have been successfully applied in many fields [2]. Such tools include model verification and validation, which ensure the math and physics, respectively, of a model are correct, as well as uncertainty analysis.

Uncertainty analysis, which encompasses methods for uncertainty quantification (UQ), mitigation (UM), and propagation (UP) [3], and often collectively termed uncertainty quantification, is an integral component of model verification and validation (V&V). It facilitates a rigorous process for ensuring the models in use have proper fidelity and a quantified confidence level. UQ is concerned with identifying all possible sources of error within the simulation and quantifying their effect on the simulation predictions. UM uses techniques to minimize the effect the sources of error have on the model output, and UP is concerned with projecting uncertainties to other models or to situations outside the bounds of calibration.

Within the scope of materials science and engineering, advancements in technology throughout history have been enabled by the discovery and design of new materials which have specific properties and novel manufacturing processes to produce these materials. However, as noted by Arro'yave [4], this is often also the bottleneck in the advancement of new technology as well. With the rate of technological discoveries and advancements ever increasing, there is a recognized importance for the acceleration of the design and discovery of new materials, emphasizing the importance of computer simulations to meet these goals. In response to this demand, over the past several decades, an emphasis has been made on Integrated Computational Materials Engineering (ICME) [5] with the aim to describe the processing → microstructure → properties → performance relationships of materials across various length and time scales. This paradigm integrates science and engineering with theory, experiment and simulations [6] to develop new materials and processes while eliminating the bulk of repetitive and time-consuming experiments traditionally required to bring new materials to the market.

A proper understanding of the multiple sources of error present, effective mitigation of these errors, and thorough analyses on their propagation throughout the simulation is critical in order to reliably develop and deploy new and improved materials. This information is needed for rigorous risk-assessment and decision-making to determine the appropriateness of new materials for design purposes. Unfortunately, materials modeling is largely deterministic and does not consider models or unknown quantities to be stochastic in nature, effectively ignoring the uncertainty which is present. It is often the case that a model is used under the assumption the model is correct, the model input parameters and system variables are known, and the data collected for calibration represent the true underlying behavior of the system. This is far from the

truth when, in fact, models are imperfect representations of the systems they are targeting, and the data collected for calibration are often sparse due to the cost of running experiments and inconsistent from test to test with additional error introduced by the machine and user. Furthermore, the practice of calibrating a model to the average of multiple experiments with the belief it will yield average input parameters and an average prediction is dangerous for highly non-linear models, which comprises the majority of models used in the computation of material properties. This is due to the fact that a non-linear transformation of the average input parameter values does not necessarily yield an average model output [7].

In deterministic approaches for model calibration, typically an approach such as a regression method is used to determine the model and set of model parameters which minimizes the residual sum of the square differences between observations and the model output [8]. This provides a single best estimate for the unknown model parameters which minimizes the difference between the experimental results and the model. It is also not uncommon to perform the calibration in an ad-hoc setting, where the analyst adjusts the parameters by hand until an acceptable agreement between the model and observations is reached.

In probabilistic approaches for model calibration, the most likely model and set of parameters is determined by conditioning on the data which is available. The result is a probability distribution over the uncertain input parameters describing the likelihood that each possible parameter value is the ‘true’ value. When this parameter uncertainty is propagated to simulation output, it produces a probability distribution over the expected behavior and the likelihood of observing each of the same. This provides very valuable information since there are oftentimes multiple models available which aim to represent the same physical system and an infinite set of possible parameter sets to consider. Many parameter and model pairs yield nearly identical model predictions, and each have a finite probability of being the pair which best captures the system being modeled. This calibration approach uses iterative methods to converge statistically on a distribution of possible parameters. This is in contrast to deterministic calibration approaches which ignore all but one possible model/parameter pair and model output. Probabilistic approaches are very beneficial in that they allow the analyst to account for a multitude of error sources affecting the material simulation and to additionally propagate this uncertainty to predictions outside the range of calibration or to linked models. These techniques aid in the assessment of simulation accuracy and reliability and have been gaining increasing attention within the materials computational community. There are a multitude of stochastic procedures for calibration and uncertainty analysis; a few popular techniques within the materials computational community involve machine learning techniques such as Gaussian processes, Bayesian inference, and polynomial chaos expansion.

The work performed for this project sought to address the problem of correctly quantifying and propagating sources of error in materials modeling using the random effects Bayesian inferential framework. Furthermore, a generalization of the procedure is presented in an attempt to make the complex analysis more accessible to the community at large and easily adaptable to different problems of interest to materials scientists. This framework accounts for the discussed sources of error in the parameter estimation and summarizes the uncertainty through posterior functionals such as quantiles, and highest posterior intervals. It also provides a method to propagate this uncertainty to the simulation output, or, alternatively, to other models

linked in a modeling chain. This framework allows for a straightforward extension to Bayesian model averaging, which allows us to combine models for more robust prediction.

IMPORTANCE OF UNCERTAINTY QUANTIFICATION IN MATERIALS SCIENCE & ENGINEERING

ICME

Recognizing the importance of materials computational-supported design has been on the rise and has been encouraged with the push for ICME and material development through the integration of multi-scale models and experiment. One key benefit stemming from this integration approach is a reduction in the abundance of experiments typically required to obtain a proper design confidence, minimizing the cost in both time and resources. However, with this reduction comes the challenge of obtaining the design confidence through simulations of the system or process. This is challenging since many of the models used to understand the physics and mechanisms at play in a material or system are deterministic in nature. This means for a single set of input parameters and system variables, the model will always produce identical results. Hence, there is no inherent understanding of the distribution of material responses or, equally important to reliable design and ensuring the material in use meets the performance requirements for the given application, the likelihood of each response being observed.

Another goal of ICME is the linking of models across scales, spanning many orders of magnitude in order to account for phenomena and mechanisms affecting the material response at multiple levels. An example, Table 2.1 from Chernatynskiy [2] highlights the various scales which govern the mechanical behavior of metals. This linking has the potential to provide a better understanding of material processing-performance relationships, allowing predictions on material properties to be established based on the processing settings. It is therefore possible for the inverse problem of determining the processing environment needed to achieve a desired property for a specific component and application. These links can be achieved by running the models in a hierarchical way whereby the output from one simulation becomes the input for the next simulation at the next scale. Alternatively, this can be achieved by running models concurrently within a simulation [9]. The associated challenge with this endeavor is not only the aforementioned task of reliably performing uncertainty analyses at each scale, but also propagating the uncertainty between models. This requires sophisticated tools as the model outputs are not formulated in such a way to naturally account for and propagate uncertainty [10, 3].

Table 2 1: Multi-scale modeling in material mechanics [2]

Unit	Length scale (m)	Timescale(s)	Mechanics
Complex structure	10^3	10^6	Structural mechanics
Simple structure	10^1	10^3	Fracture mechanics
Component	10^{-1}	10^0	Continuum mechanics
Grain microstructure	10^{-3}	10^{-3}	Crystal plasticity
Dislocation microstructure	10^{-5}	10^{-5}	Micromechanics
Single dislocation	10^{-5}	10^{-5}	Dislocation dynamics
Atomic	10^{-5}	10^{-5}	Molecular dynamics
Electron orbitals	10^{-5}	10^{-5}	Quantum mechanics

There is an additional challenge in selecting the best model when there are several models available to choose from, each representing the same physical phenomenon. For example, in crystal plasticity modeling there is a vast range of models available to determine the final microstructure and deformation of a sample undergoing some state of stress. The available models include those ranging from reduced-order, homogenized, phenomenological and empirical models [1, 11–13] to spatially resolved physics-based models with internal state variables tracking dislocation interactions and other physics influencing the final state of deformation [14–21]. As can be expected, the complexity of the physics incorporated into the model, as well as the resolution of the solver, has a strong influence on the efficiency and fidelity of the property prediction. In this scenario, the challenge of uncertainty analysis broadens not only from recovering and propagating uncertainty in the unknown model components, which includes parameter inference, but also to determining the best model to use to represent the system, where ‘best’ includes the combined required level of efficiency and fidelity for the application.

While there is not an agreed upon standard for UQ and UP in materials computational design, an understanding of its importance has been on the rise with Bayesian inferential techniques widely accepted for this task [22–32]. This approach involves updating prior knowledge about unknown model components, such as input parameters, by conditioning on available data in order to recover and propagate their associated uncertainty. The appropriate technique to use for UQ and UP is highly application-dependent. For example, reduced-order, homogenized, or phenomenological models are well-suited for sampling-based methods such as Markov chain Monte Carlo (MCMC) [33], which explicitly sample the input space. Thousands, or possibly millions, of model evaluations are necessary in order to obtain statistical convergence. Conversely, it is most likely the required number of model runs for MCMC simulations exceeds the computational budget available for physics-based and full-field models. Here, machine learning techniques such as Gaussian processes [34] or a surrogate model, which efficiently emulates the computer model/simulation [35], may be a more suitable choice. From here, inference and UQ is performed directly on the surrogate. For applications where there are multiple models available, the analyzer may choose to fuse the information from multiple models using multi-fidelity techniques [36, 37] or Bayesian model averaging [38, 39]. In multi-fidelity optimization, low and high fidelity model output from models of differing complexities are leveraged in order to maximize the accuracy of the model predictions while simultaneously reducing the cost of inference on unknown model parameters. Bayesian model averaging is a model ensemble technique where inference is performed on unknown model components as well as the models themselves to determine the probability any given model is the best. This information is then used to generate an ensemble prediction. Alternatively, model selection approaches [40] may be preferred to select the single best model for the given scenario.

In order to reliably meet the goals of ICME, which is to integrate computational tools and experiments bridging vast scales, advanced uncertainty analyses must become fully integrated into the design process. This would yield simulation predictions with a known level of fidelity with the uncertainty in model parameters accounted for and propagated to linked models to produce predictions for material properties with a known level of confidence.

MATERIALS DESIGN

Uncertainty quantification also plays an important role in the design of materials. Traditionally, the reliability of materials design has been established through repetitive experiments which is costly and time consuming. For example, Federal Aviation Administration certification of metallic structures is typically established based on “A”- basis design where at least 99% of the material population equals or exceeds minimum property values with 95% confidence [41, 42]. Slightly less stringent is “B”-basis design where at least 90% of the population is required to equal or exceed minimum property values with 95% confidence. Alternatively, with significantly less experimental demand, probabilistic risk analysis informs decision-making for determining possible failures and their level of severity, as well as the likelihood such failure will occur [43]. This analysis is critical in a reduced-experiment scenario for the robust and reliable design of materials, which must meet prescribed safety requirements for their specified purpose. Important to the support of risk analysis is sensitivity analysis, which is comprised of methods to determine the contribution of each input random variable to the overall uncertainty in the output space [44–46]. While uncertainty analysis uses techniques to determine what the overall output uncertainty is due to the uncertain inputs [44], sensitivity analysis goes one step further by measuring the strength of the relationship between specific inputs and features of the output space [2]. This allows the analyzer to target the critical model input parameters to reduce the uncertainty of in order to achieve the greatest reduction in the output uncertainty [2].

Additional tools aiding in risk analysis are reliability-based design and robust design optimization. Reliability analysis [47] is a tool for establishing the probability of failure for specific design variables. Gargama [47] provides a useful illustration in Fig. 2.1 demonstrating the value of reliability-based design optimization compared to deterministic optimization. In this illustration, the deterministic optimum of design variables X_1 and X_2 lies on the intersection of the design constraint boundary. However, once uncertainties in the design variables are considered, it is revealed the deterministic design optimum may not be feasible and does not meet the required probability of satisfying the constraint, making failure likely. Therefore, in order to achieve a safe and reliable design, the true deterministic optimum is sacrificed in order to find the reliable optimal solution in a feasible region. Details of this process can be found in [47].




Figure 2 1: Reliability-based design vs deterministic design [47]

Robust design optimization seeks to minimize the effect of propagated uncertainty to material performance in order to ensure structural safety [44]. The goal is to make the product or process as insensitive to sources of uncertainty as possible, even though the sources still exist [48]. While risk assessment and its associated strategies are not specifically considered in the work presented in this document, it underlies the importance of quantifying and propagating uncertainties in material models in a principled way.

SOURCES OF UNCERTAINTY

In order to properly account for and analyze model fidelity, clear classification of the sources and types of uncertainty is needed. The following discussion elaborates on the two primary categories, aleatoric and epistemic, and the various sources of error that are assigned these classifications.

ALEATORIC VS EPISTEMIC

Aleatoric uncertainty is due either to the physical randomness of a system or stochastic numerical algorithms. This type of uncertainty cannot be reduced but only quantified and is revealed in simulation results in the form of random error [49]. For example, aleatoric uncertainty can be observed in materials testing where the material behavior is influenced by the microstructure of the sample, such as in tensile testing. Since microstructure is stochastic in nature, repetitive tests of the same sample material would yield slight variations due to the randomness.

Epistemic uncertainty, on the other hand, is due to imperfect knowledge about the system being modeled or from diffusing incomplete knowledge into the formulation of a model. This type of uncertainty may be present if it is difficult, dangerous, or too expensive to collect sufficient observations about the system in question, resulting in poorly-characterized or sparse calibration data. Additional contributions come from simplifications or numerical methods used in the model. These contributions lead to uncertainty in calibrated model parameters as well as discrepancy between the model output and the observations [49]. However, unlike aleatoric uncertainty, epistemic sources are reducible by increasing one's knowledge about the system, incorporating more information into the model, or by improving the quantity and quality of calibration data.

A useful illustration of these different classifications, provided by Chernatynskiy [2] in Fig. 2.2, is that of a target. Aleatory uncertainty can be thought of as the precision of the shots, where a higher precision corresponds to less aleatoric uncertainty. The epistemic uncertainty corresponds to the accuracy of the shots, where shots centered around the bulls eye have less epistemic uncertainty than those not centered.




Figure 2 2: Aleatoric vs epistemic uncertainty [2]

CLASSIFICATION OF SOURCES OF ERROR

Uncertainty can further be classified as model form, model parameter, and natural uncertainty. Natural uncertainty, error in the simulation stemming from the natural process, is a form of aleatory uncertainty and, therefore, cannot be reduced but only quantified and managed. As previously described, this contributes to the model parameter uncertainty since a model calibrated to observations affected by a random process would yield slightly different optimal parameters experiment-to-experiment. Further description of the other two classifications, model form and model parameter, are discussed below.

PARAMETER

Model parameter uncertainty results from a model being calibrated against data which has inherent measurement error from either the machine, user or both, which leads to observations which have deviated from the ‘ground truth’. This introduces a bias in the calibrated parameters and contributes to the uncertainty in their true value. It is helpful to support this discussion by demonstrating how one might model the data which is observed. Suppose we have N observations of the physical system, y_n , $n = 1, \dots, N$, which we have observed at N variable inputs $x = \{x_1, \dots, x_N\}$. These observations are likely not representative of the true property values, $\zeta(x)$, and are contaminated with some additive independent Gaussian error (measurement error), which we will represent with random variable e_n . The observations are thus modeled as the true values with additive Gaussian error,

$$y_n = \zeta(x_n) + e_n \quad n = 1, \dots, N \quad (2.1)$$

This Gaussian error contaminating the observations contributes to parameter uncertainty.

MODEL FORM

Model form uncertainty is seen in the discrepancy between the simulation output and the observations. This is due to incomplete knowledge about the system in consideration, missing physics in the model, and necessary approximations, simplifications, and numerical methods needed to solve a model efficiently and accurately [49, 50]. If the model of the physical system is represented by $m(\cdot)$, which receives either a scalar or vector of input parameters θ and is evaluated at N variable inputs, $x = \{x_1, \dots, x_N\}$, then the model evaluation of the system is denoted as $m(x, \theta)$. The model of the system is related to the true underlying process $\zeta(x)$ by accounting for the discrepancy between the two at each variable input, which we denote by $\Delta(x)$,

$$\zeta(x_n) = m(x_n, \theta) + \Delta(x_n), \quad n = 1, \dots, N. \quad (2.2)$$

In other words, with the inclusion of this discrepancy term, the model output is exactly representative of the true property values. An important note is that θ is assumed to be the true but unknown parameter scalar or vector. This is significant since, even under the assumption of zero parameter uncertainty, there is still bias between the model and true system due to the imperfect representation by the model of the physical system. Altering Eq. (2.1) to include this additional information, the model for the observations now becomes,

$$y_n = m(x_n, \theta) + \Delta(x_n) + e_n, \quad n = 1, \dots, N. \quad (2.3)$$

Independence is assumed for $m(\cdot)$ and $\Delta(\cdot)$.

Kennedy and O'Hagan [35] were among the first to consider discrepancy in inference to avoid biased and over-confident parameter estimates. Not accounting for this type of uncertainty is especially problematic when the parameters have intrinsic value representing some material property which allows the model to be extrapolated to other scenarios. Without the consideration of model discrepancy, these true parameter values are unlikely to be recovered in the inference, yielding unreliable propagation of uncertain parameter inputs to model predictions.

A Gaussian process is one approach by which the discrepancy between the model and true mean property value can be accounted for [34]. This approach treats the discrepancy as a function with values equal to the amount of bias. Further modeling considerations are explored and discussed by Brynjarsdóttir and O'Hagan [51]. Additional details on Gaussian processes for accounting for model misfit will be discussed in later.

Both sources of error stemming from the model, form and parameter, are a subset of epistemic sources of uncertainty which can be reduced. Model form uncertainty can be reduced by conducting additional experiments to learn more about the system of interest and incorporating additional physics into the model. In addition to this, limiting the amount of simplifying assumptions in the model and using more accurate numerical methods to arrive at the correct solution will aid in reducing the contribution from this source. Parameter uncertainty can also be reduced by improving both the quality and quantity of calibration data.

UNCERTAINTY QUANTIFICATION & PROPAGATION

The focus of the work presented in this report hinges on applying statistical tools for uncertainty quantification and propagation in order to establish a known model and simulation fidelity, which can then be used to quantify design confidence and support risk analysis and decision making.

Uncertainty quantification refers to a collection of techniques used to provide probability distributions characterizing the uncertainty in the unknown model components as well as any correlations which may exist among them. Uncertainty propagation is the notion of repeated model evaluations at many possible input parameter values in order to develop an understanding of how the uncertainty in inputs translates to the output space of the model. Many techniques are available to accomplish these goals. A few such examples include analytical non-sampling methods such as Polynomial Chaos [52] and Kriging [53], most probable point-based methods (first-order and second-order reliability method) [54, 55], numerical sampling-based techniques such as Monte Carlo [56, 33] and importance sampling [57], as well as non-probabilistic approaches such as interval analysis [58] and fuzzy theory [59, 55]. The selection of the method is application-dependent and at the discretion of the analyzer.

In the following sections, a discussion on how statistical inference can be used to achieve the goals of UQ and UP is presented as well as a description of the methods adopted in this work for analyses in materials modeling.

STATISTICAL INFERENCE

Statistical inference is the process of establishing information about unknown components of a probability model conditioned on observed data while accounting for the associated sources of uncertainty. In this scenario, data y is assigned a probability model which consists of the set of distributions $M = \{ f(y | \theta); \theta \in \Theta \}$, where θ represents the unknown components of the model which belong to the space Θ . The θ s may be either unknown input parameters of the model or other unknown components of the system.

There are two primary paradigms to consider for statistical inference which handle the estimation and interpretation of uncertainty in very different ways: these are the classical, or frequentist, setting and the Bayesian setting. The primary deviation between the two paradigms is the use of prior knowledge to establish an understanding of parameter uncertainty [60].

The frequentist setting treats the parameters of interest as fixed but unknown quantities. A ‘best guess’ for the unknown parameters of interest is established through the selection of a function, called an estimator, which describes the observed data (i.e., maximum likelihood (MLE) or a least squares estimate) [60]. The data in this scenario are assumed to be random samples from a random variable, or single realizations from a population, which are statistically independent of one another. With infinite random samples from the population, the estimator approximates the true value of the unknown parameter and is considered unbiased if its expected value equals the population parameter. Hence, inference follows by drawing conclusions on the unknown model components based on the frequency of the data being observed. The variability,

or sampling distribution, of the chosen estimator is established under all possible data scenarios. In other words, inference on the parameters is performed by determining the likelihood of the observed data given those parameters [61]. The results of the inference is then presented through p-values and confidence intervals [60].

In contrast to the frequentist approach, the Bayesian setting formally incorporates prior information, knowledge the analyzer has before any data is observed, into the design of the inferential statistical model. The parameters of interest in this setting are treated as unknown random variables with an associated probability distribution describing the uncertainty about their values. Inference follows by directly updating the prior understanding of the unknown model components by conditioning on the available data to establish a subjective belief about the uncertainty of these unknown quantities.

Figure 2.3, provided by Skrepnek [60], illustrates the fundamental differences between the frequentist and Bayesian approach. The top figure is representative of the frequentist approach with θ being the true but unknown parameter value and $\hat{\theta}$ being the estimated value, which is informed based on the likelihood of the data being observed given the parameter values. The bottom panel is an illustration of the Bayesian paradigm, whereby the uncertainty in the unknown parameter, represented through the posterior distribution, is achieved through a combination of the likelihood *and* prior.




Figure 2.3: Frequentist vs Bayesian inference. θ is the true value and $\hat{\theta}$ is the 'best guess' provided by the estimator [60].

THE BAYESIAN PARADIGM

The Bayesian paradigm provides the framework used in this work to understand and propagate uncertainty in material models.

The initial modeling step required in the Bayesian approach is to incorporate any prior knowledge about the unknown model components, before any data are observed, into the statistical model. This a priori information is distilled into a probability model called a prior distribution and is denoted by $\pi(\theta)$. The confidence the analyzer has in the prior knowledge is reflected in the strength of the prior distribution. For example, a prior distribution which is more diffuse is less informative than one which is more concentrated. This is an important modeling step since it allows the inclusion of any hard or soft constraints to be placed on the parameters required by physics or for the analyzer to incorporate their knowledge of the expected value or ranges of the unknown parameters. Once any additional information becomes available, data for example, the prior distribution is updated by conditioning on this information in order to establish the posterior distribution via Bayes' rule:

$$\pi(\theta|y) = \frac{f(y|\theta)\pi(\theta)}{\int_{\Theta} f(y|\theta)\pi(\theta)} \quad (2.4)$$

This updated posterior distribution, equal up to a constant of proportionality to the product of the likelihood and the prior, is a complete description of the belief about θ . Estimates of the parameter value follow from the posterior mean, and a summary of the uncertainty about its value can be established through the posterior quantiles and highest posterior intervals, which describe the probability of the parameter value being within a certain range.

The Bayesian approach is appealing since it can seamlessly extend to make predictions about unobserved data, y^* . For example, applied for UQ in a crystal plasticity model, as will be demonstrated in later chapters, once the posterior over the model parameters of interest has been recovered, it can be used to predict the stress-strain response in experiments not yet conducted. This involves jointly modeling the posteriors of the unobserved data and of the unknown parameters given the data which was observed and marginalizing over θ ,

$$\pi(\tilde{y}|y) = \int_{\Theta} f(\tilde{y}|\theta, y) = \int_{\Theta} f(\tilde{y}|\theta) f(\theta|y) d\theta \quad (2.5)$$

Therefore, inference and prediction follow directly from the posterior and posterior predictive distributions, respectively.

MARKOV CHAIN MONTE CARLO TECHNIQUES

A challenge to consider is that the posterior densities are often not analytically tractable due to their high dimensionality as a result of their large parameter spaces. The state-of-the-art methods to recover uncertainty information contained in posterior summaries is through MCMC simulations [62–65]. MCMC methods are comprised of a class of techniques which can be used to sample from the target probability distribution (the true but unknown distribution of the unknown parameters) defined on some high-dimensional space. As the number of samples in the

chain tends towards infinity, the stationary distribution of the Markov Chain matches the target posterior distribution of interest [33].




Figure 2.4: Schematic of a Metropolis-Hastings Markov chain Monte Carlo simulation from [66]

The most simple sampling technique is Gibbs sampling [67], which capitalizes on choosing a likelihood and prior probability distribution such that the posterior distribution of interest is analytically tractable, also known as conjugacy [68]. The Markov Chain can then be updated with a series of full-conditional updates, thus sampling directly from the target posterior. However, this is not always possible, and a variety of other algorithms can be used which are based on the Metropolis-Hastings algorithm. Metropolis-Hastings Markov chain Monte Carlo (MH-MCMC) [69, 70] is a technique used to create a Markov chain by drawing samples, $\theta(1), \dots, \theta(M)$, in a manner such that the stationary distribution of the chain matches the posterior distribution, $\pi(\theta | y)$, of interest. At each iteration, $m \in \{1, \dots, M\}$, parameter $\theta(p)$ is proposed from a symmetric proposal distribution centered at the current chain value with density $p(\theta | \theta(m-1))$. The proposed parameter is then accepted or rejected with probability $\min(1, \pi(\theta(p) | y) / \pi(\theta(m-1) | y))$, repeating the previous parameter value, $\theta(m-1)$, otherwise. In this way, proposed parameters, $\theta(p)$, which yield a higher posterior than the previous value, $\theta(m-1)$, are always accepted into the chain. Figure 2.4 from [66] illustrates this sampling technique. The prior distribution on the left is updated to the posterior distribution on the right by sampling parameters and accepting or rejecting contingent on their relative posterior probability.

RANDOM EFFECTS HIERARCHICAL MODEL

An additional benefit of the Bayesian paradigm is that it allows for the selection from a variety of statistical models. In this work, it is proposed the random effects hierarchical model is well-suited for uncertainty quantification and propagation for a class of problems within materials modeling. In order to make this case, a comparison between the random effects model and an alternative model, the fixed effect model, is very useful. The primary difference between the two resides in the assumptions made about the distribution of the experimental data [71].

In the fixed effect model, it is assumed that there is a single underlying true property value, which we will call θ , that underlies all studies or experiments. Therefore, any deviation seen in the data from experiment to experiment is supposed to be strictly due to sampling error. In contrast, the random effects model assumes there is an additional source of variability which causes deviations in the experimental data experiment-to-experiment. Here, it is supposed the mean property value underlying each experiment, $\theta[s]$, $s = 1, \dots, S$, where S is the number of experiments performed, do not reflect the true underlying state nature, θ . Instead, they are considered as random samples from the distribution of some underlying population which describes θ (the overall effect); these are called random effects [61]. The random effects model is more flexible since it does not impose strict assumptions on the distribution of the experimental data. Parameter inference performed under the random effects model result in a higher estimated parameter variance since an additional source of uncertainty is included.

Figure 2.5 from Borenstein [71] illustrates the difference between the random effects and fixed effect model with the classical example of surveying students. If two groups of students from the same university are surveyed in two separate studies, it is reasonable to assume these two groups share the same true mean. In Fig. 2.5a, the statistics from the two studies are shown with the estimated means of each study denoted by squares. The true underlying mean of each study, marked with a circle, is assumed to be equal to the true mean and is represented by θ . If we now consider the case of surveying students across multiple campuses, it is possible and indeed likely, that the true sample mean varies from campus to campus. In this scenario, the fixed effect model would not apply and we should opt instead to use the random effects model. In Fig. 2.5b, the true sample means of each study (circles) are distributed about the true mean of the population which is denoted by μ and are considered as a random samples from the population.




Figure 2 5: Schematics comparing the (a) fixed effects (left) and (b) random effects (right) models [71]

With these two models in mind, consider now the application to uncertainty quantification in materials modeling. In many experimental setups (mechanical testing, for example) the property being evaluated is influenced by the microstructure of the specific sample being tested. As a result, samples cut from the same material will each be tested under slightly different states due to the random microstructure, producing a variability in experimental data which cannot be attributed to sampling error. Each sample has a unique true underlying property value due to its unique microstructure; however, there is a single true overall property of the material, θ , about which the sample properties are distributed that is wished to be known. In addition to this, in many cases, obtaining calibration data of sufficient quantity and quality for model calibration can be quite costly. A reasonable solution to this dilemma is to utilize data from other labs or in published work to increase the information in the analysis. In this scenario, the data is also likely to have been generated under varying states or conditions, thus demonstrating an inconsistency across sources. For these reasons and in these settings, the random effects model would likely yield more accurate estimates on the model parameters. In the event there is no variability in the true sample means, a random effects analysis will yield the same results as a fixed effect analysis, since the latter is just a special case of the former.

As a final note on the choice of model, a hierarchical statistical model is one which is written on multiple levels (or in a hierarchical nature) [72]. Under this structure, inference made about one quantity affects the inference on another since they share certain attributes and similarities. It allows probability models which contain a large number of random variables to be written as products of conditional distributions, implying a dependence on the joint probability model of the parameters [73, 74]. That this is the natural structure of the random effects model should not be surprising since the data are related to the overall effect only through the random effects. This leads to the following general notation for the posterior distribution obtained via Bayes' rule (2.4),

$$\pi\left(\theta, \{\theta^{[s]}\}_{s=1}^S | y\right) \propto \prod_{s=1}^S f(y | \theta^{[s]}) f(\theta^{[s]} | \theta) \pi(\theta) \quad (2.6)$$

The use of the Bayesian paradigm for inference has been supported as well as the selection of the statistical model used for uncertainty quantification and propagation for a class of problems in materials engineering. Specific modeling assumptions, generalizations and algorithms will be explicated in the respective sections in which they appear.

STATE OF THE FIELD OF UQ AND UP IN MATERIALS MODELING

The importance of UQ and UP in computational supported design is increasingly being acknowledged within the computational materials community and is reflected in the growing body of work in the field. In 2013, a review by Chernatynskiy et al. [2] compiled cutting edge work for UQ in materials modeling, focusing on multi-scale crystal plasticity and atomistic simulations. Also, very recently, Honarmandi et al. [75] reviewed state-of-the-art work in the field considering UQ and UP in computationally-supported materials design. In this work, the framework for statistical inference and UQ/UP is applied to two unique problems of interest: 1) mesoscale material mechanics and 2) thermodynamic calculations of phase equilibrium. While there has been significant exploration of UQ in other areas such as simulations at the atomistic scale [76–80], the review of current work exploring UQ and UP in materials modeling presented here is restricted to the general areas relevant to this research in order to prepare an understanding of the contribution to the field.

MECHANICS OF MATERIALS

As was illustrated in Table 2.1, there are many scales at which phenomena and mechanisms operate which govern the mechanical behavior of materials. Several examples of UQ being applied to models accounting for a range mechanisms affecting the final deformation of a material are detailed below.

Evidence theory (Dempster-Shafer theory [81, 82]) was used by Tang et al. [83] to quantify epistemic uncertainty in metal fatigue crack growth and also by Salehghaffari et al. [84] for UQ in the physics-based BCJ plasticity model [17]. The Bayesian inferential approach was used with a MH-MCMC algorithm by Yeratapally et al. [30] to consider parameter uncertainty and predict fatigue crack initiation in a microstructure-sensitive model and also by Honarmandi and Arroyave [85] for UQ and UP in the plastic response of multi-phase transformation-induced plasticity steels. In 2009, Kouchmeshky [29] quantified the uncertainty in macroscopic properties obtained from metal forming processes due to the randomness of microstructure and processes using a sparse grid collocation approach. Other examples of Bayesian Inferential approaches for UQ in Macroscopic properties include [86–89].

Examples of UQ applied to multiscale deformation processes include Kouchmeshky and Zabaras [31], who considered the multiscale deformation in polycrystalline materials using a sparse grid collocation strategy, and also by Koslowski and Strachan [32], who considered multiscale plasticity in FCC metals using response functions to propagate uncertainty across scales.

Other advanced techniques, such as multi-fidelity optimization [90–93], have been used by Sun [94] for the sheet metal forming process, and by Honarmandi [95] to predict precipitate

behavior in shape memory alloys using a technique presented by Kennedy and O'hagan [93]. Talapatra et al. [96] used Bayesian model averaging to efficiently explore the materials design space to identify configurations which provide optimal material properties to support materials discovery.

UQ SOFTWARE AND PROGRAMS

Finally, there are many available software and programs for UQ, and while they aren't necessarily specific to thermodynamics or material mechanics, or materials modeling for that matter, they could be applied in these areas. Dakota is one such toolkit developed at Sandia National Lab [97] which contains algorithms for least-squares uncertainty quantification with sampling, reliability analysis, parameter estimation, and sensitivity analysis, among others. This is widely used for applications spanning defense programs, climate modeling, as well as computational materials. Another example is the NESSUS probabilistic analysis software, developed at Southwest Research Institute [98], which has a suite of algorithms available for reliability and probability analysis. This software is used for probabilistic analysis of structural and mechanical components of systems in a wide range of applications, from aerospace and automotive structures to gas turbine engines and pipelines. In both examples presented here, the parameter distributions are defined prior to the analysis by the user; a variety of sampling techniques are then used to map the parameter variability to the output space.

A program developed by Hunt et al. [99], which has been largely demonstrated in materials applications, is the PRISM uncertainty quantification framework (PUQ), providing techniques to sample and propagate the uncertainty of model input parameters through the computer simulation. Sampling methods such as Monte Carlo, Latin Hypercube and sparse grids are available, and computationally inexpensive surrogate models, generated through polynomial chaos and radial basis sets, are used which relate the uncertainty of the inputs with the quantities of interest. The uncertainty of the input parameters is assumed to be known a priori, whether normal, exponential, Weibull or a custom probability density function (PDF) generated in a previous study. This tool is operated under the assumption the model incorporates the physics necessary to produce an adequate prediction of the quantity of interest [100]. This program has been used in several examples of recent work including UQ for equiaxed solidification of an aluminum alloy [100–103], and UP for high pressure die casting for an aluminum alloy [104], and UP for aluminum alloy direct chill casting [105].

While this summary is by no means exhaustive, it demonstrates state-of-the-art techniques and procedures being used today for UQ and UP in materials modeling. The works reviewed have substantially advanced the state of UQ and UP in a variety of complicated problems, significantly contributing to rigorous model predictions for ICME. A wide range of techniques have been applied with modeling assumptions and considerations varying greatly case-to-case. In some cases, only parameter uncertainty is considered, with the assumption that the model contains the correct physics to capture the phenomena of interest. While the assumption of model-correctness is certainly appropriate in some cases, this is a decision that should be made with care since it can affect parameter estimates and property prediction to the detriment if this is a faulty belief. Some analyses are performed under the assumption the model parameters have a known probability model describing their uncertainty, which is then

propagated to uncertainty in the material property. Finally, while some studies do consider aleatory uncertainty, this is not always the case and many only consider epistemic uncertainty.

Uncertainty Quantification for Parameter Estimation and Response Prediction: Generalizing the Random Effects Bayesian Inferential Framework to Account for Material and Experimental Variability

Integrated Computational Materials Engineering (ICME) is a relatively new area of focus which has received increasing attention in materials science research over the past decade [50, 106–109, 3, 5, 110]. The optimization of materials, manufacturing, and component designs is accomplished through the integration of materials information in computational tools spanning a significant range of length and timescales. This chain of information leads to material performance predictions supported by the integration of material models providing processing, microstructural, and property information [50, 111, 112, 3, 113].

Along with the many benefits of ICME come new challenges. Traditionally, a sufficient confidence in material performance has been achieved via extensive and repetitive experimentation. For example, Federal Aviation Administration certification of metallic structures is often achieved based on “A”-basis design, where at least 99% of the material population equals or exceeds minimum property values with 95% confidence [41]. A key benefit of ICME is the ability to reduce the bulk of these expensive and lengthy experimental procedures, leading to a reduction in the cost of time and resources. This in turn requires a statistical design confidence to be established from models and simulations in-lieu of experimentation. Because of this, the continued growth of ICME is dependent on the adoption of statistical approaches for model uncertainty quantification (UQ). Specifically, we focus on the uncertainty in calibrating model parameters when confronted with limited and/or noisy calibration data.

UQ is the quantification of the effect various sources of error, present in both the model and experimental data, have on the model output. These sources are categorized as either epistemic or aleatory; the former is due to imperfect knowledge, and the latter is due to the inherent physical randomness of a system or from the use of stochastic numerical algorithms [114, 49]. Aleatory uncertainty is often seen in simulation results in the form of random error, while epistemic uncertainty may reveal itself in simulation results in the form of variability primarily stemming from model-form or parameter uncertainty [49]. The presence of discrepancy between experimental data and model output is from model-form uncertainty. This is due to one or more of a variety of sources which can include but are not limited to necessary approximations, simplifications, and selected numerical methods [49, 50]. Parameter uncertainty describes our uncertainty about the model parameters when data with inherent measurement error or with sample-to-sample variability are used to calibrate the model. When uncertainty is introduced at model initiation from parameters, it will assuredly lead to error in the simulated data generated by the model [49]. Therefore, quantifying this uncertainty at the model input and tracking its effect on the model output is critical to measure the fidelity of a model result.

Uncertainty quantification is particularly important when integrating models across length scales for ICME and tracking how uncertainty propagates between them. This is very

difficult in large because many complex simulations for material design are deterministic in nature and do not inherently provide a statistical measure on uncertainty [10, 3]. Mature tools for UQ, developed through mathematical probability theory and statistical analyses to establish model fidelity, have been successfully applied in many disciplines [2]. While there is a growing understanding of the importance of UQ within the materials modeling community, a cohesive UQ framework has not been fully integrated. Models are often applied for materials performance predictions while giving surprisingly low attention to identification and quantification of the sources of uncertainty in these models, leaving model fidelity questionable. Uncertainty originating from the experimental data used for model calibration propagates to the model output and has a great effect on the material performance prediction. Therefore, improperly addressing these sources, or not at all, can lead to unreliable performance predictions.

This study considers a generalized random effects framework for quantifying parameter uncertainty when calibrating to noisy or inconsistent experimental data. We then apply the framework to two case studies to quantify model fidelity. The initial study is for framework development and applies UQ to crystal plasticity modeling, which provides key microstructural and grain-level deformation information for use within the ICME chain. The second case study demonstrates the generality of the framework by performing UQ in thermodynamic calculations of phase equilibrium.

A widely used non-sampling approach for UQ in the materials computational community is through polynomial chaos expansion (PCE) [31, 115–118, 22, 99, 102]. This method evaluates model uncertainty through the construction a truncated PCE of a stochastic surrogate model [119, 2] and is very useful in cases where the model in consideration can be reasonably approximated with linearizations [120]. An example of software utilizing this technique used within the materials modeling community is the PRISM uncertainty quantification framework (PUQ) [99, 103, 102, 101], which propagates uncertainty by using an a priori known uncertainty in model input parameters (i.e., uniform, normal, exponential, Rayleigh, Weibull, or custom probability density function (PDF)) to determine the output distribution of the quantity of interest using a variety of techniques such as Monte Carlo or Latin Hypercube. The work we present here is fundamentally different than this approach in that parameter uncertainty is a priori unknown. While the resulting parameter PDFs could in theory be used in conjunction with UQ tools such as PUQ, natural outputs of the Bayesian approach are posterior and posterior predictive distributions of the quantity of interest. We use sampling-based Markov Chain Monte Carlo (MCMC) simulations to explore parameter and predictive uncertainty by sampling from these distributions. Homogenized, coarse-grained, or reduced-order models are well suited for this class of methods, and computationally intensive models can be handled by embedding a surrogate model into the Bayesian hierarchy [121], allowing the uncertainty associated with this emulation to be seamlessly propagated into the posterior and posterior predictive uncertainty.

Uncertainty Quantification Accounting for Model Discrepancy Within a Random Effects Bayesian Framework

A critical component to advancing the goals of integrated computational materials engineering (ICME) is the quantification, propagation, and mitigation of the sources of uncertainty influencing model simulations and their predictive power, formally referred to as

uncertainty quantification in the materials computational community. This is of particular interest and importance with efforts to develop model linkages where the input to large scale component models is generated from the output of models at smaller scales. The underlying goal is to account for the multitude of mechanisms and phenomena occurring at scales ranging from chemical bonds at the quantum and atomistic levels and grain interactions at the mesoscale contributing to the macroscopic properties and performance of a design component. While the benefits of such an approach are clear (improved model fidelity and design reliability), the effectiveness is contingent on the adoption and routine use of rigorous and reliable techniques for uncertainty quantification (UQ), propagation (UP), and mitigation (UM). Unaccounted error or uncertainty at the lower-level models will lead to macroscopic model outputs with compounded uncertainty and predictions with an unknown level of reliability.

Figure 2.6 summarizes the various types of uncertainty which are grouped into two primary categories: epistemic and aleatory. Both sources play a role in the model output and quality of parameter inference and, consequently, the propagation of uncertainty for predictions (or in the case of ICME, model linkages). Aleatory uncertainty is a description of the randomness in the system, and could be due to natural physical randomness, which may manifest in the sample-to-sample variability among experimental tests, or due to the use of stochastic numerical methods. This form of uncertainty, while quantifiable, is not reducible. On the other hand, epistemic uncertainty is a result of omitted or missing knowledge or physics in the constitutive models, or from the use a variety of numerical techniques to simplify the computations. Since a model can always be improved by learning more about the system and incorporating the understanding into the model, this type of uncertainty is reducible.




Figure 2 6: Various sources of uncertainty present in computational models that must be accounted for to perform reliable inference and prediction.

Mature techniques for UQ, UP, and UM are being adopted with increased frequency within the materials modeling community [22–32, 100, 103, 94, 96]. A recent paper by Ricciardi, Chkrebtii, & Niezgoda [175] proposed using a random effects Bayesian hierarchical statistical

model for quantifying parameter uncertainty in material models. The authors argue that this statistical model reflects the inherent structure of an important class of problems in materials modeling and simulation since the aleatory uncertainty caused by microstructure heterogeneity in the material samples tested is naturally accounted for. A Bayesian inferential approach was used to recover information on the unknown model parameters. Posterior functionals were estimated from a posterior sample generated via a Metropolis-Hastings Markov Chain Monte Carlo (MH-MCMC) simulation. Posterior uncertainties were then propagated to induce distributions in the model output and predict the distribution of possible future outcomes (prediction). This approach was demonstrated for two fundamentally different modeling problems: 1) crystal plasticity modeling at the mesoscale and 2) thermodynamic calculations of phase equilibrium.

The work presented here extends the random effects framework for estimation and uncertainty quantification in the presence of model misfit, also known as model discrepancy. Model misfit is a consequence of missing physics, modeling simplifications, or numerical methods that may lead to systematic discrepancy between the model output and observations. Improperly accounting for model misfit, or simply ignoring misfit when performing parameter inference or uncertainty quantification can lead to bias in calibration that propagates as error to property or response predictions.

One crucial application is when model parameters have a clear connection to intrinsic material properties (such as diffusivities or mobilities, defect energies, elastic constants, etc.). In an ideal world, the model is a perfect representation of reality and, when properly calibrated, will exactly predict the material response under varying initial and boundary conditions. However, in practice, models are often based on some assumptions or approximations. As a thought experiment, the measured experimental material response can be decomposed into three components: 1) the part well represented by the model, 2) the portion of “true” material response not captured by the model due to model discrepancy, and 3) noise or measurement error. Mathematically the discrepancy should be independent of the portion represented by the model. When the model has enough degrees of freedom, i.e., fitting parameters, it is often possible to find a set of parameters that allow the model to closely mimic the experimental response. When fitting is performed without accounting for the discrepancy, the values of the parameters are biased by this misfit (meaning they attain a value that is systematically different than that obtained from an ideal model). Using such biased parameter estimates for response prediction under different conditions than are used for calibration will necessarily lead to a persistent error in predictions and inaccurate error bounds, leading to overly optimistic design tolerances or safety factors.

Quantifying this source of uncertainty is rendered challenging by the requirement of balancing a flexible discrepancy model with the requirement that the complete model be identifiable given the data [51]. Discussion on model discrepancy and techniques to model this form of error are provided in following sections.

3: METHODS, ASSUMPTIONS, AND PROCEDURES

Uncertainty Quantification for Parameter Estimation and Response Prediction: Generalizing the Random Effects Bayesian Inferential Framework to Account for Material and Experimental Variability.

Statistical inference refers to the goal of recovering information about unknown components of a probability model from observed data while accounting for associated uncertainty. A probability model for the data y is a collection $M = \{ f(y | \theta); \theta \in \Theta \}$ of probability distributions indexed by unknown model components, or parameters, θ that belong to a parameter space Θ . Two major paradigms for inference are the classical and the Bayesian views, which have fundamentally different approaches to the characterization and interpretation of uncertainty. The classical, or frequentist, approach consists of identifying a function of the data called an estimator which provides a “best guess” of the unknown parameters. Inference is based on establishing the distribution of the estimator (sampling distribution) under all possible data scenarios. This understanding of uncertainty answers the question: What is the variability of the chosen estimator under the assumed data-generating probability model? This information may then be used to construct hypothesis tests and build confidence regions. In contrast, the Bayesian paradigm is a framework for directly defining and updating beliefs about unknown quantities conditionally on observed data that are not centered around the choice of an estimator (though estimators can be obtained under this framework in a principled way). The Bayesian paradigm, which will be adopted in this paper, instead answers the question: What is my subjective belief about the unknown model component given the data that were observed?

The Bayesian approach requires an initial modeling step for uncertainty in the unknown parameters before any data are observed. This a priori, or prior knowledge is distilled into a probability model (i.e., a distribution) that describes the relative frequency of the unknown quantity[1]. This model is called the prior distribution and its density is denoted by $\pi(\theta)$. A prior density that is very diffuse indicates more prior uncertainty than a density that is more concentrated. Prior specification can also be helpful in constraining the parameter space in cases where a parameter cannot take on values in a given region, by choosing a prior density with a restricted domain. Therefore, an important benefit of the Bayesian approach is its ability to enforce both hard and soft constraints, including the range or most probable values, required by physical laws or obtained through understanding about the system based on previous studies. A prior is not a model of the fixed but unknown parameters, but rather of our uncertainty about the parameter before the data are observed. When data (or generally any source of information) become available, the prior model is updated by conditioning on the available information. Denoting the density of each candidate distribution in M by $f(y | \theta)$, the updated, or posterior, density is obtained via Bayes’ rule as,

$$\pi(\theta|y) = \frac{f(y|\theta)\pi(\theta)}{\int_{\Theta} f(y|\theta)\pi(\theta)d\theta} \propto f(y|\theta)\pi(\theta) \quad (3.1)$$

The posterior distribution with this density completely characterizes our updated belief about θ . Functionals of the posterior can be used for estimation (e.g., the posterior mean), and as summaries of uncertainty (e.g., posterior quantiles or probabilities of the parameter lying in a

specific region). Due to this direct characterization of uncertainty, the interpretation of credible intervals (quantiles of the posterior) and statements about beliefs are intrinsically different than those provided by confidence intervals and hypothesis tests in the classical setting.

One of the appeals of the Bayesian paradigm is that it can be seamlessly applied to any inference problem. For example, the extension to make prediction about unobserved data, y^* , involves joint posterior modeling of the unobserved data and the unknown parameters given the observed data y , which is then marginalized over θ as,

$$\pi(y^*|y) = \int_{\Theta} f(y^*, \theta|y) d\theta = \int_{\Theta} f(y^*|\theta) \pi(\theta|y) d\theta \quad (3.2)$$

The density of this posterior predictive distribution is a continuous mixture of the likelihood of the unobserved data weighted by the posterior density of θ given the observed data y . Posterior predictive intervals and probabilities may be obtained directly from this density.

Since inference and prediction follow entirely from the posterior and posterior predictive distributions, respectively, an important consideration is how to recover these when the posterior densities are not analytically tractable. In other words, although we may be able to write the posterior density up to a constant of proportionality as in Eq. (3.1), we may not be able to exactly compute integrals $\int_{\Theta} q(\theta) \pi(\theta|y) d\theta$ required to obtain, e.g., posterior means ($q(\theta) = \theta$, $(a, b) = \Theta$), posterior probabilities ($q(\theta) = 1$) within (a, b) , and other summaries required to communicate posterior uncertainty. Though numerical integration may be used, the relatively large dimensionality of the parameter space Θ leads to large numerical errors. Markov chain Monte Carlo (MCMC) simulation methods are the state of the art for this task [122, 33, 123–125]. They consist of constructing a Markov chain in such a way that its stationary distribution equals the posterior distribution of interest. This yields an approximate posterior sample, $\theta^{(1)}, \dots, \theta^{(M)}$ from which we can compute a Monte Carlo estimator:

$$\int_a^b q(\theta) \pi(\theta|y) d\theta \approx \frac{1}{M} \sum_{m=1}^M q(\theta^{(m)}) \mathbb{I}\{\theta^{(m)} \in (a, b)\} \quad (3.3)$$

which has the desirable statistical property of consistency (the mean squared error tends to zero as the number of Monte Carlo samples grows).

RANDOM EFFECTS HIERARCHICAL STATISTICAL MODEL

The flexibility of the Bayesian paradigm allows consideration of a large variety of models. We argue that a random effects hierarchical Bayesian model reflects the natural structure of an important class of problems within the field of materials modeling and simulation. A schematic representation of the proposed random effects model is shown in Fig. 3.1. Here, we review this model structure, by considering common distributional assumptions that can be easily generalized to fit different problems.

In many applications, the data collected do not directly reflect the universal underlying state of nature θ . This occurs when, for example, the parameter of interest θ is an overall property or effective property of a system, such as a material, but the data themselves are

collected from samples that, due to heterogeneity or variability, may have been generated under a different state $\theta[s]$ (e.g., scatter in material property measurements to due variability in structure or noise in the measurement). Though $\theta[s]$ itself is not of interest (a nuisance parameter), it is connected to θ by the probability model,

$$\theta[s] | \theta, \Lambda \sim f(\theta[s] | \theta, \Lambda), \quad s = 1, \dots, S, \quad (3.4)$$

where Λ denotes other quantities parameterizing this model, which may be either set by the user or unknown (see below). The sample parameters $\theta[1], \dots, \theta[S] \in \Theta \subset \mathbb{R}^D$ are called random effects and the state $\theta \in \Theta \subset \mathbb{R}^D$ is called the overall effect. Note that in this context the expression “effects” is used to indicate correlation rather than causality. In the absence of known experimental bias, a multivariate normal model is often appropriate in the context of measured material properties, implying that individual measurements are symmetrically scattered about the overall parameter and small deviations are more likely than large,

$$\theta^{[s]} | \theta, \Lambda \sim \text{MVN}_D(\theta, \Lambda^{-1}), s = 1, \dots, S \quad (3.5)$$

where the overall parameter θ is the mean and Λ is a $D \times D$ symmetric, positive-definite precision matrix that controls the amount of dispersion of the random effects about θ . Additional constraints on the overall parameters may be included in this model as illustrated in “Case Study I: Reduced-Order Crystal Plasticity” and “Case Study II: Thermodynamic Calculations” sections.

For each material sample $s \in 1, \dots, S$, the likelihood of the experimental result $y \in Y \subset \mathbb{R}^N$ is also a function of $\theta[s]$. In analogy to Eq. (3.5), it is often convenient to adopt a multivariate normal model as:

$$y^{[s]} | \theta^{[s]}, \Psi \sim \text{MVN}_N(m(\theta^{[s]}), \Psi^{-1}), \quad s = 1, \dots, S \quad (3.6)$$

where $m(\theta[s])$ is the forward model evaluated at parameter $\theta[s]$, and Ψ is a $N \times N$ symmetric, positive-definite precision matrix associated with observation error. That is, Ψ controls the amount of dispersion of the data about the forward model output $m(\theta[s])$. This corresponds to the model $y[s] = m(\theta[s]) + e$. The random variable e represents noise or error in the measurement and is modeled as an N -dimensional multivariate normal random variable with mean zero vector and precision matrix Ψ .

After the choice of a suitable prior probability model $\pi(\theta, \Lambda, \Psi)$ is selected, the posterior density of all the unknowns can be expressed up to a multiplicative constant as:

$$\begin{aligned} & \pi(\theta, \theta^{[1]}, \dots, \theta^{[S]}, \Lambda, \Psi | y^{[1]}, \dots, y^{[S]}) \\ & \propto \pi(\theta, \Lambda, \Psi) \prod_{s=1}^S f(y^{[s]} | \theta^{[s]}, \Psi) f(\theta^{[s]} | \theta, \Lambda), \end{aligned} \quad (3.7)$$

under the assumption of conditional independence of the data from material sample to sample, given the model parameters. Inference on the unknowns can proceed directly from this posterior density.

Similarly, the problem of prediction is based on the marginal posterior predictive distribution with density,

$$\begin{aligned}
& \pi \left(y^{[new]} \mid y^{[1]}, \dots, y^{[S]} \right) \\
&= \int \pi \left(y^{[new]}, \theta^{[new]} \mid y^{[1]}, \dots, y^{[S]} \right) d\theta^{[new]} \\
&= \int f \left(y^{[new]} \mid \theta^{[new]}, \Psi \right) f \left(\theta^{[new]} \mid \theta, \Lambda \right) \\
& \quad \pi \left(\theta, \Lambda, \Psi \mid y^{[1]}, \dots, y^{[S]} \right) d\theta d\theta^{[1]} \\
& \quad \dots d\theta^{[S]} d\Lambda d\Psi d\theta^{[new]}.
\end{aligned} \tag{3.8}$$

The marginalization is required because a new observation $y^{[new]}$ to be predicted is assumed to come from a new sample with random effect $\theta^{[new]}$. Therefore, a joint posterior predictive distribution for $y^{[new]}$ and $\theta^{[new]}$ is required, but since $\theta^{[new]}$ is a nuisance parameter, the posterior predictive distribution is integrated over all possible values of $\theta^{[new]}$.

PRIOR MODELING

Prior modeling is an important part of a Bayesian analysis. The degree to which the prior choice affects posterior uncertainty depends on (1) how informative the prior is and (2) how much information the data contain about the unknown parameters. Therefore, care must be taken in the selection of appropriate priors, especially in relatively data-poor situations. In many materials applications, the number of experimental measurements for a given sample is often reasonably large, so much information is available about the individual random effects $\theta[s]$, $s = 1, \dots, S$. However, the number, S , of samples tested is frequently quite small due to experimental cost constraints. For example, for a stress-strain test the sampling rate of load and displacement is quite high but the number of independent tests (random effects) is typically quite small. Since these random effects inform the overall effect θ , the amount of information propagated through to θ in the hierarchy may be limited, so that some sensitivity to the choice of $\pi(\theta)$ is expected. For this reason, it is advisable to choose a prior that is diffuse and whose domain is restricted to obey any physical laws or model constraints (as demonstrated below in our first case study).

In the absence of information to the contrary, a reasonable prior assumption is that certain parameters or groups of parameters are a priori independent. Although we often know of physical restrictions on the parameters defining the mechanistic model defining the system, it is less common to know such things about the prior precision matrices Λ and Ψ above. Furthermore, their relationship to the model parameters is, in general, not known before any data are observed, motivating a model of the form,

$$\pi(\theta, \Lambda, \Psi) = \pi(\theta)\pi(\Psi)\pi(\Lambda) \tag{3.9}$$

where $\pi(\theta, \Lambda, \Psi)$ is the joint probability density function of the unknowns and $\pi(\theta)$, $\pi(\Psi)$, $\pi(\Lambda)$ are their marginal density functions. Overall parameters θ are often also assumed a priori independent up to any hard constraints required by the physics, which will be denoted by $C(\theta) > 0$, yielding the expression,

$$\pi(\theta) \propto \mathbb{I}\{\mathcal{C}(\theta) > 0\} \prod_{d=1}^D \tilde{\pi}(\theta_d). \quad (3.10)$$

Specific forms of the prior densities $\tilde{\pi}(\theta_d)$ and the function $C(\theta)$ will be motivated by the application. Note that the assumptions of prior independence above do not preclude posterior dependence between the parameters after the data are observed.

Another consideration in the choice of priors is a practical matter, though its importance has decreased with the advent of sophisticated MCMC schemes. Conjugacy refers to the case in which the prior and likelihood are such that the resulting posterior is an analytically tractable distribution. Though this is not always possible (and indeed, not necessary), this idea can be used to simplify posterior simulation via Gibbs sampling [67]. The idea is to define the Markov chain as a sequence of full-conditional updates, one parameter or block of parameters at a time. These full conditional updates are derived in “Appendix A.2” for the following prior choices. The prior on the random effects inverse variance-covariance (precision) matrix Λ is assumed to follow a Wishart distribution, a probability distribution on positive-definite matrices, with density,

$$\begin{aligned} \pi(\Lambda) &= \text{Wishart}_D(\Lambda; V_o, \nu_o) \\ &\propto |V_o|^{-\nu_o/2} |\Lambda|^{(\nu_o-D-1)/2} \exp \left\{ -\frac{1}{2} \text{tr}(V_o^{-1} \Lambda) \right\}, \\ \nu_o &\geq D+1. \end{aligned} \quad (3.11)$$

Although this prior model is computationally convenient and performs well for the examples presented here, the Wishart prior is often restrictive due to the complex relationship between the hyperparameters and features of the distribution, which can introduce strong unintended prior information. Some more flexible alternatives include modeling the components of a decomposition of the covariance matrix instead [126, 127]. The same prior assumption may be made for the precision matrix Ψ of the observation error process e . But because observation error in many experimental settings is assumed to be independent from observation to observation, we make the further simplifying assumption that $\Psi := \delta^2 \mathbf{I}_N$, where $\mathbf{I}_N$ is the N -dimensional identity matrix. Because the Wishart is a matrix generalization of the Chi-squared distribution, the analogous prior on the precision δ^2 takes the form,

$$\begin{aligned} \pi(\delta^2) &= \text{Gamma}(\delta^2; a, b) \propto (\delta^2)^{a-1} \exp\{-b\delta^2\}, \\ a, b &> 0. \end{aligned} \quad (3.12)$$

Prior specification of the model parameters θ is typically application-specific, and will be discussed. This prior model may incorporate our knowledge about the absolute magnitude, relative sizes, and physical constraints for these quantities. Another important role for prior modeling is the ability to induce regularization when the inverse problem is ill-posed (the posterior is multimodal).




Figure 3 1: Visualization of a Bayesian random effects model. Arrows represent the direction of conditional dependence. Clear nodes represent unobserved/unknown quantities while shaded nodes indicate observations.

POSTERIOR SIMULATION VIA MARKOV CHAIN MONTE CARLO

This section describes the tools used to sample from the posterior (Eq. 3.7) and posterior predictive (Eq. 3.8) densities under the modeling assumptions described in “Random Effects Hierarchical Statistical Model” and “Prior modeling” sections. For this, it will be helpful to define the following short-hand notation for the squared £2 distance between two vectors b_1 and b_2 weighted by the matrix C ,

$$\|b_1 - b_2\|_2^2(C) := (b_1 - b_2)^\top C (b_1 - b_2) \quad (3.13)$$

which can be simplified to $\|b_1 - b_2\|_2^2$ when the weighting matrix is the identity.

The decision about which MCMC sampling method to use typically follows a sequence, beginning with, e.g., Gibbs steps, which are often efficient, but require closed-form updates that can be easily simulated, and escalating to a more general approach when such updates are not available. Under the assumptions outlined in this work, we are able to compute efficient Gibbs samplers for the precision, Λ , of the random effects distribution,

$$\begin{aligned} & f(\Lambda \mid y^{[1]}, \dots, y^{[S]}, \theta, \theta^{[1]}, \dots, \theta^{[S]}) \\ & \propto \text{Wishart}_D \left(\Lambda; \left\{ V_o^{-1} + \sum_{s=1}^S \left\| \theta^{[s]} - \theta \right\|_2^2 \right\}^{-1}, v_o + S \right) \\ & v_o \geq D + 1, \end{aligned} \quad (3.14)$$

as well as the observation error precision, δ^2 ,

$$\begin{aligned} & f(\delta^2 \mid y^{[1]}, \dots, y^{[S]}, \theta, \theta^{[1]}, \dots, \theta^{[S]}) \\ & \propto \text{Gamma} \left(\delta^2; \frac{SN}{2} + a, b + \frac{1}{2} \sum_{s=1}^S \left\| y^{[s]} - m(\theta^{[s]}) \right\|_2^2 \right), \\ & a, b > 0. \end{aligned} \quad (3.15)$$

The derivation of the above full-conditional densities is provided in the appendix.

In the absence of such full-conditional updates, a variety of algorithms based on the Metropolis-Hastings algorithm may be used. The Metropolis-Hastings Markov chain Monte Carlo (MH-MCMC, see for e.g., [125, 128]) aims to construct an ergodic discrete-time Markov chain $\theta(1), \dots, \theta(M)$ in such a way that its stationary distribution is the posterior density $\pi(\theta \mid y)$ of interest (Eq. 3.7). This can be accomplished at iteration $m \in \{1, \dots, M\}$ by proposing a parameter $\theta(p)$ from a symmetric proposal distribution with density $p(\theta \mid \theta(m-1))$ and then accepting the proposed value with probability,

$$\min\{1, \pi(\theta^{(p)} \mid y) / \pi(\theta^{(m-1)} \mid y)\} \quad (3.16)$$

and replicating the previous value $\theta^{(m-1)}$ otherwise. This procedure can be implemented on the entire vector of unknown model parameters, or in subsets (called blocks) conditional on previously sampled values of the remaining parameters. Blocking can increase the efficiency of the sampler, but is balanced by the fact that each evaluation of the acceptance probability requires a likelihood evaluation. In the end, the decision of how many blocks to use and which parameters to include is up to the analyst. In our analysis, overall model parameters were sampled in a single block, as were each of the sets of random effects parameters. This sampler may be extended to additionally generate samples from the posterior predictive density $\pi(y[\text{new}] | y)$. To see this, notice that Eq. (3.8) is a continuous mixture of likelihoods of the new data given the parameters weighted by their posterior probability density. Therefore, at every MCMC iteration described above, one may also generate a new datum $y[\text{new}]$ from the likelihood given $\theta^{(m)}$, thus resulting in a Markov chain $y[\text{new}](1), \dots, y[\text{new}](M)$ with stationary distribution given by the posterior predictive density. The implementation of this sampler under the modeling assumptions described in “Random Effects Hierarchical Statistical Model” and “Prior modeling” sections is detailed in Algorithm 1 in the appendix.

CASE STUDY I: REDUCED-ORDER CRYSTAL PLASTICITY

In general, crystal plasticity (CP) models are used to predict the mechanical behavior of polycrystalline materials, such as the stress-strain response and the texture evolution under specific deformation modes such as compression, shearing, or rolling [129–136, 15, 137, 11, 138, 139, 13, 140, 141, 12, 142–145, 14, 146–153, 20, 154, 16]. These models vary greatly in their complexity according to the constitutive laws (details of physical mechanisms and interactions included) and numerical formulations.

While the three-dimensional spatially resolved nature of models such as the full-field FFT model [145] or crystal plasticity finite-elements (CPFE) [14] makes them appealing for scientific inquiry, they are computationally too expensive to establish the statistical confidence required for design and optimization applications. This leads to the reliance on reduced-order, homogenized, coarse-grained, or mean-field models within ICME [155–157, 133] which are computationally much less expensive to optimize, making them preferential for UQ where thousands, or at times millions, of model evaluations are required to establish parameter and model uncertainty. The estimated model fidelity of reduced-order models is an important step in the development of ICME since these model predictions eventually become input for component scale or system-level models. Therefore, it is important to quantify and propagate such uncertainty between models by eliciting model predictive distributions.

As an initial case study, UQ was performed on the reduced-order viscoplastic self-consistent (VPSC) model developed at the Los Alamos National Laboratory by C.N. Tome and R.A. Lebensohn [1]. VPSC uses a mean-field solver and considers the plastic contribution to deformation due to externally applied stresses and strains under the visco-plastic mechanical regime[1]. VPSC represents the polycrystal with volume fractions of crystal orientations, or statistically representative (SR) grains. Each SR grain is treated as an ellipsoidal inclusion within an infinite homogeneous effective medium (HEM) representing all other SR grains in the system. The HEM represents the average environment seen by each grain with identical orientation [158]. The simplifications of this model, such as the linearization procedure used to describe the

grain/matrix interaction [1, 158], grain rotation, microstructure representation, and the empirical nature of the hardening law (more details found in [1]) provide sources of uncertainty within the model. These inherent model uncertainties in addition to the sample-to-sample uncertainty of training data, and the computational time to run VPSC make this an ideal model for exploration and development of the technique.

The local behavior within each SR grain is described with the following constitutive relationship,

$$\dot{\epsilon}^p = \sum_{\alpha} m^{\alpha} \dot{\gamma}^{\alpha} = \dot{\gamma}_0 \sum_{\alpha} m^{\alpha} \left(\frac{m^{\alpha} : \sigma^{\alpha}}{\tau_0^{\alpha}} \right)^n, \quad (3.17)$$

which is the sum of the local behavior over the set of slip systems α within the SR grain. The initial threshold stress is represented by τ_0 , m is the Schmid tensor, ϵ^p and σ are the deviatoric plastic strain rate and stress, respectively, and $\dot{\gamma}$ is the local shear rate on slip system α . $\dot{\gamma}_0$ is a reference shear rate and n is the rate-sensitivity exponent.

The evolution of the CRSS with deformation is described in VPSC by the Voce law, first proposed by Voce [159, 160] and modified by Kocks [161]. The Voce hardening law depends on four parameters τ_0^{α} , τ_1^{α} , ξ_0^{α} , and ξ_1^{α} . These are, respectively, the initial CRSS, the back-extrapolated increase in CRSS, the initial hardening rate, and the asymptotic hardening rate. This hardening law empirically captures the dynamics of hardening by considering the evolution of the CRSS, τ^{α} , as a function of the total accumulated shear over all the slip systems, $\Gamma = \sum_{\alpha} \Delta\gamma^{\alpha}$, in the grain,

$$\tau^{\alpha} = \tau_0^{\alpha} + (\tau_1^{\alpha} + \xi_1^{\alpha} \Gamma) \left(1 - \exp \left\{ -\Gamma \left| \frac{\xi_0^{\alpha}}{\tau_1^{\alpha}} \right| \right\} \right). \quad (3.18)$$

The increase in the CRSS, $\Delta\tau^{\alpha}$, is then determined by,

$$\Delta\tau^{\alpha} = \frac{d\tau^{\alpha}}{d\Gamma} \sum_{\alpha'} h^{\alpha\alpha'} \Delta\gamma^{\alpha'}, \quad (3.19)$$

where $d\Gamma$ is the shear increment, and α iterates through all slip systems within the grain. The inclusion of coefficients $h^{\alpha\alpha'}$ from the interaction matrix, h , allows for the consideration of 'self' and 'latent' hardening by empirically accounting for obstacles in slip systems α affecting the propagation of the activity on slip system α . For the case where the latent and self-hardening are indistinguishable, $h^{\alpha\alpha} = 1$ and $\Delta\tau^{\alpha}$ reduces to being just a function of the Voce hardening law,

$$\Delta\tau^{\alpha} = \frac{d\tau^{\alpha}}{d\Gamma} \Delta\Gamma. \quad (3.20)$$

For this work, we follow the convention [1] of imposing constraints on the Voce parameters to maintain a physically consistent hardening behavior for the room- temperature deformation of

copper with a positive hardening rate which decreases monotonically with additional strain: (1) $\dot{\xi}[s] - \xi[s] > 0$, (2) $\tau[s] \geq 0$ and (3) $\dot{\tau}[s] \geq 0$.

CALIBRATION DATA

For this case study, we adopted a virtual calibration data consisting of simulated experimental data. The virtual data were generated with a dislocation-based elasto-viscoplastic model within a full-field FFT (EVP-FFT) solver first developed by Lebensohn and later expanded by Ma et al. and Zhao et al. [16, 15, 162]. The use of virtual ground truth data was convenient for several reasons. First, it allows the evaluation of the MCMC algorithm and results since the virtual data have an a priori known amount of sample-to-sample variability as well as noise. Secondly, this allows control over data quality (i.e., noise level and sample-to-sample variability) and quantity, which can be used to consider different scenarios of data availability and noise level for a particular experimental setup.

The data were generated by operating the EVP-FFT model at optimized parameters reported by Zhao et al. [15] for polycrystalline copper under compression. Material response variability due to microstructure heterogeneity, as well as experimental variability (i.e., not performing a perfect test) was simulated by perturbing the model parameters. The degree of parameter perturbation was selected to agree with the results presented by Kiener (2008) [163] to mimic the natural sample-to-sample variability for miniature single-crystal copper samples in tension. The study of Kiener was chosen since the testing results from miniature samples have a much higher degree of variability in stress-strain response than standard tensile tests performed to ASTM E8 [164]; thus, it represents a worst-case scenario for sample-to-sample variability. We make the simplifying assumption the testing error is small compared to sample-to-sample variability and should be approximately zero. Therefore, noise was not added to the simulated data. However, the statistical model still estimates the error precision term (δ^2 in Eq. (3.15)). The estimation of this term is used as a diagnostic tool since we know a priori the true value of the noise, assuming that the reduced-order model is correct. Future work will incorporate a model error term that will account for model-form error.




Figure 3 2: Calibration data generated via dislocation-based EVP model for the stress-strain response up to 9% strain of miniature copper samples under tension. Six experimental data sets were created to simulate sample-to-sample and experimental variability

DETAILS OF THE STATISTICAL MODEL

For the CP VPSC case study, the random effects precision is a measure of the variability in the stress-strain response among the different material samples tested, and the error precision is a measure of the noise in the data (Fig. 3.2). The random effect for the s^{th} sample consists of the Voce hardening parameters associated with that sample. We denote these by the vector of random effects,

$$\theta^{[s]} := (\tau_0^{[s]}, \tau_1^{[s]}, \xi_0^{[s]}, \xi_1^{[s]})^T, s = 1, \dots, S. \quad (3.21)$$

We further assume that the Voce hardening parameters for each sample vary around a set of overall parameters,

$$\theta := (\tau_0, \tau_1, \xi_0, \xi_1)^T, \quad (3.22)$$

which we consider to represent the true average behavior of the system. This assumption is modeled after Eq. (3.5), with the vector $D = 4$ dimensional random effect vector $\theta[s]$ distributed as a multivariate normal random variable centered at the overall parameter vector θ , with precision matrix Λ . Here, we additionally incorporate any known constraints $C(\theta) = (\tau_0, \tau_1, \xi_1 - \xi_0) > 0$ on θ that are motivated by our understanding of the physical system as follows:

$$\begin{aligned}
f(\boldsymbol{\theta}^{[s]} | \boldsymbol{\theta}, \Lambda) &\propto MVN_D \left(\boldsymbol{\theta}^{[s]; \boldsymbol{\theta}, \Lambda^{-1}} \right) \mathbb{I}\{\mathcal{C}(\boldsymbol{\theta}) > 0\} \\
&\propto \det\{\Lambda\}^{1/2} \exp \left\{ -\frac{1}{2} \left\| \boldsymbol{\theta}^{[s]} - \boldsymbol{\theta} \right\|_2^2 (\Lambda) \right\} \cdot \mathbb{I}\{\mathcal{C}(\boldsymbol{\theta}) > 0\}, \\
s &= 1, \dots, S
\end{aligned} \tag{3.23}$$

resulting in a truncated multivariate normal model.

The observational error model for the s^{th} sample follows Eq. (3.6), which encodes the assumption:

$$y[s] = \sigma_{vp\text{sc}}(\boldsymbol{\theta} [s]) + \mathbf{e}, \quad \mathbf{e} \sim MVNN(0, \delta^{-2} \mathbb{I}_N), \tag{3.24}$$

where $\sigma_{vp\text{sc}}(\boldsymbol{\theta} [s])$ represents the output (Von Mises stress) from the VPSC model under Voce parameters $\boldsymbol{\theta} [s]$ for strain levels $\epsilon_1, \dots, \epsilon_N$. In other words, the observed Von Mises stress $y[s]$ corresponding to strains $\epsilon_1, \dots, \epsilon_N$ for the s^{th} sample forms an N -dimensional random vector which follows a multivariate normal distribution centered at the VPSC model evaluations $\sigma_{vp\text{sc}}(\boldsymbol{\theta} [s])$ with $N \times N$ variance-covariance matrix $\Psi^{-1} = \delta^{-2} \mathbb{I}_N$. Specifically,

$$\begin{aligned}
f(y^{[s]} | \boldsymbol{\theta}^{[s]}, \delta^2) &= MVN_N \left(y^{[s]; \sigma_{vp\text{sc}}(\boldsymbol{\theta}^{[s]}), \delta^{-2}} \right) \\
&\propto \delta^N \exp \left\{ -\frac{1}{2} \left\| y^{[s]} - \sigma_{vp\text{sc}}(\boldsymbol{\theta}^{[s]}) \right\|_2^2 (\delta^2 \mathbb{I}_N) \right\}, \\
s &= 1, \dots, S.
\end{aligned} \tag{3.25}$$

We next discuss the prior model in Eq. (3.10) in the context of this case study. The first two Voce hardening parameters, τ_0 and τ_1 are constrained to take values on the positive real line by a probability model that enforces this constraint with probability 1. Specifically, marginal Gamma prior densities,

$$\begin{aligned}
\tilde{\pi}(\tau_0) &= \text{Gamma}(\tau_0; a_{(\tau_0)}, b_{(\tau_0)}), \quad a_{(\tau_0)}, b_{(\tau_0)} > 0 \\
\tilde{\pi}(\tau_1) &= \text{Gamma}(\tau_1; a_{(\tau_1)}, b_{(\tau_1)}), \quad a_{(\tau_1)}, b_{(\tau_1)} > 0
\end{aligned} \tag{3.26}$$

were chosen with shape and rate parameters that yield reasonable expected values and moderate variances, detailed in Table 3.1. The expected value for τ_0 , the y-intercept at zero strain, was chosen visually. Parameter τ_1 was assigned a prior with a mean at the asymptotic back-extrapolated intercept of the stress-strain curves. In contrast, the last two parameters, ξ_0 and ξ_1 , which describe the initial and asymptotic hardening rates, do not require absolute range restrictions. Therefore, diffuse normal priors were chosen for each,

$$\begin{aligned}
\tilde{\pi}(\xi_0) &= N(\xi_0; \mu_{(\xi_0)}, \sigma_{(\xi_0)}^2), \quad \sigma_{(\xi_0)} > 0 \\
\tilde{\pi}(\xi_1) &= N(\xi_1; \mu_{(\xi_1)}, \sigma_{(\xi_1)}^2), \quad \sigma_{(\xi_1)} > 0.
\end{aligned} \tag{3.27}$$

The prior means were elicited by averaging the slope over the entire plastic region from previous studies. Since the hardening rate is expected to be significantly higher than the asymptotic hardening rate, the former was scaled by 1.25 of the average slope and the latter was scaled by .75. We also imposed a hard relative constraint, requiring the hardening rate to be always greater than the asymptotic hardening rate by setting $C(\theta) = \xi_0 - \xi_1$.

Priors on the measurement noise and random effects precisions were defined as in Eqs. (3.12) and (3.11), respectively. Hyperparameters were chosen to be weakly informative. For example, choosing the scale matrix hyperparameter V_0 to the identity matrix and the degrees of freedom to $D + 1 = 5$, leads to a full-conditional distribution (Eq. 3.14) for Λ whose mean is similar to the sample precision (a data-driven outcome). These choices are detailed in Table 3.1.

Table 3 1: Prior hyperparameters for case study I with corresponding prior means and variances.

Parameter	Model	E_θ	V_θ	a_{shape}	b_{rae}
τ_0	Gamma	70	5000	.98	.014
τ_1	Gamma	70	5000	.98	.014
ξ_0	Normal	1800	1×10^6	—	—
ξ_1	Normal	600	1×10^6	—	—
δ^2	Gamma	10	100	1.00	0.100
Λ	Wishart	$5 \cdot \mathbb{I}_4$	$5 \cdot \mathbb{I}_4 + 5 \cdot \mathbb{I}_4$	—	—

CASE STUDY II: THERMODYNAMIC CALCULATIONS

The generality of the framework is illustrated in a second case for an unrelated problem of computational thermodynamics for studying Gibbs free energy using a Redlich-Kister model for a binary solid solution, Ag-Cu. Though the two case study applications are different, they share a similar structure, in that the model parameters are assumed to vary from sample to sample around an overall set of parameters. This makes the random effects framework natural for both applications.

The molar Gibbs free energy of any phase ϕ can be expressed as the sum of six Gibbs energy types,

$$G_m^\phi = G_{ref}^\phi + G_{id}^\phi + G_{ex}^\phi + G_{mag}^\phi + G_{pres}^\phi + G_{surf}^\phi. \quad (3.28)$$

is the weighted sum of the standard Gibbs free energy in the reference state of each constituent under the same temperature, structure, and pressure as the constituent being considered. For a substitutional solid solution with random mixing, this term can be expressed as,

$$G_{ref}^\phi = \sum_{i=1}^n x_i G_i^\phi \quad (3.29)$$

with the elemental Gibbs free energy, G^ϕ , empirically represented as,

$$G_i^\phi(T) = a + bT + cT \ln(T) + \sum_{i=1} d_i T^n. \quad (3.30)$$

Here, x_i is the atomic fraction of each constituent, n is a set of integers usually taking on the values of $[2, 3, -1]$ [165] and $(a, \dots, d_i)$ are adjustable coefficients. G^ϕ is the Gibbs free energy of an ideal solution and is written as,

$$G_{id}^\phi = RT \sum_{i=1}^n x_i \ln(x_i). \quad (3.31)$$

G_{ex}^ϕ is the excess Gibbs free energy and accounts for any contributions not considered by the other terms and is expressed in the following way for a binary system:

$$G_{ex}^\phi = \sum_{\substack{i,j=1: \\ i \neq j}}^n x_i x_j \sum_{z=0}^m L_z (x_i - x_j)^z \quad (3.32)$$

where parameters L_z are the pressure-dependent interaction parameters, also known as the Redlich-Kister parameters [166]. These parameters may also be expressed in a temperature-dependent form,

$$L_z = \sum_{q=0}^2 A_{zq} T^q. \quad (3.33)$$

This form for writing G_{ex}^ϕ uses a polynomial expansion to empirically represent the interaction of phases and does not account for the underlying physics in the system. G_{mag}^ϕ , G_{pres}^ϕ , and G_{surf}^ϕ are, respectively, the magnetic, pressure, and surface contributions to the Gibbs energy and will be neglected for this example [167]. The Gibbs free energy expression from Eq. (3.28) can be reduced to a subregular solid solution model with a reference term,

$$\begin{aligned} G_m^\phi &= G_{ref}^\phi + G_{id}^\phi + G_{ex}^\phi \\ &= \sum_{i=1}^n x_i \left(a_i + b_i T + c_i T \ln(T) + \sum_{i=1} d_i T^n \right) + \\ &\quad \left(RT \sum_{i=1}^n x_i \ln(x_i) \right) + \dots \\ &\quad \left(\sum_{\substack{i,j=1: \\ i \neq j}}^n x_i x_j \sum_{z=0}^m L_z (x_i - x_j)^z \right). \end{aligned} \quad (3.34)$$

CALIBRATION DATA

Ag-Cu forms a simple eutectic system with a reported eutectic temperature of 1051K at 40% atomic fraction Cu with solubility limits of approximately 14 atomic % Cu and 3 atomic % Ag [168, 169] and will serve as our model system. Values from Scientific Group Thermodata Europe (SGTE) are widely used by the CALPHAD community to determine the Gibbs free energy values of pure elements as a function of temperature and pressure [170]. This database was used to establish the elemental Gibbs free energy for pure Cu and Ag. The thermodynamic phase data were collected from the Thermocalc databases ssol and tccu and also the Pandat database PanAl2018 as a proxy for experimental measurement.

The next section details how UQ is used to assess the uncertainty in the unknown Redlich-Kister interaction parameters for a subregular Ag-Cu solution and then propagated via posterior predictions for the Gibbs free energy as a function of composition for 1000 K (≈ 50 K below eutectic temperature).

DETAILS OF THE STATISTICAL MODEL

The random effects for the Gibbs free energy calculations consist of the Redlich-Kister parameters,

$$\theta^{[s]} := (L_0^{[s]}, L_1^{[s]})^T, s = 1, \dots, S, \quad (3.35)$$

where $S = 3$. The first two random effects, $s = 1, 2$, model each of the two Thermocalc databases, while the random effect $s = 3$ corresponds to the Pandat database. Despite their different sources, we assume that these models are related to the overall effect parameters,

$$\theta := (L_0, L_1)^T. \quad (3.36)$$

The distribution of the random effects around the overall effect is modeled directly after Eq. (3.5), without including any restrictions on the parameters, which define a mainly empirical part of the model. The observed and calibration data in this case are the Gibbs free energy as a function of the atomic fraction of copper and are denoted by G , assumed to be contaminated with independent additive Gaussian noise,

$$y^{[s]} = G_{fm}(\theta^{[s]}) + e, \quad e \sim MVN_N(0, \delta^{-2} \mathbb{I}_N), \quad (3.37)$$

resulting in the likelihood model,

$$\begin{aligned} f(y^{[s]} | \theta^{[s]}, \delta^2) &= MVN_N(y^{[s]}; G_{fm}(\theta^{[s]}), \delta^{-2}) \\ &\propto \delta^N \exp \left\{ -\frac{1}{2} \|y^{[s]} - G_{fm}(\theta^{[s]})\|_2^2 (\delta^2 \mathbb{I}_N) \right\}, \\ &s = 1, \dots, S. \end{aligned} \quad (3.38)$$

The joint prior on the overall effect parameters θ and the error precision is modeled as the product of the marginals, following the prior independence assumption used above. L_0 and L_1 are polynomial coefficients without a strict physical meaning, except at zero, which represents an ideal solution where there is no interaction between the constituents. Therefore, we choose a prior that is defined on the entire real line and is centered at 0. A large prior variance further reflects our a priori uncertainty about the value of these parameters. The prior hyperparameters on Λ and δ^2 were chosen to be the same as in the first case study for similar reasons. All prior models along with hyperparameters, prior expectations, and variances are provided in Table 3.2.

Table 3 2: Prior hyperparameters for case study II with corresponding prior means and variances.

Parameter	Model	$\mathbb{E}_\theta$	$\mathbb{V}_\theta$	a_{shape}	b_{rate}
L_0, L_1	Normal	0	1×10^9	—	—
δ^2	Gamma	10	100	1	.1
Λ	Wishart	$3 \cdot \mathbb{I}_2$	$3 \cdot \mathbf{1}_2 + 3 \cdot \mathbb{I}_2$	—	—

METHODS - Bayesian Computer Model Emulation Using a Gaussian Process

Bayesian inference is a flexible approach for modeling and quantifying uncertainty from different sources, and propagating it through to prediction of unknown states. The Bayesian framework is based on modeling degrees of belief about unknown material properties, $\theta = (\theta_1, \dots, \theta_D)^T \in \Theta$, initially by a prior probability model with density $\pi(\theta)$. Beliefs are then updated by conditioning on observed data $y = (y_1, \dots, y_N)^T$, with likelihood $f_m(y | \theta)$, via Bayes' rule. The resulting posterior probability model has density,

$$\pi(\theta | y) = \frac{f_m(y | \theta)\pi(\theta)}{\int_{\Theta} f_m(y | \theta)\pi(\theta)d\theta} \propto f_m(y | \theta)\pi(\theta). \quad (3.39)$$

This general formulation admits conditioning on multiple sources of data, and more complex conditional dependence structures, called Bayesian hierarchical models (BHM, [e.g., ref Berliner]). Having obtained the posterior density, $\pi(\theta | y)$, predictions on new data $\tilde{y}$ are based on the posterior predictive distribution, which for an exchangeable model has density,

$$\pi(\tilde{y} | y) = \int_{\Theta} \pi(\theta | y)f_m(\tilde{y} | \theta)d\theta. \quad (3.40)$$

The data-generating model $f_m(y | \theta)$ in equation (3.39) can be as simple as a probability density centered at θ , or more complex when the output of a physical model is combined with a stochastic error structure (e.g., Figure 3.3). Here, we will focus on a stochastic data-generating mechanism underlying a deterministic materials model $m : \Theta \rightarrow X$ where the model output X lives in a finite-dimensional function space. In this setting, an important practical challenge of performing Bayesian inference is the inability to obtain a closed-form expression for the posterior density (3.39). A related challenge is computing summaries of the posterior distribution that are used for estimation and UQ, such as posterior means, quantiles and highest posterior density intervals. Modern computational methods overcome this issues via Monte Carlo simulation from the posterior, which enables simulation-based estimates of any posterior functionals required for inference. A more general approach is the class of Markov chain Monte Carlo methods, which construct a Markov process with stationary distribution equal to the posterior distribution. However, MCMC sampling requires repeated forward evaluation of the physical model, which is not possible when the model is expensive. A standard approach to overcome this issue is to consider a surrogate for model m , which is learned based on evaluating the model over a fixed set of parameter settings resulting in the input-output pairs $(\eta_1, m_1), \dots, (\eta_J, m_J)$ where $m_j := m(\eta_j)$ and $m = (m_1, \dots, m_J)^T$. Properly characterizing and propagating the epistemic uncertainty in this surrogate model or emulator through to estimation allows for UQ over θ that is exact up to Monte Carlo error.

BHM FORMULATION

Our contribution in this paper is to formulate the surrogate modeling approach as a BHM, which allows us to generalize the methodology to different observation models and to use with physical models with function-valued output. Since the experimental data y and the model evaluations m are conditionally independent given the model m and the parameters θ , we can write the likelihood in the product form $f(y | m, \theta) f(m | m, \psi)$, where ψ is a vector of hyperparameters defining the surrogate model. Assuming that m lives in some finite-dimensional subspace, inference on the unknown material properties θ is based upon the marginal posterior distribution with density,

$$\pi(\theta | y, m) = \int_{\mathcal{X}} \pi(\theta, m | y, m) dm \propto \int_{\mathcal{X}} f(y | m, \theta) f(m | m, \psi) \pi(\theta) \pi(m) dm. \quad (3.41)$$

where the functional form of $f(y | m, \theta)$ does not include the unknown model m , in contrast to the likelihood $f_m(y | \theta)$ in (3.39). The prior $\pi(m)$ over m and the likelihood of the model evaluations $f(m | m, \psi)$ effectively determine the surrogate model that we choose to fit to the model evaluations. Figure 3.3 (b) illustrates this inferential framework, where the difference from the conventional Bayesian method in panel (a) is indicated by the grey color.

MODELS WITH A FUNCTION-VALUED OUTPUT

Many materials models produce function-valued outputs [refs here, introduce VPSC model as an example]. In this section we apply the BHM framework introduced in earlier to this type of model and later consider a specific case study.




Figure 3 3: (a) Diagram illustrates the data generation and computation of posterior distribution for the conventional Bayesian problems. The dashed directed lines indicates the model executions. (b) Visualization of data generation and derivation of posterior distribution for expensive simulators. The grey elements indicate difference from the conventional Bayesian inference problems.

Assume a deterministic materials model $m : \Theta \times T \rightarrow \mathbb{R}$ where $m(\theta, \cdot) \in CK$, where CK is a K -dimensional subspace. Data are observed at locations $t_1^{obs}, \dots, t_n^{obs}$ according to the model,

$$y_n = m(\theta, t_n^{obs}) + \varepsilon_n, \quad (3.42)$$

$$\varepsilon_n \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2), \quad n = 1, \dots, N, \quad (3.43)$$

where θ represents the true unknown input vector. In (3.43) the model evaluation is contaminated with additive zero-mean Gaussian error along the dimension of the variable spatial input $t \in T$. We note that this basic observation structure can be straightforwardly generalized. Equation (3.43) implies a Gaussian likelihood centered around $m(\theta, \cdot), \dots$. A suitable prior probability model will allow us to write the posterior density up to a constant as in (3.41).

Due to the flexibility of basis splines, we can choose a B-spline space that well represents the physical model output space. With the suitable choice of basis spline space, the projection allows us to deal with a finite dimensional objects that represents infinite dimensional curves. With K number of basis spline functions, $B_1(t), \dots, B_K(t)$, the physical model output from the parameter θ formulated as below.

$$m(\theta, t) = \sum_{k=1}^K c_k(\theta) B_k(t), \quad \theta \in \Theta, t \in \mathcal{T}, \quad (3.44)$$

where T is the continuous domain of variable input.

The function-values output, $m(\theta, t)$, is represented by K basis coefficients, $c_1(\theta), \dots, c_K(\theta)$ after the projection. Given that the training output data are evaluated at variable input points $t_1, \dots, t_S$, the generated basis function values are formulated by the matrix B . When the training outputs are generated at the calibration parameter η , the corresponding vector of basis coefficients is a linear transformation of the physical model output vector, $m(\eta, t_1), \dots, m(\eta, t_S)$. The inverse mapping from the basis coefficients to the evaluated model output is a linear projection as shown below.

$$\mathbf{B} = \begin{bmatrix} B_1(t_1) & \dots & B_1(t_S) \\ \vdots & & \vdots \\ B_K(t_1) & \dots & B_K(t_S) \end{bmatrix}, \quad (3.45)$$

$$(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \begin{bmatrix} m(\eta, t_1) \\ \vdots \\ m(\eta, t_S) \end{bmatrix} = \begin{bmatrix} c_1(\eta) \\ \vdots \\ c_K(\eta) \end{bmatrix}, \quad \mathbf{B}^\top \begin{bmatrix} c_1(\eta) \\ \vdots \\ c_K(\eta) \end{bmatrix} = \begin{bmatrix} m(\eta, t_1) \\ \vdots \\ m(\eta, t_S) \end{bmatrix}. \quad (3.46)$$

Given the ground true calibration parameter θ , the true process may be represented by basis coefficients $c_1(\theta), \dots, c_K(\theta)$. As the true process is unknown, we use the projected experimental data, $(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{y}$, which is constructed with the true basis coefficients and projected experimental errors.

$$(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \begin{bmatrix} y(t_1^{obs}) \\ \vdots \\ y(t_N^{obs}) \end{bmatrix} = \begin{bmatrix} c_1(\theta) \\ \vdots \\ c_K(\theta) \end{bmatrix} + (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \begin{bmatrix} \epsilon_1 \\ \vdots \\ \epsilon_N \end{bmatrix}, \quad (3.47)$$

Figure 3.4 illustrates the basis spline projection used in the case study of a crystal plasticity model. We project the function-valued output to a finite dimensional space spanned by five basis functions which are illustrated on the upper right plot. The upper left plot shows function-valued outputs from the ground truth parameter value, which is indicated by the dashed curve, and the pre-generated training parameter values. The five dimensional basis coefficients then define the physical model output. The basis coefficients of the true process are stars on the five plots on the bottom. The correspondence between the model outputs and the projected basis coefficients are distinguished through their colors.

GAUSSIAN PROCESS PRIOR OVER CONTINUOUS DOMAIN

To formulate the Bayesian hierarchical model and calibrate the parameters, the prior distribution of the response surface needs to be specified. As the basis coefficients are functions of calibration parameters, we need a probability model over functions with continuous domains. A Gaussian process is a well-known probability model on function spaces that is adaptable to broad class of the function spaces. Given that $c_1(\theta)$ is a function of θ , a Gaussian process constructs a probability model on $c_1(\theta)$ that is fully defined by a mean function and a covariance function of θ . At a calibration parameter, θ , the distribution of $c_1(\theta)$ is evaluated by the mean function, $\mu : \Theta \rightarrow \mathbb{R}$, and the covariance function, $v : \Theta \times \Theta \rightarrow \mathbb{R}$. The covariance of two different function input values, θ and θ^* , is obtained from the covariance function. There may exist unknown quantities in the mean and covariance functions other than calibration parameter. We call those parameters hyperparameters and denote the vector of hyperparameters by ψ .

$$c_1(\theta)|\theta, \psi \sim \mathcal{N}(\mu(\theta), v(\theta, \theta)), \quad (3.48)$$

$$Cov(\theta, \theta^*)|\theta, \theta^*, \psi = v(\theta, \theta^*). \quad (3.49)$$

For simplicity, we assume that $c_1(\theta), \dots, c_K(\theta)$ are independently and identically distributed as follows.

$$c_k(\theta)|\theta, \psi \stackrel{iid}{\sim} \mathcal{G} \mathcal{P}(\mu(\theta), v(\theta, \theta)), \forall \theta \in \Theta. \quad (3.50)$$

While the prior distribution of training basis coefficients is derived using the Gaussian processes, the distribution of experimental error also has to be concerned for the distribution of projected experimental data. Since we assume a Gaussian distribution for the error term, the projected output also follows a normal distribution.

$$(\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \begin{bmatrix} y(t_1^{obs}) \\ \vdots \\ y(t_N^{obs}) \end{bmatrix} \Bigg| \theta, \psi \sim \mathcal{N}(\mu, v(\theta, \theta) \otimes \mathbb{I}_K + (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \sigma^2), \quad (3.51)$$

where μ is the K dimensional vector $(\mu(\theta), \dots, \mu(\theta))^T$, $\otimes$ indicates the Kronecker product of two matrices, and $\mathbb{I}_K$ is the $K \times K$ identity matrix.

The joint distribution of the training basis coefficients and projected experimental data is in (3.52). Denote a vector of basis coefficients of specific calibration parameter point η_i by $c(\eta_i) = (c_1(\eta_i), \dots, c_K(\eta_i))^T$, $i = 1, \dots, J$ and the projected experimental data by

$$\mathbf{y}^{proj} = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top [y(t_1^{obs}), \dots, y(t_N^{obs})]^\top.$$

$$\begin{aligned} & \begin{bmatrix} \mathbf{c}(\eta^1) \\ \vdots \\ \mathbf{c}(\eta^J) \\ \mathbf{y}^{proj} \end{bmatrix}_{K(J+1) \times 1} \left| \begin{bmatrix} \eta^1 \\ \vdots \\ \eta^J \\ \theta \end{bmatrix}_{(J+1) \times 1} \right., \psi \sim \mathcal{N}(\mu, \Sigma) \text{ where } \mu = \begin{bmatrix} \mu(\eta^1) \\ \vdots \\ \mu(\eta^J) \\ \mu(\theta) \end{bmatrix}_{(J+1) \times 1} \otimes \mathbb{I}_K, \\ \Sigma &= \begin{bmatrix} v(\eta^1, \eta^1) & \dots & v(\eta^1, \eta^J) & v(\eta^1, \theta) \\ \vdots & \ddots & \vdots & \vdots \\ v(\eta^J, \eta^1) & \dots & v(\eta^J, \eta^J) & v(\eta^J, \theta) \\ v(\theta, \eta^1) & \dots & v(\theta, \eta^J) & v(\theta, \theta) \end{bmatrix}_{(J+1) \times (J+1)} \otimes \mathbb{I}_K + \begin{bmatrix} 0 & \dots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & 0 & 0 \\ 0 & \dots & 0 & \sigma^2 [\mathbf{B}(t^{obs})^\top \mathbf{B}(t^{obs})]^{-1} \end{bmatrix}_{K(J+1) \times K(J+1)}, \end{aligned} \quad (3.52)$$

where $\mathbb{I}_K$ is a K -dimensional vector of where all the elements are one, and subscripts of matrix is the dimension of a square matrix or a vector. This joint distribution is a principal tool to calibrate the physical model by training data and derive the statistical inference of the true parameter.

To fully define the Gaussian process prior, a covariance function structure has to be chosen. The choice of covariance function that well explores the function space is critical to calibrate the model of input-output relationship Figure 3.5 illustrates three different covariance functions. The x-axis of plots are Euclidean distances between two input points. The smoothness of covariance function differs based on covariance structures. In the neighborhood of the Euclidean distance close to zero, the exponential covariance function is sharper than the other two covariance functions. While the Matern covariance function and squared exponential covariance function look like a smooth hill, the exponential covariance function is more like a mountain with a sharp edge near zero distance. The exponential covariance function would only consider a very small neighborhood of zero to have high correlation. Outside such neighborhood, the other data are regarded to have rather small covariance. Then, such rough covariance function is likely to poorly explore the function space than smooth functions which takes relatively high correlation between the data in the moderate sized neighborhood. Thus, smooth covariance structures, such as a Matern covariance function with $\nu = 3/2$ or a squared exponential covariance function, captures the function space better than a rough covariance structures, exponential covariance function.




Figure 3 4: The upper left plot shows the model outputs from eight training calibration parameter values (solid curves) and the true process (blue dashed curve) generated in the case study. The upper right plot is the basis functions constructing the basis spline space. The bottom plots show five basis coefficient of eight training data (circles) and the true process (star).




Figure 3 5: Plots of Matern covariance function with $\nu = 3/2$, squared exponential covariance function and exponential covariance function where the x-axis is the Euclidean distance of two input points.




Figure 3 6: Bayesian hierarchical model of calibration for expensive simulators. Directed connection between nodes represent the direction of conditional dependence. The shaded nodes are the observed data and the clear nodes are the unknown quantities.

Figure 3.6 summarizes the Bayesian hierarchical model we presented. The Bayesian model consists of the experimental data, the ground true basis coefficients, the experimental error, and the parameters and the Gaussian process prior of basis coefficients. Conditional on the true process, the experimental data, $y(\text{tobs}), \dots, y(\text{tobs})$, follow a normal distribution where the mean is the true process and the variance is σ^2 . The true process, which is a output of physical model at the true parameter value, θ , is represented by the basis coefficients $c(\theta)$. The crucial element for the surrogate model is the observed paired set of calibration parameter and the response, $(\eta_1, c(\eta_1)), \dots, (\eta_J, c(\eta_J))$. Since we have the correspondence between $\eta_1, \dots, \eta_J$ and $c(\eta_1), \dots, c(\eta_J)$, we use the paired training data to calibrate the relationship between the basis coefficients, and the parameter values. The calibrated model is used to perform uncertainty quantification of the ground true parameter space given the experimental data.

DERIVATION OF POSTERIOR DISTRIBUTION

This section demonstrates the derivation of posterior distribution of calibration parameters and basis coefficients conditional on the training and experimental data. The posterior distributions are estimated through the Markov chain Monte Carlo (MCMC) computation method to quantify the uncertainty by sampling from the posterior distribution. Since there are more than one unknown components in the Bayesian hierarchical model, we adopt the Gibbs sampling method to sample from the joint posterior distribution of the unobserved components, $\pi(\theta, c(\theta) | y, c(\eta_1), \dots, c(\eta_J))$. For each Gibbs sampling step, individual components of interest is sampled conditioning on all the other components by setting the unknown quantities as the latest sampled values. The full conditional distribution of the basis coefficients, $c(\theta)$, has the analytic solution obtained from the joint distribution of basis coefficients of the ground truth state and training data as in the equation (3.52) and is in the appendix.

Figure 3.7 visualizes the full conditional distribution of five basis coefficients in the case study where there are four parameters. In each i^{th} column, only the i^{th} calibration parameter of the training data, θ_i , varies whereas the other calibration parameters are fixed. The k^{th} row represents the full conditional distribution of each k^{th} basis coefficients. The curves represent the point-wise mean and 95% credible bound from the full conditional distribution of each basis coefficient. The full conditional distribution has the credible bound which is close to zero in the neighborhood of each training data point and increases as it gets further from the training data points. The plots illustrate how the physical function space is interpolated accounting for uncertainties through the training data set.




Figure 3 7: Plots illustrate the full conditional distributions of basis coefficients. The blue curves are the mean of full conditional distribution and the red and yellow curves represents the 95% credible bounds.

The full conditional distribution of the calibration parameter, however, does not have an analytic solution. In such cases, the Metropolis-Hastings (MH) algorithm is often used instead of the Gibbs step. The MH method enables sampling from the posterior distribution when the distribution is only known up to a constant of proportionality as below.

$$\begin{aligned} \pi(\theta|c(\theta), y, c(\eta^1), \dots, c(\eta^J)) &= \frac{f(y, c(\eta^1), \dots, c(\eta^J)|\theta, c(\theta))\pi(\theta, c(\theta))}{f(y, c(\eta^1), \dots, c(\eta^J))} \\ &\propto f(y, c(\eta^1), \dots, c(\eta^J)|\theta, c(\theta)) \cdot \pi(c(\theta)|\theta) \cdot \pi(\theta) \end{aligned} \quad (3.53)$$

When we do not have the expression of the normalizing constant, we may generate samples from the proportional distribution using the MH approach. As in the Gibbs step, we only sample the vector θ and consider the other unknown components, $c(\theta)$, as fixed to the previous samples. Denoting $\theta[i]$ the MCMC sample at i th iteration, the method samples a next candidate, θ_{cand} , from a proposal distribution symmetric to the previous sample $\theta[i]$, e.g., a normal distribution

with mean $\theta [i]$. The candidate sample is accepted with the probability of likelihood ratio between the candidate sample and the previous sample as follows.

$$\min \left\{ 1, \frac{\pi(\theta^{cand} | c(\theta)^{[i]}, y, c(\eta^1), \dots, c(\eta^J))}{\pi(\theta^{[i]} | c(\theta)^{[i]}, y, c(\eta^1), \dots, c(\eta^J))} \right\} \quad (3.54)$$

Applying the equation (3.53), the acceptance ratio is computed by the posterior distribution known up to a constant of proportionality. When rejected, the next sample is recorded as the same value of the previous sample, $\theta [i+1] = \theta [i]$.

The proposal distribution for sampling a candidate sample needs to be specified and could influence the estimation of posterior distribution. The adaptive MH MCMC method updates the proposal distribution during the sampling based on the previous samples. That is, for every I iterations, we evaluate the covariance of the past I samples and adapt the proposal covariance based on the previous sample covariance matrix, in the case study simulation we set $I = 1,000$. Moreover, we tune the variance of each parameters by evaluating their acceptance rate so that the MCMC samples well explores the support of the posterior distribution. By setting a target range of acceptance rate, the variance of parameters can be corrected every I iterations; increase (or decrease) the variance if the acceptance rate is lower (or higher) than the target acceptance rate. The full algorithm of the MCMC sampling is demonstrated in the appendix.

ESTIMATION OF HYPERPARAMETERS

To run the MCMC algorithm, the hyperparameters defining the Gaussian process, ψ , need to be specified. The general Bayesian framework would account for the uncertainty of hyperparameters and derive the joint posterior distribution, $\pi(\theta, c(\theta), \psi | y, c(\eta^1), \dots, c(\eta^J))$. Such computation, however, is mostly impractical as the response surface is not known. We adopt the maximum likelihood (ML) approach proposed by Kennedy and O'Hagan. Prior to the MCMC computation, the hyperparameters are estimated through the ML method. The ideal likelihood to maximize would be the conditional probability density of the ground truth basis coefficients given the training coefficients and the hyperparameters. However, derivation of the likelihood is intractable as we do not have the exact input-output relationship of the ground truth basis coefficients. Thus, we use the known input-output correspondence of the training data set estimate the response surface space.

To find estimates that would fit the training data, we apply the cross-validation approach. We randomly split the training data to two groups. One group of data, say group A, is then set to be the training data to estimate hyperparameters where as the other group, group B, is the test data set. The statistic to maximize would be the conditional probability density of the basis coefficients of group A given those of group B and hyperparameters. When the $\{\text{test}_1, \dots, \text{test}_H\}$ indicates the indexes of test data and $\{\text{train}_1, \dots, \text{train}_{J-H}\}$ denotes the indexes of training data for hyperparameter estimation, the estimator of hyperparameters is

$$\arg \max_{\psi} f(c(\eta^{\text{test}_1}), \dots, c(\eta^{\text{test}_H}) | c(\eta^{\text{train}_1}), \dots, c(\eta^{\text{train}_{J-H}}), \psi). \quad (3.55)$$

Methods - Uncertainty Quantification Accounting for Model Discrepancy Within a Random Effects Bayesian Framework

We begin with a brief discussion on the Bayesian paradigm, statistical model, assumptions and simulation methods applied in this work. The interested reader is referred to [175] for a full discourse on the theory and techniques extended in this work. For consistency, the same notation will be adopted.

Bayesian inference is used in this study to recover unknown components of the data generation process, taking into account various sources of uncertainty. The analysis yields a posterior probability distribution describing our understanding about the unknown model components or model parameters, θ , given the available data, y . A statistical model for the data (not to be confused with material or constitutive model of material response) is a collection of probability distributions $M = \{f(y|\theta); \theta \in \Theta\}$, where Θ is the parameter space to which θ belongs. The first ingredient in a Bayesian analysis is a prior distribution with density $\pi(\theta)$ over the unknown components, which is simply an epistemic model of our subjective uncertainty before any data is observed. The prior distribution is then updated by conditioning on observed data to establish the posterior distribution through Bayes' rule,

$$\pi(\theta|y) = \frac{f(y|\theta)\pi(\theta)}{\int_{\Theta} f(y|\theta)\pi(\theta)} \propto f(y|\theta)\pi(\theta) \quad (3.56)$$

The posterior distribution is the updated belief about θ after data has been observed and is proportional to the prior belief about θ and the likelihood of observing data y given θ , $f(y|\theta)$. The same construct can be extended to make predictions about future observations, y^* to establish the posterior predictive distribution given observed data y . This is a continuous mixture of the likelihood of the unobserved data y^* weighted by the posterior of the unknown model parameters θ ,

$$\pi(y^*|y) = \int_{\Theta} f(y^*|\theta)\pi(\theta|y)d\theta \quad (3.57)$$

The posterior and posterior predictive distributions are all that is required for inference and prediction, respectively. Although posterior densities are often not analytically tractable, they can be estimated from a Monte Carlo sample. Markov chain Monte Carlo (MCMC) [62–65] is a class of simulation techniques for constructing a Markov chain with stationary distribution equal to the posterior distribution being targeted. Monte Carlo estimates of desired posterior functionals, such as the mean, mode, and highest posterior intervals (HPIs), can then be computed from the resulting sample.

Our MCMC sampling approach consists of an adaptive Metropolis-Hastings algorithm [69, 70] with conditional Gibbs updates [67] where full conditional distributions are available.

MODELING ASSUMPTIONS

We consider a random effects model, which naturally accounts for the variability observed between multiple material samples tested. This statistical model is an excellent candidate for cases where the data which are observed do not directly represent the true underlying property, θ , which can also be called the overall effective property. Instead, each material sample is tested under a state which deviates slightly from the true state and may therefore be modelled as randomly distributed about the overall effect. The sample states are called random effects and are denoted by $\theta[s]$, $s = 1, \dots, S$, where S is the number of samples tested. Though they are not of interest in and of themselves (they are often referred to as

nuisance parameters), they must be estimated in order to infer the overall effect of interest θ .

The random effects, $\theta [s]$, $s = 1, \dots, S$, and overall effect, θ , are related through the following probability model,

$$\theta^{\{[s]\}} | \theta, \Lambda \sim f(\theta^{\{[s]\}} | \theta, \Lambda), \quad s = 1, \dots, S, \quad (3.58)$$

which also depends on additional parameters Λ to define the distribution. Though the choice of this distribution is application-specific, a multivariate Normal (or Gaussian) distribution is often a flexible choice to model the distribution of random effects around the overall effect since it implies that 1) The random effects are symmetrically distributed about the overall effect and 2) The random effects are more likely to lie in a region that is close to the overall effect θ . Thus, for D -dimensional parameters θ , $\theta [s]$, $s = 1, \dots, S$, we assume,

$$\theta^{\{[s]\}} | \theta, \Lambda \sim MVN_D(\theta, \Lambda^{-1}), \quad s = 1, \dots, S, \quad (3.59)$$

where Λ is a $D \times D$ symmetric positive-definite inverse variance-covariance (precision) matrix, which controls how tightly dispersed the random effects are about the overall effect θ .

A commonly used prior distribution for the precision Λ is the Wishart, a distribution on symmetric positive-definite matrices. Since this choice does not allow precise control over the prior model through the hyperparameters, we take a more flexible prior modeling approach. The random effects precision Λ is decomposed into two components: a $D \times D$ positive-definite and symmetric correlation matrix, R , and a $(1 \times D)$ vector of standard deviations, $t = (t_1, \dots, t_D)^T$, which are then modelled separately following [126, 127]. The decomposition,

$$\Lambda := \text{diag}(t) R \text{diag}(t), \quad t = (t_1, \dots, t_D)^T, \quad (3.60)$$

allows extra flexibility in modeling known features of the distribution while allowing enough dispersion. We assign a Wishart prior to the correlation matrix, R , which preserves conjugation and allows R to be sampled via Gibbs updates,

$$\begin{aligned} \pi(R) &= \text{Wishart}_D(R; R_0, r_0) \\ &\propto |R_0|^{-r_0/2} |R|^{(r_0-D-1)/2} \exp\left\{-\frac{1}{2} \text{tr}(r_0^{-1} R)\right\}, \\ &r_0 \geq D + 1, \end{aligned} \quad (3.61)$$

where R_0 is a $D \times D$ symmetric positive-definite scale matrix and r_0 represents the distribution's degrees of freedom. For computational convenience, the variances, t_d^2 , $d = 1, \dots, D$ rather than the standard deviations are modeled as,

$$t_d^2 \sim \text{Gamma}(a_t, b_t), \quad d = 1, \dots, D, \quad (3.62)$$

where a_t and b_t are fixed hyperparameters.

For each of the S samples tested, the likelihood of the N -dimensional data $y^{[s]}$ given the physical model $m(\theta^{[s]})$ evaluated at the corresponding random effect is modeled as a multivariate Normal distribution,

$$y^{[s]} | \theta^{[s]}, \Psi \sim MVN_N(m(\theta^{[s]}), \Psi^{[s]}), \quad s = 1, \dots, S, \quad (3.63)$$

centered about $y^{[s]}$ with dispersion defined through the $N \times N$ error precision matrix Ψ .

MODEL MISFIT

So far, a key assumption has been that the physical model m adequately describes the mean of the data-generating process. In reality this is either not known with certainty, or is known and the structure of this model discrepancy is understood to some degree. To capture both these scenarios, we denote by $\zeta(x)$ the 'ground truth', or true state of the system, at location x and we make the assumption that this signal has been contaminated with additive Gaussian error e_n , $n = 1, \dots, N$ corresponding to observation locations $x_1, \dots, x_N$, representing observation error or some other type of variability,

$$y_n = \zeta(x_n) + e_n, \quad n = 1, \dots, N. \quad (3.64)$$

Under this assumption the model of the system, m , which accepts a vector (or scalar) of parameters θ and is evaluated at location x , is representative of the true process such that $m(x, \theta) = \zeta(x)$. This assumption was made in [175]. The present work diverges here and will adopt a different, more realistic, data-generating model which accounts for the effect of model discrepancy. If we take into account the discrepancy between the model representation of the system, $m(x, \theta)$, and the true mean process, $\zeta(x)$, and denote it by Δ , we can relate these terms as:

$$\zeta(x) = m(x, \theta) + \Delta(x). \quad (3.65)$$

The model for the observations now becomes,

$$y_n = m(x_n, \theta) + \Delta(x_n) + e_n, \quad n = 1, \dots, N, \quad (3.66)$$

where $m(\cdot)$ and $\Delta(\cdot)$ are assumed independent and θ is the true but unknown vector of calibration parameters.

Among the first to incorporate model discrepancy into an inferential framework, is the work of Kennedy and O'Hagan [35] in order to avoid estimation bias and over-confidence in the estimates when conducting inference by constructing reduced-order versions of expensive simulators. [51] They illustrate similar issues arising from a failure to take into account systematic model discrepancy in general inference problems. Importantly, the role of the discrepancy model and its link to identifiability is explored.

With all components of the statistical model now introduced, a schematic of the statistical model can be found in Fig. 4.2, in which the direction of the arrows represents the direction of conditional dependence between model components. Shaded nodes indicate that model components are observed, with everything else fixed and unknown. Under the assumption of conditional independence from sample to sample of the observations $y^{[s]} = (y_1^{[s]}, \dots, y_N^{[s]})^T$ given the random effects $\theta^{[s]} \in \mathbb{R}^D$, the density of the posterior distribution is,

$$\begin{aligned} \pi(\theta, \{\theta^{[s]}\}_{s=1}^S, R, t^2, \Psi, \Delta | \{y^{[s]}\}_{s=1}^S) \\ \propto \pi(\theta, R, t^2, \Psi, \Delta) \prod_{s=1}^S f(y^{[s]} | \theta^{[s]}, \Psi, \Delta) f(\theta^{[s]} | \theta, \Lambda), \end{aligned} \quad (3.67)$$

where $f(y^{[s]} | \theta^{[s]}, \Psi, \Delta)$ is the density of,

$$y^{[s]} | \theta^{[s]}, \Psi, \Delta \sim \text{MVN}_N(m(\theta^{[s]}) + \Delta, \Psi^{-1}). \quad (3.68)$$

A flexible prior model for the discrepancy term Δ is discussed in detail in Section 4.2.3.

Propagating this posterior uncertainty through to prediction of the data leads to the posterior density,

$$\begin{aligned} \pi(y^{[new]} | \{y^{[s]}\}_{s=1}^S) \\ = \int f(y^{[new]} | \theta^{[new]}, \Psi, \Delta) f(\theta^{[new]} | \theta, \Lambda) \\ \pi(\theta, \Lambda, \Psi, \Delta | \{y^{[s]}\}_{s=1}^S) \\ d\theta d\theta^{[1]} \dots d\theta^{[S]} d\Lambda d\Psi d\Delta d\theta^{[new]} \end{aligned} \quad (3.69)$$

GAUSSIAN PROCESS PRIOR FOR MODEL DISCREPANCY

As an unknown model component, the discrepancy function Δ must be assigned a prior model within the Bayesian hierarchy. The importance of carefully incorporating all prior knowledge about Δ into this probability model was demonstrated by Brynjarsdóttir and O'Hagan [51], where bias and over-confidence in the estimates resulted from using priors that

were too flexible. The problem, called lack of identifiability in the statistical literature, results from the ability of different model components (in this case the discrepancy function and model parameters) to trade off in ways that can describe the data equally well. When these regions are disjoint or lie on a low-dimensional manifold in the parameter space, this hinders our ability to obtain useful estimates of the model components. One way to overcome lack of identifiability is by defining sufficiently informative prior models, subject of course to the availability of such information. Informally, an informative prior choice serves to penalize different configurations or trade-offs and guide the posterior towards regions that we know a-priori to be more probable or physically reasonable.

A Gaussian process (GP) is a stochastic process, which is a collection of random variables such that any finite sub-collection of those variables have a jointly Gaussian distribution. For an accessible introduction to these models, the reader is referred to [34]. GPs can be used to define prior distributions over functions, such as our discrepancy function $\Delta : D \rightarrow \mathbb{R}$. They can be thought of as a functional generalization of the multivariate Gaussian distribution, which is characterized by a mean vector and positive definite covariance matrix. Likewise, a GP is fully characterized by a mean function, $\mu : D \rightarrow \mathbb{R}$ and a positive definite covariance function, $c : D \times D \rightarrow \mathbb{R}$. Therefore, we can specify a GP prior model for the discrepancy function Δ as,

$$\Delta \sim GP(\mu, c) \quad (3.70)$$

It is often realistic to model the discrepancy as having prior mean function $\mu(x) = 0$, as this incorporates the neutral and often reasonable assumption that the magnitude of the systematic deviation from the physical model across the domain is not known a-priori. The choice of the covariance function is an important modeling consideration, as it controls the degree to which correlation decays with distance between inputs over the input domain. As such, the covariance controls the smoothness of realizations (called sample paths) of a GP. More information on this choice is described in another section.




Figure 3 8: Visualization of the Bayesian random effects model with discrepancy. Arrows represent the direction of conditional dependence. Clear nodes represent unobserved/unknown quantities while shaded nodes indicate observations.

A convenient feature of GPs is the joint Gaussianity of any finite sub-collection of points. While the function Δ is in theory continuous, it can be probed at a finite set of variable inputs since

$$\Delta(x_1), \dots, \Delta(x_N) \sim \text{MVN}_N(0, C),$$

(3.71)

where 0 is the N -dimensional zero vector and,

$$C = Cov \left((\Delta(x_1), \dots, \Delta(x_N))^T, (\Delta(x_1), \dots, \Delta(x_N))^T \right), \quad (3.72)$$

and the covariance between the process at any two points x_i and x_j is related to the covariance function through $c(x_i, x_j) = Cov(\Delta(x_i), \Delta(x_j))$, for $x_i, x_j \in \{x_1, \dots, x_N\}$.

CHOICE OF COVARIANCE STRUCTURE

Since the covariance structure controls the correlation scale and smoothness of the stochastic process, careful choice of this function is required to specify a sufficiently informative prior process for the model discrepancy Δ . The more slowly the covariance between $\Delta(x_i)$ and $\Delta(x_j)$ decays with distance between x_i and x_j , the smoother are the sample paths of the resulting Gaussian process. One example of a smooth covariance is the squared exponential function,

$$\begin{aligned} Cov(\Delta(x_i), \Delta(x_j)) &= c(x_i, x_j) \\ &= \sigma^2 \exp\left\{-\frac{1}{2} \left(\frac{x_i - x_j}{w}\right)^2\right\}, \end{aligned} \quad (3.73)$$

where σ^2 is a variance parameter and w is the length-scale. A GP with this covariance structure has infinitely differentiable sample paths. Other popular covariance functions include the Matern [176], whose smoothness can be controlled by the choice of model parameters, and the exponential [177], which yields relatively rough sample paths.

When the GP is used as a prior model, the variance parameter describes prior uncertainty about the unknown function through the variability of resulting sample paths. A large variance parameter corresponds to a diffuse GP prior distribution, while a smaller variance produces a tighter distribution appropriate when more prior information is available. The length-scale parameter controls how dependence scales with the distance between nearby points. As with all prior hyperparameters, these must either be set by the user or can be treated as unknown and estimated from the data. Fig. 4.4 illustrates the GP sample paths under a selection of different hyperparameter choices.

A stationary covariance function is one that depends only on the distance between the inputs and is independent of their relative location. Non-stationary covariance models, on the other hand, allow for changes to the dependence structure across the input domain.

Fig. 4.3 illustrates the difference between the behavior of sample paths from a GP with a stationary (left) and a non-stationary (right) covariance. The samples in Fig. 4.3b initially have stronger dependence, which decreases as x moves away from 0.




Figure 3 9: Multiple sample functions drawn from a GP prior with a smooth covariance structure. The cases shown are a) a stationary covariance function and b) a non-stationary covariance function.

TOY EXAMPLE

As a simple example, suppose we have measured the mass of an ice block (100 kg), apply an unknown force to it, and then collect observations of the position of the block at different time increments. The ice block accelerates from an initial velocity and position of zero, and we wish to infer from the observations the force which was applied using the fundamental equation,

$$m(t, \theta) = \frac{1}{2} \left(\frac{\theta}{mass} \right) t^2, \quad \ni t, \theta > 0 \quad (3.74)$$

where the unknown force applied is denoted by θ .




Figure 3 10: Multiple sample functions drawn at random from the prior specified by a GP which favors smooth functions and a stationary covariance function with parameters a) $w = .025, \sigma^2 = 10$ b) $w = .25, \sigma^2 = 10$, c) $w = .025, \sigma^2 = 100$, d) $w = .25, \sigma^2 = 100$

However, suppose that in reality the ice block is melting and the mass is changing with time, and the position of the block under a constant applied force follows the relationship,

$$\zeta(t) = \frac{1}{2} \left(\frac{\theta_{true}}{mass \cdot \exp(-\lambda t)} \right) t^2, \quad \exists, \theta > 0 \quad (3.75)$$

where mass indicates the initial mass of the block and λ is a factor controlling the rate of the melt. In this example, there is knowledge missing from the model of the physical system (4.19), the changing mass of the ice block, which will induce discrepancy in the inference on the force which is being applied. A comparison of the true behavior, the model of the behavior, and the observations is provided in Fig. 4.5. The observant reader might notice some observations indicate a negative position due to the noise; while this example is not perfect, it serves the purpose of emphasizing the importance of considering model misfit in the analysis. By comparing the model of the system and the true behavior, it is clear that at the initial time of 0, the true behavior and the model are identical and the discrepancy between them grows as time goes on. Analyses conducted with and without consideration of the model discrepancy reveal the importance of including this component in the statistical model.





Figure 3 11: Schematic of the toy problem showing the true behavior ζ (blue line), the model of the behavior $m(t, \theta)$ (magenta line) and the observations y_n (open circles).

Figure 3.12 shows marginal histograms and posterior draws (blue markers) from two different analyses performed to estimate the unknown model parameter (θ , which is the applied force) by conditioning on the available data (red markers). A comparison of the two illustrates the need to account for model discrepancy in the statistical model. If discrepancy between the model and the true process is not considered, as in Fig. 3.12a, posterior inference suffers from both bias (a systematic difference between the estimator and the true value) and under-coverage (over-confidence in the results as measured by the amount of spread in the posterior distribution over θ). In comparison, Fig. 3.12c shows results from an analysis in which discrepancy was incorporated through an additive function which was assigned a GP prior with non-stationary covariance. In this case, our belief about the value of θ as described by the location and spread of the posterior density is much closer to the ground truth, which is now covered by the central 95% credible interval. It is interesting to note that this bias-variance trade-off [178] is a commonly observed phenomenon, whereby a predictive model with lower bias has a higher variance and vice versa. This dilemma is brought on by the conflict of trying to simultaneously reduce both sources of error (bias and variance) and finding a balance between over-fitting and under-fitting the training data.


With a clear understanding of the importance of considering model misfit in the inference problem, we now progress to discuss the materials application to which we will be applying these techniques.

CASE STUDY


Reduced-order, phenomenological and homogenized models are well-suited for Bayesian inferential analyses using MCMC simulations due to their low computational cost. In exchange for their simplified formulations and reduced computational cost (when compared to 3D, physics-based and full-field models), these models may leave out certain physics affecting the material behavior or process of interest or may make other simplifying assumptions. It may also be the case that a model of higher fidelity may not be available for a particular application. The bottom line is that while these models may be very reliable over a subset of the domain in their representation of the process or behavior, they may not be reliable over the entire domain (e.g., a model which captures plastic deformation well but does not consider elastic deformation). Consequently, this may introduce a systematic discrepancy over a restricted domain, or a systematic discrepancy from the true process over the entire domain.




(a)



(b)



(c)



(d)

Figure 3 12: A comparison of inference on the toy example without (a & b) and with (c & d) the contribution of model misfit.

However, this discrepancy need not prevent the use of reduced- order models as reliable tools in the prediction of material behavior. With the incorporation of model-form error into the statistical inferential problem, their parameters can be reliably inferred without the need to restrict the domain of the model.

Take, for example, the phenomenological visco-plastic self-consistent (VPSC) crystal plasticity model developed by Tome´ and Lebensohn at Los Alamos National Laboratory [1]. This model is commonly used to understand deformation and texture evolution at large plastic strains, however it does not include physics to consider the elastic regime of deformation. Because of this, a systematic discrepancy is present between the model and experimental observations of stress-strain response as they are collected. Reliable calibration of this model without restricting the domain requires the consideration of model-form error.

The VPSC crystal plasticity model uses the Voce hardening law [159, 160] to describe the evolution of the critical resolved shear stress at the level of the grain. This phenomenological relationship evolves the critical resolved shear stress (CRSS) as a function of the accumulated shear across all slip systems, α , in the grain. Here, the incremental shear is denoted by $d\Gamma$, and unknown parameters in the model are the initial critical resolved shear stress (CRSS), τ_0 , the asymptotic CRSS increase τ_1 , the initial hardening rate ξ_0 and the asymptotic hardening rate ξ_1 . 4 Accurate description of hardening in the model is dependent on reliable parameter estimation within the law. The hardening rates are traditionally represented by θ_0 and θ_1 in the Voce hardening law. However, to avoid confusion with parameters in the statistical model, this notation is adopted as in [175]

$$\hat{\tau}^\alpha = \tau_0^\alpha + (\tau_1^\alpha + \xi_1^\alpha \Gamma) \left(1 - \exp \left\{ -\Gamma \left| \frac{\xi_0^\alpha}{\tau_1^\alpha} \right| \right\} \right). \quad (3.76)$$

based on physics but do have physical meaning and whose values can be extrapolated outside the bounds of calibration to consider various cases of deformation. The VPSC model produces the final texture evolution and the stress-strain response of a material undergoing deformation by homogenizing the local response to obtain the macroscopic response of the material.

One of the two case studies presented in Ricciardi et al. [175] consider inference on the unknown Voce hardening parameters and propagate this uncertainty to induce a predictive distribution in the stress-strain response enabling prediction for experimental data not yet observed. Observational error as well as sample-to-sample variability were considered in performing inference on the parameters while conditioning on observations from only the plastic regime of deformation. Conditioning on this restricted domain allowed the reasonable assumption that the VPSC model was correct despite its drawbacks in the elastic regime of deformation.

However, in general, we wish to condition on the full experimental data collected without restricting the observation domain (i.e., with elastic and plastic data), which may risk introducing bias and under-coverage in any parameter estimates when model misfit is not explicitly considered. The missing physics in the VPSC model produces a systematic discrepancy between the model and observations. Therefore, it can no longer be calibrated under the assumption that it forms the mean of the observation process across both elastic and plastic regions. Including model misfit in the analysis will account for this discrepancy, reducing bias in the estimated parameters and propagated model output.

To demonstrate the necessity of including the model discrepancy in the inference problem, an analysis was performed by conditioning the VPSC model on both the elastic and plastic regimes of deformation without consideration of model misfit. As pointed out by Brynjardo'ttir [51], the posterior distribution of θ centers around a 'posterior best fit' value which minimizes the mean residual error between the model output and the observations. Thus, ignoring model misfit introduces bias in the parameter estimates and consequently the model trajectory, as illustrated in Figure 3.13. The observations of the state are in red and 200 samples drawn from the posterior distribution and propagated through the model to induce a distribution on the true stress-strain behavior given the data (shown in black). The evident bias is particularly problematic in cases such as this where model parameters are not simply tuning parameters but

whose values are of intrinsic interest and necessary for reliable prediction, in both cases of interpolation and extrapolation.




Figure 3 13: Analysis for VPSC conditioned on the elastic + plastic regimes of deformation not accounting for model discrepancy. The red points are observations and the black lines show approximately 200 samples from the posterior distribution.

STATISTICAL MODEL

We assume the true underlying behavior of the system is represented by overall effect, $\theta = \{\theta_1, \dots, \theta_D\}$, which is comprised of a D -dimensional vector of unknown model parameters. Yet, due to the random nature of microstructural heterogeneity, observations from each sample are generated under slightly different states which are therefore modeled as random effects. Model parameters corresponding to sample $s \in \{1, \dots, S\}$ (where S is the total number of samples) are denoted by $\theta^{[s]} = \left(\theta_1^{[s]}, \dots, \theta_D^{[s]}\right)^T$. A D -dimensional multivariate normal distribution is used to model the relationship between the random effects and the overall effect,

$$\begin{aligned} f(\theta^{[s]} | \theta, \Lambda) &\propto MVN_D \left(\theta^{[s]}; \theta, \Lambda^{-1} \right) \mathbb{I}\{\mathcal{C}(\theta) > 0\} \\ &\propto \det\{\Lambda\}^{1/2} \exp \left\{ -\frac{1}{2} \left\| \theta^{[s]} - \theta \right\|_2^2 (\Lambda) \right\} \\ &\cdot \mathbb{I}\{\mathcal{C}(\theta) > 0\}, \quad s = 1, \dots, S. \end{aligned} \tag{3.76}$$

The decomposition of the precision matrix $\Lambda = \text{diag}(t) R \text{diag}(t)$, allows us to define a flexible prior model separately on the correlation matrix, R , and the standard deviations, $t = (t_1, \dots, t_D)^T$. Physical constraints on the model parameters are introduced through the indicator function $I\{C(\theta) > 0\}$, where $C(\theta) = (\tau_0, \tau_1, \theta_0 - \theta_1)^T$ ensures a condition where there is no softening in the material, and the asymptotic hardening rate reaches a limiting value.

The model for the observations follow Eq. (3.65) and we assume i.i.d. Gaussian error at each variable input represented by δ^2 such that, $\Psi = \delta^2 I_N$. In other words, we assume the observational error contribution is the same across all variable inputs. The model for the observations becomes,

$$\begin{aligned} y^{[s]}(x_n) &= m(x_n, \theta^{[s]}) + e_n + \Delta(x_n), \\ e_n &\sim MVN_N(0, \Psi^{-1}), \\ n &= 1, \dots, N, \quad s = 1, \dots, S, \end{aligned} \tag{3.77}$$

where Δ is an unknown discrepancy function. A multivariate normal distribution is assumed for the likelihood of the observations centered at the sum of the model plus the discrepancy with covariance matrix Ψ^{-1} .

$$\begin{aligned} f(y^{[s]} | m(\theta^{[s]}), \Psi, \Delta) &\propto MVN_D(y^{[s]}; m(\theta^{[s]}) + \Delta, \Psi^{-1}) \\ &\propto \det\{\Psi\}^{1/2} \exp\left\{-\frac{1}{2} \|y^{[s]} - m(\theta^{[s]}) - \Delta\|_2^2(\Psi)\right\}, \\ s &= 1, \dots, S. \end{aligned} \tag{3.78}$$

Diffuse priors were assigned to the overall effective parameters in θ , reflecting our prior uncertainty about the spread of the sample parameters from their center. The priors on τ_0 and τ_1 were chosen to be Gamma distributed with hyperparameters specifying a mean of 70 for both and a large variance,

$$\begin{aligned} \pi(\tau_0) &= \text{Gamma}(a_{\tau_0}, b_{\tau_0}) \\ \pi(\tau_1) &= \text{Gamma}(a_{\tau_1}, b_{\tau_1}). \end{aligned} \tag{3.79}$$

Priors on ξ_0 and ξ_1 were modeled as diffuse normal distributions,

$$\begin{aligned} \pi(\xi_0) &= \text{Normal}(\mu_{\xi_0}, \sigma_{\xi_0}^2) \\ \pi(\xi_1) &= \text{Normal}(\mu_{\xi_1}, \sigma_{\xi_1}^2). \end{aligned} \tag{3.80}$$

Prior means were chosen based on previous studies. The error precision, δ^2 is modeled with a Gamma distribution in order to ensure that only positive real values are possible. A Wishart prior is chosen for the correlation matrix, R , since it is a symmetric positive-definite matrix and the variances, t^2 , $d = 1, \dots, D$ are assigned Gamma distributions. The prior on the unknown discrepancy function, Δ , is zero-mean Gaussian process with a custom covariance structure discussed below and shown in Fig. 4.8. Since the likelihood model (3.78) only requires Δ to be evaluated at N discrete points, this prior becomes a multivariate normal distribution centered at zero with variance-covariance matrix Γ , obtained by evaluating the GP covariance structure at the vector of observation locations $(x_1, \dots, x_N)^T$. Summaries of the prior distributions are provided in Table. 3.3

Table 3 3: Prior hyperparameters for the case study with corresponding prior means and variances.

Parameter	Model	$\mathbb{E}_{\theta}$	$\mathbb{V}_{\theta}$	a_{shape}	b_{rate}
τ_0	Gamma	70	5000	.98	.014
τ_1	Gamma	70	5000	.98	.014
ξ_0	Normal	1800	1×10^6	–	–
ξ_1	Normal	600	1×10^6	–	–
δ^2	Gamma	10	100	1.0	.100
Δ	MVN_N	0	Γ	–	–
R	$Wishart_D$	$(D+2) \cdot \mathbb{I}_D$	$(D+2) \cdot 1_D + (D+2) \cdot \mathbb{I}_D$	–	–
t_d^2	Gamma	10	100	1.0	.100

Since the VPSC model uses a constitutive equation that only accounts for plastic deformation and does not consider the elastic regime of deformation, a large discrepancy is expected in the elastic region, decreasing as we get closer to the elastic-plastic transition, and the model is expected to have little to no systematic misfit in the plastic region. Because of the contrast in the expected structure of the discrepancy for these two different regimes, we select a piece-wise formulation for the prior model. The strain at the elastic-plastic transition also marks the transition in the piece-wise covariance function,

$$\Gamma_{i,j} = \begin{cases} \text{Cov}_{el}(x_i, x_j), & i, j \in \mathbf{N}_{el}, \\ \text{Cov}_{pl}(x_i, x_j), & i, j \in \mathbf{N}_{pl}. \end{cases} \quad (3.81)$$

We denote the strain at the elastic-plastic transition by t . This value is traditionally the strain corresponding to the yield stress at .2% offset. While this value will differ between samples, we take this value to be fixed at .0025. In Eq. (3.81), $\mathbf{N}_{el} = \{1, \dots, t\}$ is the set of indices within the elastic regime, and $\mathbf{N}_{pl} = \{t+1, \dots, N\}$ is the set of indices of plastic strain.

For Cov_{el} we choose a non-stationary smooth covariance function since the greatest discrepancy is expected to occur at $\varepsilon = 0$ and to taper to zero at ε_t . While the length scale and variance of the covariance function can be included as unknown components of the statistical model to be inferred, here we assign values $w_{el} = 1.4 \times 10^{-4}$ and $\sigma_{el}^2 = 1.0 \times 10^{14}$.

A stationary covariance structure is chosen for the plastic region with a length scale of $w_{pl} = .02$ and a variance of $\sigma_{pl}^2 = 1.0 \times 10^{-4}$. A small prior variance is chosen since we expect very little misfit in the plastic regime. Essentially, the prior penalized deviations of the model from the observations in this region. Figure 3.14 shows multiple draws from the prior of Δ for both the elastic and plastic regions as well as over the full domain.

Bayesian Computer Model Emulation Using a Gaussian Process: CASE STUDY: VPSC MODEL

This section describes a simulation study of phenomenological visoplastic self-consistent (VPSC) CP model.

STATISTICAL MODEL

There are four calibration parameters in the model, $\theta = (\theta_1, \dots, \theta_4)^T = (\tau_0, \tau_1, \xi_0, \xi_1)^T$. Experimental data was generated with a set of parameter value, $\theta = (55, 55, 700, 200)$, at equally spaced strain values with an experimental error term.

$$y(t_n^{obs}) = m(\theta, t_n^{obs}) + \varepsilon_n, \quad (3.82)$$

$$\varepsilon_n \stackrel{iid}{\sim} \mathcal{N}(0, 0.5^2), n = 1, \dots, 200. \quad (3.83)$$

For simplicity, we assume a constant mean function and an isotropic covariance function with a small nugget; the covariance structure is the same for all the locations. The covariance function is assumed to be a squared exponential covariance function.

$$\mu(\theta) = 0, \forall \theta \in \Theta, \quad (3.84)$$

$$v(\theta, \theta^*) = \begin{cases} \lambda^2 \exp\left(-\sum_{i=1}^4 \frac{(\theta_i - \theta_i^*)^2}{\zeta_i}\right), & \text{if } \theta \neq \theta^*, \forall \theta, \theta^* \in \Theta. \\ \lambda^2 + \tau^2, & \text{if } \theta = \theta^*, \end{cases} \quad (3.85)$$

The 400 training data were generated from a wide range of parameter values and they were evaluated at 200 equally spaced strain points from 0 to 0.5. The training output data and the experimental data were projected to 5-dimensional basis spline function space. Using the training data, we estimate the hyperparameters, $\psi = (\lambda, \zeta, \tau)^T$, by the maximum likelihood method. Applying the ML estimate of the hyperparameters, the MCMC sampling method is implemented as the algorithm explained in the appendix. During the sampling, the covariance of the proposal distribution is adjusted based on the acceptance rate and the correlation of the past 1,000 samples. The MCMC output provides the estimate of posterior distribution of calibration parameters and the basis coefficients. To decrease the influence of initial values and the previous sample, we disregard a portion of the samples and record only 5th iterations.




Figure 3 14: The black dots are the experimental data. The red line is the basis spline and the blue line is the VPSC output where the corresponding iteration generates the mode posterior. The subplot is the a part of the plot zoomed in.

MCMC Implementation

A Bayesian framework for uncertainty quantification implementing a random effects model was applied to various modeling applications in materials science. These include parameter estimation and response prediction within the phenomenological VPSC model[1] without and with the consideration of model-form error, as well as UQ for calculations of phase equilibrium.

These studies were carried out with MCMC simulation methods by which approximate posterior samples were generated through the construction of a Markov chain with a stationary distribution equal to the posterior distribution of interest. This allowed posterior densities, which were not analytically tractable, to be recovered and for the extraction of useful summaries about the distribution (i.e., posterior modes, means, and quantiles) needed to communicate posterior uncertainty. The MCMC simulations were used to approximate the posterior distribution over the model parameters of interest and to also make predictions about unobserved data.

The appendix contains details about the MATLAB code used for the MCMC simulation which considered UQ and UP in the phenomenological VPSC model with the consideration of discrepancy, modeled with a Gaussian process. The main file is contained in the appendix, and functions for the log likelihood, log posterior, posterior and posterior predictive draws, and proposal covariance adaptation can be found in the appendix.

The model and general modeling assumptions are the same as those explicated earlier. However, for easy reference, the schematic showing this information is included in here in Figure 3.15. The reader is referred to Section 3 for full details on the parameters and model details.

Parameters for which conjugacy exists between the likelihood and prior, thus making the posterior analytically tractable, were sampled in Gibbs updates. Full-conditional distributions for these parameters were derived and are reported in the appendix, and include the discrepancy, Δ , in Eq. (B.4), the decomposed correlation of the random effects precision matrix, R , in Eq. (B.3), and the error precision, δ^2 in Eq. (B.1). Parameters for which full-conditional updates were not possible were sampled with an adaptive Metropolis-Hastings algorithm. These include the random effects parameters, $\theta^{[s]}$, $s = 1, \dots, S$, the overall effect parameters, θ , and the decomposed variance of the random effects precision matrix, t^2 . Figure 2.4 from Section 2 shows a schematic of the Metropolis-Hastings algorithm, which is included within this chapter in Fig. 3.16.

The first modeling consideration was to define all priors and hyperparameters. Next, the vector of unknown model parameters which were updated with the MH algorithm was initiated, $\Phi^{(0)} = (\theta^{[1](0)}, \dots, \theta^{[S](0)}, \theta, t^{2(0)})$ and partitioned into $B=S+2$ blocks denoted by $\Phi^{1(0)}, \dots, \Phi^{B(0)}$. There was one block for each of the random effects, one for the overall effect and one for the decomposed variance of the

random effects precisionmatrix. Parameters updated with Gibbs sampling were also initiated at this point, $R^{(0)}$, $\Delta^{(0)}$, and $\delta^{2(0)}$, as well as the proposal covariance matrices, Σ^b , $b = 1, \dots, B$, one for each block. At each iteration, $m = 1, \dots, M$, one block of parameters at a time was proposed, $\Phi^{b(p)}$, $b = 1, \dots, B$, via sampling from a symmetric distribution centered at the current state of the chain, $\Phi^{b(p)} \sim N(\Phi^{b(m-1)}, \Sigma^b)$, $b = 1, \dots, B$.

The proposed block of parameters was then accepted or rejected based on the ratio of the posterior of the proposed parameters and the posterior at the current state of the chain, $\min \left\{ 1, \frac{\pi(\Phi^{b(p)}, \Phi^{-b(m)}, \delta^{2(m-1)}, \Delta^{(m-1)}, R^{(m-1)} | y)}{\pi(\Phi^{(m-1)}, \delta^{2(m-1)}, \Delta^{(m-1)}, R^{(m-1)} | y)} \right\}$.

Here, $\Phi^{-b(m)}$ denotes the vector $\Phi^{(m)}$ with the b^{th} block removed. Accepted parameters were included in the Markov chain, and the previous values, $\Phi^{b(m-1)}$, were replicated otherwise.




Figure 3 15: Visualization of the Bayesian random effects model with discrepancy. Arrows represent the direction of conditional dependence. Clear nodes represent unobserved/unknown quantities while shaded nodes indicate observations.




Figure 3 16: Schematic of Metropolis-Hastings algorithm

The proposal covariance matrices were updated every 2000 iterations based on the acceptance rates over the past 2000 iterations and the target acceptance rates. The algorithm for the adaptation is included in the appendix. Parameters which were sampled in Gibbs updates were accepted with probability 1 since the full-conditional distributions allowed for direct sampling from the target posterior distributions.

Samples from the posterior predictive distribution were generated every 500 iterations, and the algorithm is included in the appendix. The output of the algorithm was a number of realizations, defined by the number of iterations M , of a Markov chain targeting the posterior distribution as well as a Markov chain targeting the posterior predictive density.


In comparison to the MCMC algorithm outlined here which was used in one of the case studies, the other studies use the same general algorithm with slight modifications. The most obvious difference is that the other two studies do not consider model form errors as a source of uncertainty. Therefore, the discrepancy, Δ , is not included as an unknown parameter in the statistical model in either of the studies. In one section, UQ and UP was performed on the VPSC model, and the random effects error precision matrix, Λ , was sampled instead of the decomposed correlation and variances. Another section made the assumption that the observation error was unique for each experiment. Therefore, a unique $\Psi^{[s]}$ was estimated for each of the random effects, $s = 1, \dots, S$. Besides these minor differences, the algorithms were identical to the one presented in the appendix.

4: RESULTS AND DISCUSSION


Uncertainty Quantification Accounting for Model Discrepancy within a Random Effects Bayesian Framework

Case Study I: Reduced-Order Crystal Plasticity

Monte Carlo samples targeting the posterior (Eq. 3.7) and the posterior predictive (Eq. 3.8) distributions were obtained via Algorithm 1. A total of 1.2×10^5 samples were produced, 20,000 of which were discarded as burn-in. Convergence was monitored visually by examining trace plots and correlation plots. The bivariate and marginal density estimates of a representative random effect are shown in the top panel of Fig. 4.1. These suggest uni-modal, approximately symmetric marginal posterior densities and moderate to very strong posterior correlations between parameters. Here, we remind the reader that posterior correlations do not necessarily imply direct physical relationships between parameters. In fact, posterior correlation is indicative only of a trade-off relationship between parameters, meaning that a change in the parameters in the direction of the correlation is associated with a model output that remains relatively stable. This results in a nearly constant likelihood given the observed data and leads to a stable posterior if the prior is diffuse.




(a) Random Effect




(b) Overall Effect

Figure 4 1: Correlation plots for a) a random effect and b) the overall effect. Marginal posterior (black) and prior (red) density estimates for each parameter are shown on the diagonal; bi-variate posterior (colors by intensity) and prior (red) contours



(a) Marginal Posterior Contours for (τ_1, ξ_1)



(b) Marginal Posterior Contours for (τ_0, ξ_0)

Figure 4 2: Posterior marginal contours for a) (τ_1, ξ_1) and b) (τ_0, ξ_0) . Solid contours are used for random effects; dashed contours are used for overall effects.

Table 4 1: Posterior summaries for the overall effect parameters and error precision.

Parameter	MAP	$\mathbb{E}_{\theta y}$	$\mathbb{V}_{\theta y}$	95% HPI
τ_0	53.66	53.93	2.43	(50.71, 56.96)
τ_1	56.43	55.75	4.74	(51.61, 60.32)
ξ_0	688.80	691.21	570.88	(642.76, 738.33)
ξ_1	151.99	155.22	93.49	(135.67, 174.35)
δ^2	23.78	22.99	.87	(21.18, 24.91)




Figure 4 3: 95% HPI (outer dashed lines) and posterior mean (central dashed line) for the posterior predictive distribution. 200 sampled predictive stress-strain curves (grey lines) are shown, and the MAP VPSC evaluation (solid red line).

It is helpful to recall that following a posterior contour corresponds to a path in the parameter space where the posterior density remains unchanged. For example, while the nearly perfect negative posterior correlation between $\tau[s]$ and $\xi[s]$ in Fig. 4.2a is not immediately a priori obvious, this result makes sense considering the physical meaning of these parameters. A larger asymptotic hardening rate is associated with a lower back-extrapolated intercept of the y-axis, so that trading off these values results in essentially the same model output.


The lower panel of Fig. 4.1 shows bivariate and marginal density estimates for the overall parameters. The marginal priors and bivariate prior contours, shown in red, are much more diffuse than the resulting posterior contours, indicating that the data are informative for all four overall model parameters. Posterior correlations among the overall parameters are mostly mild. Their interpretation as defining the distribution of the random effects is illustrated in Figs. 4.2a, 4.2b, and A.1. Note that the posterior over the overall effects is quite diffuse relative to that of the individual random effects. Indeed, the rich calibration data provide direct information only

about each individual random effect, while the overall effects are informed only by the relatively small number ($S = 6$) and location of the random effects themselves. An example is the opposite direction of the correlation between the random effects $\tau[s]$ and $\xi[s]$ versus that of their respective random effects (Fig. 4.2b) which proceeds from the location of the random effects. That there are different parameter correlation effects for the individual random effects versus the overall ensemble behavior is motivated by their different roles within the hierarchical model. The marginal posterior and prior for the error precision δ^2 are shown in Fig. A.3 in the appendix. Again, the contrast of the diffuse non-informative prior and highly peaked posterior marginals is indicative of informative data (Tables 4.1, 4.2).


The output from Algorithm 1 can also be used to predict future observations by generating samples from the posterior predictive distribution with density Eq. (3.8). This is the posterior distribution of new (unobserved) data given the calibration data we have observed. Figure 4.3 shows 200 thinned samples from the posterior predictive distribution over the observed stress-strain curves. The distribution is further summarized by displaying 95% highest posterior density intervals (HPI), computed in a point-wise fashion at each different value of strain. The point-wise posterior mean as well as the maximum a posteriori (MAP) reflects the center of the predictive distribution and its mode, respectively. We first notice that the predictive draws do not capture the behavior of the data in the elastic region. This is expected, as the VPSC model does not consider the elastic region and was calibrated only with the plastic portion of the training data. The most noticeable feature of the point-wise predictive uncertainty (the spread of the posterior predictive distribution) is that it reflects not only the measurement noise for a given sample, but also the spread in the calibration sample trajectories themselves, as illustrated in Fig. 4.1. Following the pattern of the sample-to-sample stress-strain trajectories, the spread of the posterior uncertainty increases with the strain (from left to right). The spread, as well as the presence of several outliers, reflect the contribution in predictive uncertainty from uncertainty in the overall effects, which are directly informed by only a small number ($S = 6$) of random effects. This provides the following insight: If one wishes to control the predictive uncertainty for this model by designing an experiment, the highest reduction in prediction uncertainty will be achieved by testing more samples (i.e., measuring more stress-strain trajectories) rather than trying to further decrease the measurement noise for each single trajectory. This insight follows from a well-known principle of Bayesian inference: The richer the information on which unknown quantities are conditioned, the smaller the posterior uncertainty about their values. The random effect distributions were updated conditionally on data containing approximately 200 evenly spaced observations, leading to relatively low marginal posterior uncertainty as shown in the contour plots of Fig. 4.2 as well as Fig. A.1. In contrast, the marginal posterior histograms for the overall effects, which were updated based on only 6 observations (each of the random effects), yield comparatively much greater uncertainty. Since the posterior predictive distribution accounts for the uncertainty in the overall effects, decreasing the posterior spread of the overall parameter distribution by conditioning on more samples is expected to have the greatest impact on predictive uncertainty. This would be expected even if the observations within the samples were sparse or noisy since i) sample-to-sample variability is in general expected to be greater than observation error, and ii) the overall effects are directly dependent only on the samples themselves and not each observation within the samples.

Case Study II: Thermodynamic Calculations

Monte Carlo samples targeting the posterior distribution were obtained with Algorithm 1. In total, 2×10^8 samples were obtained, 3×10^7 were discarded as burn-in, and the remaining samples were thinned by a factor of 400. Evidence of chain convergence was assessed through trace plots and correlation plots. The marginal and bivariate marginal estimates for a random effect for the model at 1000 K are shown in the top panel of Fig. 4.3. These densities are approximately symmetric and uni-modal with no apparent posterior correlation between interaction parameters. The lower panel in Fig. 4.3b shows the marginal and bivariate density estimates for each overall effect parameter along with the prior density contours plotted in red. In all cases, the prior distribution is very diffuse when compared to that of the posteriors. As in the first case study, this is indicative of informative training data. Figure 4.4 shows the relative size and location of the density contours for the overall effect (dashed line) and the $S = 3$ random effects (solid lines, note also that the contours for random effects $s = 1$ and $s = 2$ overlap). The strong correlation between the overall effect interaction parameters is a result of the distribution of the random effects which is further summarized in Fig. A.2. The element-wise correlation for all parameter pairs in one of the random effects as well as the overall effect is reported in Tables A.3 and A.4. Even though the data informing each random effect are plentiful, there are only $S = 3$ data points informing the overall effect, leading to posterior densities with much higher spread. Forward model evaluations for approximately 500 draws from the posterior distribution of the overall parameters are shown in Fig. 4.5. The scale of the Gibbs free energy makes the posterior appear deceptively narrow, and magnification in Fig. 4.5b shows a much wider distribution than is evident in the top panel. The outer and center dashed red lines indicate the 95% HPI and the posterior mean, respectively. The expected value and variance of the overall parameters, the MAP parameter values, and the 95% HPI were estimated and are provided in Table 4.2. As expected, uncertainty increases moving away from each of the boundaries.



(a) Random Effect



(b) Overall Effect

Figure 4 4: Correlation plots for a) random effect and b) the overall effect. The diagonal subplots contain the marginal histogram for each parameter with the overall parameter histograms being plotted with each respective prior density (red line).

Table 4 2: Posterior summaries for the overall effect parameters and error precision at 1000K.

Parameter	MAP	$\mathbb{E}_{\theta y}$	$\mathbb{V}_{\theta y}$	95% HPI
L_0	25617	25601	903.83	(25538, 25658)
L_1	-4386	-4425	6761.2	(-4594, -4265)
$\delta^2 \times 10^{-3}$	2.1	2.0	1.3×10^{-5}	(1.8, 2.3)






Figure 4 5: Contour Plot of estimated bi-varaite distribution between L_0 and L_1 . The random effect contours are shown with solid lines and the overall effect countours are shown with dashed lines. Note that the contours of the random effect in the bottom left corner overlap.



(a) Posterior



(b) 0 – .1 atomic fraction Copper

Figure 4 6: Forward model evaluations based on 530 samples from the posterior distribution for the Gibbs free energy as a function of atomic fraction copper at 1000K: a) the entire range of 0-1 atomic fraction copper b) a magnified range of 0-.1 atomic fraction copper is shown. This distribution can be used to propagate parameter uncertainty to phase diagrams.

Uncertainty Quantification Accounting for Model Discrepancy

The MCMC algorithm detailed in [175] was used for this work with appropriate adjustments (i.e., the decomposition of the random effects precision Λ as well as the inclusion of the model discrepancy Δ) and can be found in Appendix B.1. A total of 1.25×10^5 Markov chain Monte Carlo samples were simulated targeting the desired posterior (4.12) and posterior predictive (4.14) distributions with the first 2.5×10^3 samples being discarded as burn-in. Trace plots as well as correlation plots were used to monitor convergence.




Figure 4.7: Random draws from the prior of the discrepancy in a) the elastic region and b) the plastic region c) over the full domain. Note the difference in the scale of the vertical axes.

Marginal posterior histograms as well as bi-variate kernel density estimates of marginal posterior contours are shown in Fig. 4.7 for a representative random effect and the overall effect. In Fig. 4.7b the red lines show the marginal as well as bi-variate priors placed on the overall effect parameters for comparison. Notably, there is an almost perfect positive posterior correlation between random effect parameters τ_1 and ξ_1 , also provided in Table B.1 in Appendix B.3. Likewise, there is moderate posterior correlation in the overall version of these parameters. A bi-variate kernel density estimate showing the (τ_1, ξ_1) parameter correlation for both the random effects and overall effect is displayed in Fig. 4.8a, with the overall effect being denoted by the dashed line. Figure 4.8b shows another bi-variate kernel density estimate overlay for parameters (τ_0, ξ_0) which shows a moderate correlation in both the random effects as well as the overall effect. Here, the reader is reminded that correlations between parameters are features of the posterior distribution and not directly of the physical system. Parameters combinations on a given contour of the posterior distribution simply have the same posterior probability; a more thorough discussion can be found in [175].





Figure 4 8: Marginal posterior histograms and parameter correlations for a) a representative random effect and b) the overall effect,with prior draws included in red for the overall effect.



(a) Marginal posterior contours for (τ_1, ξ_1) (b) Marginal posterior contours for (τ_0, ξ_0)

Figure 4.9: Marginal posterior contours for a) (τ_1, ξ_1) and b) (τ_0, ξ_0) . Random effect marginal contours are indicated by solid lines and the overall effect contours are indicated by a dashed line.

Posterior expectations for the model parameters are provided in Table 4.3, along with the associated uncertainty about these values, summarized through posterior variances and %95 highest posterior density intervals (HPIs). Parameter uncertainty was propagated by sampling from the posterior over the parameters to generate posterior and posterior predictive distributions; 200 representative samples are shown in Figs. 4.9b and 4.9c, respectively. In both plots, highest 95% posterior intervals of the stress-strain response are represented with red dashed lines and the mean stress-strain response is shown with a red line.

Table 4.3: Posterior summaries for the overall effect parameters and error precision.

Parameter	$\mathbb{E}_{\theta y}$	$\mathbb{V}_{\theta y}$	95% HPI
τ_0	55.74	0.90	(53.80, 57.60)
τ_1	65.35	16.27	(59.05, 75.18)
ξ_0	711.97	498.76	(670.30, 759.63)
ξ_1	139.33	297.76	(102.51, 171.54)
δ^2	12.71	0.30	(11.60, 13.79)


Here, we take a moment to remind the reader of the subtle difference in the interpretation of these two distributions. While the posterior distribution conveys our subjective belief about the true underlying property, the posterior predictive distribution is our belief about how future

experiments will behave, thus resulting in a property distribution which is more diffuse. Figure 4.9a shows the mean posterior draw as well as its decomposition into the VPSC model evaluated at the expected value of the parameters and the mean discrepancy. Also included is an inset of the elastic region.


An informative comparison can be made between the analyses operated under the assumption that discrepancy is present (Fig. 4.9b), and the assumption that it is not (Fig. 4.5). The strong bias evident in the ‘no discrepancy’ analysis is not present when model misfit is accounted for.

When to Include Discrepancy


While the incorporation of model discrepancy was appropriate and necessary for the study presented here, its inclusion is highly application dependent. In cases such as this where 1) The model is missing physics, resulting in systematic discrepancy between the model output and the observations and 2) There is strong prior knowledge about the misfit between the model and true process (details on this in the following paragraph), then it is appropriate to incorporate discrepancy for inference and prediction.



(a) Mean Summaries



(b) Posterior HPI



(c) Posterior predictive HPI

Figure 4 10: The mean posterior draw (blue), the VPSC model evaluated at the mean parameter values (red) and the mean discrepancy (green) are shown in a) with an inset showing the elastic region. 200 draws from the a) posterior and b) posterior predictive distributions are shown with the mean and 95% highest posterior intervals shown.

On the other hand, with acknowledgement that no model is perfect, it may be reasonable to exclude discrepancy from the analysis if the model is flexible enough to fit the data and the model is empirical in nature, where parameters do not have physical significance. The condition on the empirical nature is important since even if a model is flexible enough to fit the data (as in the toy example), the parameter estimates will be strongly biased with a strong likelihood the HPs will not cover the true parameter values.

Identifiability & Sensitivity of Modeling Discrepancy

The sensitivity of modeling discrepancy stems from the property of indentifiability. This describes a desirable property of parameters in statistical models by which their true values can be inferred from the data. Parameters that are non-identifiable can be considered as sharing a marginal posterior distribution; they are not independent of each other, and so even with an infinite amount of calibration data, the true values of the individual parameters cannot be determined [51]. In a paper by Brynjarsdóttir [51], this concept is succinctly expressed by rewriting the model for the observations with discrepancy as,

$$\Delta(x) = \zeta(x) - m(x, \theta). \quad (4.1)$$

Note, even if the true process ζ is known perfectly, for every value of θ , there is a corresponding Δ that satisfies the condition in (4.1). Therefore, θ and Δ are not identifiable from the observations.

As discussed in [51], the key to reliable learning about the parameters is to incorporate as much prior information about the discrepancy (and θ) as possible. The more realistic the prior information is (i.e., assigning higher prior probability to the true θ and Δ), the higher the likelihood the credible intervals of the posterior distribution will cover the true parameters [51]. In doing so, the posterior distribution over θ and Δ becomes tighter as a result of stronger prior information, allowing for more reliable inference on the unknown parameters.

Therefore, if there is no understanding of the missing physics, incorporating discrepancy into the statistical model will not allow reliable learning about the parameters (and likely lead to unreliable extrapolation).

Discrepancy in Applications for ICME

A question which arises from the importance of prior modeling is perhaps this: If reliable learning of the parameters requires such careful modeling of the discrepancy prior, why not just improve the constitutive model to include the physics and thus eliminate the need to account for the discrepancy? Often accounting for model misfit is both computationally and philosophically easier than improving a model. As an example, take the case study presented in this paper. Inference was performed on the VPSC model while accounting for discrepancy since the model does not include physics for the elastic regime of deformation. As demonstrated, the misfit due to a visco-plastic model can be handled in a statistical sense with some basic prior knowledge related to the applicable domain and effective magnitude of the elastic effect. Incorporation of an elastic component to the model, on the other hand, would require a single-crystal elastic constitutive law, numerical values of the components of the 4th order elasticity tensor, a schema for partitioning the total deformation into the elastic and plastic contributions (typically $\mathbf{F} = \mathbf{F}^e \cdot \mathbf{F}^p$), and significant updates to the numerical framework.

Even supposing a more complete model is available, it may still be beneficial to work with a simpler but discrepant model. As an alternative to VPSC, inference could be performed on its more complete counterpart, the elasto-viscoplastic self-consistent model [179] or alternatively full-field elasto-viscoplastic fast Fourier transform [11, 145, 15] or finite element simulations [19, 14], which do account for elastic deformation. However, the computational cost of these models is such that a technique such as MCMC could not be employed for direct inference on the model. Still, in many cases, a more complete model is not available for the application due to truly unknown effects, in which case the discrepant model offers the best representation of the quantity of interest. As a result, if a model is very good at representing a material process or behavior on a restricted domain, but not the entire domain (resulting in systematic discrepancy), the consideration of model-form error may still allow reliable inference and prediction on the unknown modeling components.

In a broad sense, while 3D, physics-based and full-field models may have a higher fidelity as a result of their complex formulations, their high parameter dimensionality makes them very expensive and difficult to calibrate. In effect, MCMC simulations are not a feasible option for this class of models where a great number of evaluations are necessary in order to achieve statistical convergence on the target posterior. Of course, these high-fidelity models can make use of MCMC simulations in tandem with the adoption of a surrogate model or other techniques such as multi-fidelity optimization [36] to recover information on unknown modeling components. However, the framework we present in this work is with the purpose of its application for inferential analysis on models which are favorable for MCMC simulations (although, it can easily be extended to include surrogate models). While a comparison of the techniques of emulation vs. discrepancy modeling is outside the scope of this paper, it is worthwhile to note that emulation and surrogate inference comes with its own set of challenges which may yet lead to difficulty in learning the true model parameters.

‘Lower’-Fidelity Models for Design

The inclusion of model misfit into the inferential and predictive problem has important implications for reduced-order, homogenized, and phenomenological models. Accounting for model discrepancy allows uncertainty propagation for reliable posterior predictions even when the model is incomplete or our physics isn’t perfect. This has great value for design and engineering problems since it opens the door for the reliable use of these predictions in design and optimization applications and is a main focus of ongoing work from the authors.

Inference versus Prediction

We propose a modeling framework for connecting physical models and data while accounting for different types of variability in materials and experimental conditions. The Bayesian inferential approach propagates this uncertainty to the prediction of various material behaviors. “Methods” section distinguishes the related goals of inference and prediction, which are based on the posterior and posterior predictive distributions, respectively. The distribution induced in the forward model by propagating posterior uncertainty in the parameters from Eq. (3.7) is a representation of the true underlying average material response given the observed data. Prediction of this material response must also consider sample variability and experimental error associated with making a new measurement, which means that the posterior predictive response

is always more uncertain than the posterior response. Though the notions of inference and prediction are fundamentally different, both have their place in the propagation of uncertainty in materials engineering, as demonstrated in the two case studies.

Prediction for the VPSC Model

The first case study was concerned with inference for the Voce hardening parameters defining the CP VPSC model, which simulates the stress- strain response of a material, and prediction of the response. The propagation of the uncertainty in parameter values and stress-strain response to predict the behavior of new samples not yet tested was based on the posterior predictive distribution.

Since the Voce hardening parameters are a material property independent of strain rate, deformation conditions and texture, the calibrated parameters can be used to make predictions under different conditions than the ones they were calibrated with. The Voce hardening parameters were trained with data for simple compression up to 9% strain and an initial uniform random texture. Figure 4.10a shows the predictive distribution of copper under ideal rolling boundary conditions up to 50% strain and a uni-modal “copper” texture (Bunge Euler angles $\phi_1 = 90^\circ$, $\Phi = 35^\circ$, $\phi_2 = 45^\circ$). The distinction between inference on the model output and its prediction are illustrated in Fig. 4.10a and 4.10b, showing 250 response trajectories drawn from the posterior and posterior predictive distributions, respectively. Because the posterior predictive distribution accounts for the expected sample variability and observation errors, it is ideally suited for design purposes to establish the bounds of expected observed material response.

Thermodynamic Posterior Uncertainty

The goal of the second case study was to propagate the posterior uncertainty in the Redlich-Kister interaction parameters in the excess Gibbs free energy to establish confidence bounds in the phase boundaries. Unlike the first case study, we are not concerned with how new test samples are going to behave, but rather are interested in the true underlying system behavior to determine the bounds of phase stability. Figure 4.11 shows the distribution of the atomic fraction of copper at the boundaries with the center dashed line denoting the mean value and the outer dashed lines bounding the 95% HPI. Values obtained from experimental work and databases are plotted along with our predictions for comparison. Five hundred samples were drawn from the posterior distribution, and minimization of the potential difference of the constituents in each phase was used to determine the equilibrium atomic fraction of copper of the phase boundaries (α , $\alpha + \beta$) and ($\alpha + \beta$, β) at 1000 K. Posterior expected values, variances, and 95% HPI intervals are reported in Table 4.4. A full study of uncertainty in phase diagrams should be informed by multiple data sets at a significant number of temperatures. For the purpose of this study, which was to simply demonstrate the generality of this statistical framework, the sparse data at a single temperature were sufficient.




Figure 4 11: Stress-strain response predictions (a) and posterior draws (b) up to 50% strain under ideal rolling boundary conditions and a 'copper' initial texture. Samples of size 200 are shown.

Table 4 4: Expected values, variances, and 95% highest posterior density intervals of the phase boundaries in terms of atomic fraction of Cu in the Binary AgCu system.

Phase Boundary	$E_{x_{Cu} y}$	$V_{x_{Cu} y}$	95% HPI
$\alpha, \alpha + \beta$	.106	1.52×10^{-6}	(.104, .109)
$\alpha + \beta, \beta$	.967	7.64×10^{-8}	(.966, .967)

Interpreting the Output of a Bayesian Model

Posterior uncertainty in the parameters takes the form of a distribution, so it is important to interpret its implications correctly. This posterior uncertainty often manifests in strong correlations between certain model parameters. The correct interpretation of these correlations is as the direction in which trading off these parameter results in very little change to the posterior density. More formally, a posterior contour simply shows a path in the parameter space which yields model evaluations that are equally likely given the calibration after accounting for prior uncertainty. A benefit of using sampling-based methods such as MCMC is that they fully explore this distribution, rather than just yielding a point estimate. This is key in describing the uncertainty (both its spread and direction) about the quantity of interest since there are often many parameter combinations that result in approximately the same model output.




Figure 4 12: Histograms showing the propagated uncertainty in atomic fraction Cu (x_{Cu}) for phase boundaries a) $\alpha, \alpha + \beta$ and b) $\alpha + \beta, \beta$. The outer and middle dashed lines (red) bound the 95% HPI and mean posterior value, respectively. As comparison, boundaries reported by other sources [171–173, 169, 174] are plotted as well.

Informative Data

In our hierarchical model, the posterior on the overall effect parameters is somewhat sensitive to the specification of a strongly-informative prior choice. This is because the overall effect is conditioned on a sparse collection of samples ($S = 6, 3$ for our case studies), and the effect of the prior model increases in such "data-poor" scenarios. We therefore recommended assigning diffuse priors in cases where there is little known about the parameter values a priori. This was exemplified in our case studies; the priors on the overall effect parameters as well as the error precision terms were chosen to be relatively flat and thus only weakly informative. Despite the small number of samples observed, the posterior densities in both case studies were noticeably less diffuse than the corresponding priors, implying that the data, though sparse, were very informative for the unknown quantities.

Computationally Expensive Models

The models used in both case studies were chosen in part for their computational convenience. Because of this, the MCMC algorithm was used without needing to further reduce or simplify the models. 1 million samples were obtained in the thermodynamic case study every 30 min, and 1000 samples were obtained for the crystal plasticity study approximately every 2 h. The MCMC algorithm would quickly become very burdensome, however, for more computationally intensive models. In order to accommodate such models, a surrogate or other model reduction (such a machine learning) could be built directly into the hierarchical model formulation we presented. Inference and prediction would then proceed in a similar fashion, though more uncertainty in the posterior would reflect the presence of this additional model uncertainty.

Insights for Experimental Design

It is often the case that only a small number of material samples are measured due to the high costs of experiments. While the number of data points and accuracy of the measurements corresponding to each sample are typically high, the small number of samples tested means that there is still substantial uncertainty in the overall parameters, which gives rise to most of the posterior predictive uncertainty we obtain. Therefore, for both case studies, increasing the number of data points per test and minimizing the noise may have less of an impact on improving prediction than increasing the number of observed sample trajectories.

Description of Random Effects

The two case studies presented here are scenarios where there are often few specimens tested, limiting the number of random effects informing the overall effect, and providing little information on their distribution. In cases such as this, where there is little known about a distribution a priori, using a distribution with high entropy as the probability model is desirable in order to reduce the information contained within the distribution. In both case studies, this makes the selection of a multivariate normal distribution (which has the maximum entropy among real-valued distributions) appropriate to describe the sample-to-sample variability of the random effects. However, in cases where there are more data available and the random effects are not well described by a normal distribution, the hierarchical framework presented here can be extended to any distribution which best describes our understanding of the relationship between the random and overall effects. It is useful to note that such a change in the likelihood model must also be accounted for in the derivation of any full-conditional distributions.

Modeling the Observation Error

A general description of the observation error was presented in terms of a full covariance matrix, Ψ^{-1} , in Eq. 3.25. However, in both applications presented, the observation error was assumed to be independent and identically distributed. While this simplifying assumption was convenient and appropriate for the work here, this may not always be the case. In cases where the observation error distribution varies along with the domain, or where errors are spatially correlated, the full error covariance matrix should be inferred. As an example, the derivation of the full-conditional distribution for Ψ is provided in Eq. (A.2) in the appendix. This derivation places a Wishart prior on the variance-covariance matrix.

Bayesian Computer Model Emulation Using a Gaussian Process

Figure 4.12 shows the observed data and calibrated model. The black dots are the generated experimental data including the error term. The blue curve is the VPSC model evaluated at the MCMC sample where the posterior likelihood is maximized, which is called the mode posterior. The red curve is the basis spline curve from the mode posterior of basis coefficient sample. The method well calibrates the input-output relationship of the materials model leading to almost identical curved on the Figure 4.12.

MCMC samples of calibration parameters are summarized on the Figure 4.13 and Table 4.5. The blue curves on the diagonal plots are the estimated posterior distribution from the MCMC output. The red and black vertical lines are the true parameter value and the mode posterior parameter respectively. The contour plots indicate the correlation between the calibration parameters. There seems to be positive but weak correlation between τ_0 and ξ_1 , and τ_0 and ξ_0 .

The Figure 4.14 illustrates the point-wise 95% credible bound of basis spline using a small number of training data. The upper plots shows where the basis coefficients are located and the bottom plot is the 95% point-wise credible bound of basis coefficients. The true process lies between the 95% point-wise credible bound for every stain points and the bound size is very small (there is a box which enlarged a part of the plot).

Table 4 5: The summary table of posterior distributions.

Parameters	True Value	Posterior Mode	Posterior Mean	Posterior Variance
τ_0	55	54.907	54.985	0.146
ξ_0	55	55.088	55.050	0.285
τ_1	700	699.485	699.471	2.398
ξ_1	200	200.1863	200.029	1.464




Figure 4 13: The posterior histogram and correlation contour plot of the posterior distributions.




Figure 4 14: The upper plots show the basis coefficients of the training data (dots) and of the true process (blue star). The bottom plot is the true projected spline (red) and 95% point-wise credible bound of basis splines (blue).

5: CONCLUSIONS

Uncertainty Quantification for Parameter Estimation and Response Prediction

A UQ framework was developed using the principles of Bayesian inference and demonstrated on two case studies with different inferential goals. In the VPSC study, exploration of the posterior predictive distribution established the uncertainty in the mechanical behavior of samples not yet tested. In the thermodynamic study, the posterior distribution was used to propagate the uncertainty in the Gibbs free energy as a function of atomic fraction copper in the binary Ag-Cu system to the uncertainty in phase boundaries at 1000 K. Both case studies demonstrate the benefit and reasonable assumptions made in implementing a random effects model for uncertainty quantification in forward materials models, which provides a method for various sources of uncertainty to naturally be accounted for.

This work suggests and demonstrates a UQ tool to establish confidence bounds in material performance predictions. It has been illustrated that even with a limited number of test samples and a deterministic model simulating material behavior, a distribution can be created which contains the range and likelihood of material responses, not excluding outliers. This result has great potential for the promotion of ICME, the reduction of experimental costs, and the propagation of uncertainty between models to reliably extract important material properties for design work.

Uncertainty Quantification Accounting for Model Discrepancy

There is a great significance to the goals of ICME in accounting for model misfit in uncertainty analyses. Two primary goals of ICME are 1) The linking of models across vast scales in order to account for mechanisms and phenomena affecting the material behavior and performance on many different levels, and 2) A reduction in the volume of experiments needed for design. This relies on being able to first quantify and then propagate uncertainty in and between models in the modeling chain in order to obtain a reliable processing $\leftrightarrow$ performance predictions and to establish a design confidence. However, many models which are used for design purposes are computationally intensive, with run times exceeding resources typically available to perform advanced UQ techniques without emulation. Being able to incorporate model misfit in to the UQ analysis opens the door to the use of reduced- order, phenomenological, or homogenized models in reliable design work.

In this work, a UQ analysis was performed on the phenomenological VPSC model under the framework presented in [175], with appropriate adjustments to incorporate model misfit. The VPSC model, which accounts only for the plastic regime of deformation, was calibrated with data from *both* the elastic and plastic region. Uncertainties about unknown model parameters were established and propagated through the model for inference and prediction on the stress-strain response. Posterior summaries such as the MAP and HPI parameters and model evaluations were used to summarize the analysis.

Bayesian Computer Model Emulation Using a Gaussian Process

This paper demonstrates the Bayesian calibration method for expensive or complex materials models. The presented Bayesian model calibrated a response surface through a finite number of input-output pairs of a physical model without running materials simulations during the MCMC step. This work accounted for function-valued output by projecting to a basis spline function space for a low dimensional presentation. To model the response surface in the Bayesian framework, the Gaussian process prior was used. The Bayesian hierarchical model simultaneously captured the various sources of uncertainty, such as experimental error variability, parameter uncertainty and code uncertainty. Thus, the resulting statistical inference accounts for such uncertainties. This study can be extended by a more sophisticated model, integrating the previous works and establishing Bayesian emulation method using a random effects model and a discrepancy component

- The likelihood component of our hierarchical model can be straightforwardly replaced with a more sophisticated model.
- A BHM allows us to simultaneously model all sources of uncertainty, which is then automatically accounted for inference and prediction.

6.0 REFERENCES

- [1] CN Tome´ and RA Lebensohn. Manual for code visco-plastic self-consistent (vpssc). *Los Alamos National Laboratory, New Mexico, USA*, 2009.
- [2] Aleksandr Chernatynskiy, Simon R Phillpot, and Richard LeSar. Uncertainty quantification in multiscale simulation of materials: A prospective. *Annual Review of Materials Research*, 43:157–182, 2013.
- [3] Jitesh H Panchal, Surya R Kalidindi, and David L McDowell. Key computational modeling issues in integrated computational materials engineering. *Computer-Aided Design*, 45(1):4–25, 2013.
- [4] Raymundo Arróyave and David L McDowell. Systems approaches to materials design: Past, present, and future. *Annual Review of Materials Research*, 49, 2019.
- [5] Mark F Horstemeyer. *Integrated Computational Materials Engineering (ICME) for metals: using multiscale modeling to invigorate engineering design with science*. John Wiley & Sons, 2012.
- [6] John Allison. Integrated computational materials engineering: A perspective on progress and future steps. *Jom*, 63(4):15, 2011.
- [7] Thomas Tomkins Warner. *Numerical weather and climate prediction*. Cambridge University Press, 2010.
- [8] Donald W Boyd. *Systems analysis and modeling: a macro-to-micro approach with multidisciplinary applications*. Elsevier, 2000.
- [9] Jacob Fish. Bridging the scales in nano engineering and science. *Journal of Nanoparticle Research*, 8(5):577–594, 2006.
- [10] Bradford Cowles, Dan Backman, and Rollie Dutton. Verification and validation of icme methods and models for aerospace applications. *Integrating Materials and Manufacturing Innovation*, 1(1):2, 2012.
- [11] Ricardo A Lebensohn, Anand K Kanjarla, and Philip Eisenlohr. An elasto-viscoplastic formulation based on fast fourier transforms for the prediction of micromechanical fields in polycrystalline materials. *International Journal of Plasticity*, 32:59–69, 2012.
- [12] Geoffrey Ingram Taylor. Plastic strain in metals. *J. Inst. Metals*, 62:307–324, 1938.
- [13] Sachs G. Zur ableitung einer fliessbedingung. *Z VDI*, 72:734–736, 1928.
- [14] Franz Roters, Philip Eisenlohr, Luc Hantcherli, Denny Dharmawan Tjahjanto, Thomas R Bieler, and Dierk Raabe. Overview of constitutive laws, kinematics, homogenization and multiscale methods in crystal plasticity finite-element modeling: Theory, experiments, applications. *Acta Materialia*, 58(4):1152–1211, 2010.

- [15] Pengyang Zhao, Thaddeus Song En Low, Yunzhi Wang, and Stephen R Niezgoda. An integrated full-field model of concurrent plastic deformation and microstructure evolution: Application to 3d simulation of dynamic recrystallization in polycrystalline copper. *International Journal of Plasticity*, 80:38–55, 2016.
- [16] A Ma, F Roters, and D Raabe. A dislocation density based constitutive model for crystal plasticity fem including geometrically necessary dislocations. *Acta Materialia*, 54(8):2169–2179, 2006.
- [17] Douglas J Bammann. An internal variable model of viscoplasticity. *International Journal of Engineering Science*, 22(8-10):1041–1053, 1984.
- [18] LP Evers, WAM Brekelmans, and MGD Geers. Non-local crystal plasticity model with intrinsic ssd and gnd effects. *Journal of the Mechanics and Physics of Solids*, 52(10):2379–2401, 2004.
- [19] F Roters, P Eisenlohr, C Kords, DD Tjahjanto, M Diehl, and D Raabe. Damask: the düsseldorf advanced material simulation kit for studying crystal plasticity using an fe based or a spectral numerical solver. *Procedia Iutam*, 3:3–10, 2012.
- [20] Athanasios Arsenlis and David M Parks. Modeling the evolution of crystallographic dislocation density in crystal plasticity. *Journal of the Mechanics and Physics of Solids*, 50(9):1979–2009, 2002.
- [21] Ke-Shen Cheong and Esteban P Busso. Discrete dislocation density modelling of single phase fcc polycrystal aggregates. *Acta materialia*, 52(19):5665–5675, 2004.
- [22] Pejman Honarmandi, Noah Paulson, Raymundo Arroyave, and Marius Stan. Uncertainty quantification and propagation in calphad modelling. *Modelling and Simulation in Materials Science and Engineering*, 2019.
- [23] Noah H Paulson, Brandon J Bocklund, Richard A Otis, Zi-Kui Liu, and Marius Stan. Quantified uncertainty in thermodynamic modeling for materials design. *Acta Materialia*, 174:9–15, 2019.
- [24] Walter Olbricht, Niranjana D Chatterjee, and Klaus Miller. Bayes estimation: A novel approach to derivation of internally consistent thermodynamic data for minerals, their uncertainties, and correlations. part i: Theory. *Physics and Chemistry of Minerals*, 21(1-2):36–49, 1994.
- [25] Erich Königsberger. Improvement of excess parameters from thermodynamic and phase diagram data by a sequential bayes algorithm. *Calphad*, 15(1):69–78, 1991.
- [26] Dmitri V Malakhov. Confidence intervals of calculated phase boundaries. *Calphad*, 21(3):391–400, 1997.
- [27] Richard A Otis and Zi-Kui Liu. High-throughput thermodynamic modeling

- and uncertainty quantification for icme. *JOM*, 69(5):886–892, 2017.
- [28] M Stan and BJ Reardon. A bayesian approach to evaluating the uncertainty of thermodynamic data and phase diagrams. *Calphad*, 27(3):319–323, 2003.
 - [29] Babak Kouchmeshky and Nicholas Zabaras. The effect of multiple sources of uncertainty on the convex hull of material properties of polycrystals. *Computational Materials Science*, 47(2):342–352, 2009.
 - [30] Saikumar R Yeratapally, Michael G Glavicic, Christos Argyrakis, and Michael D Sangid. Bayesian uncertainty quantification and propagation for validation of a microstructure sensitive model for prediction of fatigue crack initiation. *Reliability Engineering & System Safety*, 164:110–123, 2017.
 - [31] Babak Kouchmeshky and Nicholas Zabaras. Microstructure model reduction and uncertainty quantification in multiscale deformation processes. *Computational Materials Science*, 48(2):213–227, 2010.
 - [32] Marisol Koslowski and Alejandro Strachan. Uncertainty propagation in a multiscale model of nanocrystalline plasticity. *Reliability Engineering & System Safety*, 96(9):1161–1170, 2011.
 - [33] Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I Jordan. An introduction to mcmc for machine learning. *Machine learning*, 50(1-2):5–43, 2003.
 - [34] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.
 - [35] Marc C Kennedy and Anthony O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
 - [36] Alexander IJ Forrester, András Sóbester, and Andy J Keane. Multi-fidelity optimization via surrogate modelling. *Proceedings of the royal society a: mathematical, physical and engineering sciences*, 463(2088):3251–3269, 2007.
 - [37] Shawn E Gano, John E Renaud, Jay D Martin, and Timothy W Simpson. Update strategies for kriging models used in variable fidelity optimization. *Structural and Multidisciplinary Optimization*, 32(4):287–298, 2006.
 - [38] Jennifer A Hoeting, David Madigan, Adrian E Raftery, and Chris T Volinsky. Bayesian model averaging: a tutorial. *Statistical science*, pages 382–401, 1999.
 - [39] Jacob M Montgomery and Brendan Nyhan. Bayesian model averaging: Theoretical developments and practical applications. *Political Analysis*,

- 18(2):245–270, 2010.
- [40] Larry Wasserman. Bayesian model selection and model averaging. *Journal of mathematical psychology*, 44(1):92–107, 2000.
 - [41] RC Rice, JL Jackson, J Bakuckas, and S Thompson. Metallic materials properties development and standardization. *Metallic Materials Properties Development and Standardization, Report No. MMPDS-01 DOT/FAA*, 2003.
 - [42] Jana Jackson. Definition of design allowables for aerospace metallic materials. *AeroMat Presentation*, 2007.
 - [43] Tim Bedford, Roger Cooke, et al. *Probabilistic risk analysis: foundations and methods*. Cambridge University Press, 2001.
 - [44] Carlos Conceição António and Luísa N Hoffbauer. An approach for reliability- based robust design optimization of angle-ply composites. *Composite Structures*, 90(1):53–59, 2009.
 - [45] A Saltelli, K Chan, and EM Scott. Sensitivity analysis., volume new york. *J. Wiley*, 2000.
 - [46] CF Wong and H Rabitz. Sensitivity analysis and principal component analysis in free energy calculations. *The Journal of Physical Chemistry*, 95(24):9628–9630, 1991.
 - [47] Heeralal Gargama, Sanjay Kumar Chaturvedi, and Awalendra Kumar Thakur. Reliability-based design optimization of broadband microwave absorbers. *Journal of Electromagnetic Waves and Applications*, 27(11):1407–1418, 2013.
 - [48] C Zang, MI Friswell, and JE Mottershead. A review of robust optimal design and its application in dynamics. *Computers & structures*, 83(4-5):315–326, 2005.
 - [49] Yan Wang. Uncertainty in materials modeling, simulation, and development for icme. *Proc. 2015 Materials Science and Technology*, 2015.
 - [50] National Research Council et al. *Integrated computational materials engineering: a transformational discipline for improved competitiveness and national security*. National Academies Press, 2008.
 - [51] Jenný Brynjarsdóttir and Anthony O'Hagan. Learning about physical parameters: The importance of model discrepancy. *Inverse problems*, 30(11):114007, 2014.
 - [52] Sergey Oladyshekin and Wolfgang Nowak. Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion. *Reliability Engineering & System Safety*, 106:179–190, 2012.
 - [53] Richard Dwight and Zhong-Hua Han. Efficient uncertainty quantification using gradient-enhanced kriging. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 17th*

- AIAA/ASME/AHS Adaptive Structures Conference 11th AIAA No*, page 2276, 2009.
- [54] Xiaoping Du. Unified uncertainty analysis by the first order reliability method. *Journal of mechanical design*, 130(9):091401, 2008.
 - [55] Junfu Zhang and Xiaoping Du. A second-order reliability method with first-order efficiency. *Journal of Mechanical Design*, 132(10):101006, 2010.
 - [56] Nicholas Metropolis and Stanislaw Ulam. The monte carlo method. *Journal of the American statistical association*, 44(247):335–341, 1949.
 - [57] Peter W Glynn and Donald L Iglehart. Importance sampling for stochastic simulations. *Management Science*, 35(11):1367–1392, 1989.
 - [58] Ramon E Moore. *Interval analysis*, volume 4. Prentice-Hall Englewood Cliffs, NJ, 1966.
 - [59] Kwang Hyung Lee. *First course on fuzzy theory and applications*, volume 27. Springer Science & Business Media, 2004.
 - [60] Grant H Skrepnek. The contrast and convergence of bayesian and frequentist statistical approaches in pharmacoeconomic analysis. *Pharmacoeconomics*, 25(8):649–664, 2007.
 - [61] Julia Koricheva, Jessica Gurevitch, and Kerrie Mengersen. *Handbook of meta- analysis in ecology and evolution*. Princeton University Press, 2013.
 - [62] Alan E Gelfand and Adrian FM Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American statistical association*, 85(410):398– 409, 1990.
 - [63] Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov chain Monte Carlo in practice*. Chapman and Hall/CRC, 1995.
 - [64] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of markov chain monte carlo*. CRC press, 2011.
 - [65] Christian Robert and George Casella. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
 - [66] Jaewook Lee, Woosuk Sung, and Joo-Ho Choi. Metamodel for efficient estimation of capacity-fade uncertainty in li-ion batteries for electric vehicles. *Energies*, 8(6):5538–5554, 2015.
 - [67] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6):721–741, 1984.
 - [68] Persi Diaconis, Kshitij Khare, and Laurent Saloff-Coste. Gibbs sampling, conjugate priors and coupling. *Sankhya A*, 72(1):136–169, 2010.
 - [69] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth,

- Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- [70] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970.
 - [71] Michael Borenstein, Larry V Hedges, Julian PT Higgins, and Hannah R Rothstein. A basic introduction to fixed-effect and random-effects models for meta-analysis. *Research synthesis methods*, 1(2):97–111, 2010.
 - [72] Greg M Allenby and Peter E Rossi. Hierarchical bayes models. *The handbook of marketing research: Uses, misuses, and future advances*, pages 418–440, 2006.
 - [73] L Mark Berliner, Ralph F Milliff, and Christopher K Wikle. Bayesian hierarchical modeling of air-sea interaction. *Journal of Geophysical Research: Oceans*, 108(C4),2003.
 - [74] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*, volume 2. CRC press Boca Raton, FL, 2014.
 - [75] Pejman Honarmandi and Raymundo Arróyave. Uncertainty quantification and propagation in computational materials science and simulation-assisted materials design. *Integrating Materials and Manufacturing Innovation*, pages 1–41.
 - [76] Søren L Frederiksen, Karsten W Jacobsen, Kevin S Brown, and James P Sethna. Bayesian ensemble approach to error estimation of interatomic potentials. *Physicalreview letters*, 93(16):165501, 2004.
 - [77] Fabien Cailliez and Pascal Pernot. Statistical approaches to forcefield calibration and prediction uncertainty in molecular simulation. *The Journal of chemical physics*, 134(5):054124, 2011.
 - [78] Jens Jørgen Mortensen, Kristen Kaasbjerg, Søren Lund Frederiksen, Jens Kehlet Nørskov, James P Sethna, and Karsten Wedel Jacobsen. Bayesian error estimation in density-functional theory. *Physical review letters*, 95(21):216401, 2005.
 - [79] Felix Hanke. Sensitivity analysis and uncertainty calculation for dispersion corrected density functional theory. *Journal of computational chemistry*, 32(7):1424–1430, 2011.
 - [80] Francesco Rizzi, Maher Salloum, YM Marzouk, R-G Xu, ML Falk, TP Weihs, G Fritz, and Omar M Knio. Bayesian inference of atomic diffusivity in a binary ni/al system based on molecular dynamics. *Multiscale Modeling & Simulation*, 9(1):486–512, 2011.
 - [81] Arthur P Dempster. A generalization of bayesian inference. *Journal of the Royal Statistical Society: Series B (Methodological)*, 30(2):205–232, 1968.

- [82] Glenn Shafer. *A mathematical theory of evidence*, volume 42. Princeton university press, 1976.
- [83] Hesheng Tang, Dawei Li, Jingjing Li, and Songtao Xue. Epistemic uncertainty quantification in metal fatigue crack growth analysis using evidence theory. *International Journal of Fatigue*, 99:163–174, 2017.

- [84] S Salehghaffari, M Rais-Rohani, EB Marin, and DJ Bammann. A new approach for determination of material constants of internal state variable based plasticity models and their uncertainty quantification. *Computational Materials Science*, 55:237–244, 2012.
- [85] P Honarmandi and R Arroyave. Using bayesian framework to calibrate a physically based model describing strain-stress behavior of trip steels. *Computational Materials Science*, 129:66–81, 2017.
- [86] F Rizzi, RE Jones, JA Templeton, JT Ostien, and BL Boyce. Plasticity models of material variability based on uncertainty quantification techniques. *arXiv preprint arXiv:1802.01487*, 2017.
- [87] SD Oehler, DJ Hartl, R Lopez, RJ Malak, and DC Lagoudas. Design optimization and uncertainty analysis of sma morphing structures. *Smart Materials and Structures*, 21(9):094016, 2012.
- [88] A Martowicz, J Bryła, and T Uhl. Uncertainty quantification for the properties of a structure made of sma utilizing numerical model. In *Proceedings of the Conference on Noise and Vibration Engineering ISMA 2016 & 5th edition of the International Conference on Uncertainty in Structural Dynamics USD*, 2016.
- [89] ABM Islam and Ernur Karadoğan. Sensitivity and uncertainty analysis of one- dimensional tanaka and liang-rogers shape memory alloy constitutive models. *Materials*, 12(10):1687, 2019.
- [90] Douglas Allaire and Karen Willcox. Fusing information from multifidelity computer models of physical systems. In *2012 15th International Conference on Information Fusion*, pages 2458–2465. IEEE, 2012.
- [91] William D Thomison and Douglas L Allaire. A model reification approach to fusing information from multifidelity information sources. In *19th AIAA non-deterministic approaches conference*, page 1949, 2017.
- [92] Paris Perdikaris, Daniele Venturi, Johannes O Royset, and George Em Karniadakis. Multi-fidelity modelling via recursive co-kriging and gaussian-markov random fields. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2179):20150018, 2015.
- [93] Marc C Kennedy and Anthony O’Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.

- [94] Guangyong Sun, Guangyao Li, Shiwei Zhou, Wei Xu, Xujing Yang, and Qing Li. Multi-fidelity optimization for sheet metal forming process. *Structural and Multidisciplinary Optimization*, 44(1):111–124, 2011.
- [95] P Honarmandi, L Johnson, and R Arroyave. Bayesian probabilistic prediction of precipitation behavior in ni-ti shape memory alloys. *Computational Materials Science*, 172:109334, 2020.
- [96] Anjana Talapatra, Shahin Boluki, Thien Duong, Xiaoning Qian, Edward Dougherty, and Raymundo Arróyave. Autonomous efficient experiment design for materials discovery with bayesian model averaging. *Physical Review Materials*, 2(11):113803, 2018.
- [97] Brian M Adams, WJ Bohnhoff, KR Dalbey, JP Eddy, MS Eldred, DM Gay, K Haskell, Patricia D Hough, and Laura P Swiler. Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: version 5.0 user’s manual. *Sandia National Laboratories, Tech. Rep. SAND2010-2183*, 2009.
- [98] Ben H Thacker, David S Riha, Simeon HK Fitch, Luc J Huyse, and Jason B Pleming. Probabilistic engineering analysis using the nexus software. *Structural Safety*, 28(1-2):83–107, 2006.
- [99] Martin Hunt, Benjamin Haley, Michael McLennan, Marisol Koslowski, Jayathi Murthy, and Alejandro Strachan. Puq: A code for non-intrusive uncertainty propagation in computer simulations. *Computer Physics Communications*, 194:97–107, 2015.
- [100] K Fezi and MJM Krane. Uncertainty quantification of modelling of equiaxed solidification. In *IOP Conference Series: Materials Science and Engineering*, volume 143, page 012028. IOP Publishing, 2016.
- [101] Kyle S Fezi. *Modeling transport phenomena and uncertainty quantification in solidification processes*. PhD thesis, Purdue University, 2016.
- [102] Kyle Fezi and Matthew John M Krane. Uncertainty quantification in modeling metal alloy solidification. *Journal of Heat Transfer*, 139(8):082301, 2017.
- [103] A Plotkowski and MJM Krane. Quantification of epistemic uncertainty in grain attachment models for equiaxed solidification. *Metallurgical and Materials Transactions B*, 48(3):1636–1651, 2017.

- [104] Jiahong Fu, John Coleman, Gregory Poole, Matthew John M Krane, and Amy Marconnet. Uncertainty propagation through a simulation of industrial high pressuredie casting. *Journal of Heat Transfer*, 141(11), 2019.
- [105] Kyle Fezi and Matthew John M Krane. Quantification of input uncertainty propagation through models of aluminum alloy direct chill casting. *Metallurgical and Materials Transactions A*, 49(10):4759–4770, 2018.
- [106] John Allison, Dan Backman, and Leo Christodoulou. Integrated computational materials engineering: a new paradigm for the global materials profession. *Jom*, 58(11):25–27, 2006.
- [107] Daniel G Backman, Daniel Y Wei, Deborah D Whitis, Matthew B Buczek, Peter M Finnigan, and Dongming Gao. Icme at ge: accelerating the insertion of new materials and processes. *JOM*, 58(11):36–41, 2006.
- [108] Greg B Olson and CJ Kuehmann. Materials genomics: from calphad to flight. *Scripta Materialia*, 70:25–30, 2014.
- [109] Wei Xiong and Gregory B Olson. Integrated computational materials design for high-performance alloys. *MRS Bulletin*, 40(12):1035–1044, 2015.
- [110] Georg J Schmitz and Ulrich Prahl. *Integrative computational materials engineering: Concepts and applications of a modular simulation platform*. John Wiley & Sons, 2012.
- [111] John Allison, Dan Backman, and Leo Christodoulou. Integrated computational materials engineering: a new paradigm for the global materials profession. *JOM Journal of the Minerals, Metals and Materials Society*, 58(11):25–27, 2006.
- [112] John Allison. Integrated computational materials engineering: A perspective on progress and future steps. *Jom*, 63(4):15–18, 2011.
- [113] Gregory B Olson. Computational design of hierarchically structured materials. *Science*, 277(5330):1237–1242, 1997.
- [114] Christopher J Roy and William L Oberkampf. A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Computer methods in applied mechanics and engineering*, 200(25-28):2131–2144, 2011.
- [115] Avik Mahata, Tanmoy Mukhopadhyay, and Sondipon Adhikari. A polynomial chaos expansion based molecular dynamics study for probabilistic strength analysis of nano-twinned copper. *Materials Research Express*, 3(3):036501, 2016.
- [116] Maher Salloum and Jeremy 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), 2014.
- [117] Mikhail Khadyko, Jacob Sturdy, Stéphane Dumoulin, Leif Rune Hellevik,

- and Odd Sture Hopperstad. Uncertainty quantification and sensitivity analysis of material parameters in crystal plasticity finite element models. *Journal of Mechanics of Materials and Structures*, 13(3):379–400, 2018.
- [118] Mengfei Yuan, Nicholas Galbincea, Ayman Salem, and Stephen R Niezgoda. Validation and uncertainty quantification for manufacturing design accounting for material variability. In *2018 AIAA Non-Deterministic Approaches Conference*, page 0926, 2018.
 - [119] Anthony O’Hagan. Polynomial chaos: A tutorial and critique from a statistician’s perspective. *SIAM/ASA J. Uncertainty Quantification*, 20:1–20, 2013.
 - [120] Andrew Dienstfrey, Frederick R Phelan, Stephen Christensen, Alejandro Strachan, Fadil Santosa, and Ronald Boisvert. Uncertainty quantification in materials modeling. *JOM*, 66(7):1342–1344, 2014.
 - [121] Marc C. Kennedy and Anthony O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society Series B*, 63(3):425–464, 2001.
 - [122] Jun S Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.
 - [123] Adrian FM Smith and Gareth O Roberts. Bayesian computation via the gibbs sampler and related markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(1):3–23, 1993.
 - [124] Bernd A Berg and Alain Billoire. Markov chain monte carlo simulations. *Wiley Encyclopedia of Computer Science and Engineering*, 2007.
 - [125] Dani Gamerman and Hedibert F Lopes. *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. Chapman and Hall/CRC, 2006.
 - [126] John Barnard, Robert McCulloch, and Xiao-Li Meng. Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*, pages 1281–1311, 2000.

- [127] A James O'Malley and Alan M Zaslavsky. Domain-level covariance analysis for multilevel survey data with structured nonresponse. *Journal of the American Statistical Association*, 103(484):1405–1418, 2008.
- [128] Adrian E Raftery and Steven M Lewis. Implementing mcmc. *Markov chain Monte Carlo in practice*, pages 115–130, 1996.
- [129] Paul M Naghdi. A critical review of the state of finite plasticity. *Zeitschrift für angewandte Mathematik und Physik ZAMP*, 41(3):315–394, 1990.
- [130] Jean-Louis Chaboche. A review of some plasticity and viscoplasticity constitutive theories. *International journal of plasticity*, 24(10):1642–1693, 2008.
- [131] SR Agnew, MH Yoo, and CN Tome. Application of texture simulation to understanding mechanical behavior of mg and solid solution alloys containing lior y. *Acta Materialia*, 49(20):4277–4289, 2001.
- [132] W Wen, M Borodachenkova, CN Tomé, G Vincze, EF Rauch, F Barlat, and JJ Grácio. Mechanical behavior of low carbon steel subjected to strain path changes: Experiments and modeling. *Acta Materialia*, 111:305–314, 2016.
- [133] Ricardo A Lebensohn and CN Tomé. A self-consistent anisotropic approach for the simulation of plastic deformation and texture development of polycrystals: application to zirconium alloys. *Acta metallurgica et materialia*, 41(9):2611–2624, 1993.
- [134] O Engler, CN Tomé, and M-Y Huh. A study of through-thickness texture gradients in rolled sheets. *Metallurgical and materials transactions A*, 31(9):2299–2315, 2000.
- [135] IJ Beyerlein, RA Lebensohn, and CN Tome. Modeling texture and microstructural evolution in the equal channel angular extrusion process. *Materials Science and Engineering: A*, 345(1-2):122–138, 2003.
- [136] Saiyi Li, Irene J Beyerlein, and Mark AM Bourke. Texture formation during equal channel angular extrusion of fcc and bcc materials: comparison with simple shear. *Materials Science and Engineering: A*, 394(1-2):66–77, 2005.
- [137] Pengyang Zhao, Yunzhi Wang, and Stephen R Niezgoda. Microstructural and micromechanical evolution during dynamic recrystallization. *International Journal of Plasticity*, 100:52–68, 2018.

- [138] Sylvain Dancette, Laurent Delannay, Krystel Renard, MA Melchior, and PJ Jacques. Crystal plasticity modeling of texture development and hardening in twip steels. *Acta Materialia*, 60(5):2135–2145, 2012.
- [139] B Liu, D Raabe, F Roters, P Eisenlohr, and RA Lebensohn. Comparison of finite element and fast fourier transform crystal plasticity solvers for texture prediction. *Modelling and Simulation in Materials Science and Engineering*, 18(8):085005, 2010. JFW Bishop and Rodney Hill. Cxxviii. a theoretical derivation of the plastic properties of a polycrystalline face-centered metal. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 42(334):1298–1307, 1951.
- [140] JFW Bishop and Rodney Hill. Xlvi. a theory of the plastic distortion of a polycrystalline aggregate under combined stresses. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 42(327):414–427, 1951.
- [141] Kröner E. On the plastic deformation of polycrystals. *Acta Metall*, 9:155–161, 1961.
- [142] A Molinari, GR Canova, and S Ahzi. A self consistent approach of the large deformation polycrystal viscoplasticity. *Acta Metallurgica*, 35(12):2983–2994, 1987.
- [143] Surya R Kalidindi, Curt A Bronkhorst, and Lallit Anand. Crystallographic texture evolution in bulk deformation processing of fcc metals. *Journal of the Mechanics and Physics of Solids*, 40(3):537–569, 1992.
- [144] RA Lebensohn, P Ponte Castañeda, R Brenner, and O Castelnau. Full- field vs. homogenization methods to predict microstructure–property relations for polycrystalline materials. In *Computational Methods for Microstructure-Property Relationships*, pages 393–441. Springer, 2011.
- [145] Dierk Raabe, Michael Sachtleber, Hasso Weiland, Georg Scheele, and Zisu Zhao. Grain-scale micromechanics of polycrystal surfaces during plastic straining. *Acta Materialia*, 51(6):1539–1560, 2003.
- [146] AA Salem, SR Kalidindi, and SL Semiatin. Strain hardening due to deformation twinning in α -titanium: Constitutive relations and crystal-plasticity modeling. *Acta Materialia*, 53(12):3495–3502, 2005.
- [147] Florent Bridier, David L McDowell, Patrick Villechaise, and Jose´ Mendez. Crystal plasticity modeling of slip activity in ti–6al–4v under high cycle fatigue loading. *International Journal of Plasticity*, 25(6):1066–1082, 2009.

- [148] Dhyanjyoti Deka, Deepu S Joseph, Somnath Ghosh, and Michael J Mills. Crystal plasticity modeling of deformation and creep in polycrystalline ti-6242. *Metallurgical and materials transactions A*, 37(5):1371–1388, 2006.
- [149] Anaïs Gaubert, Yann Le Bouar, and Alphonse Finel. Coupling phase field and viscoplasticity to study rafting in ni-based superalloys. *Philosophical Magazine*, 90(1-4):375–404, 2010.
- [150] Mark A Miodownik. A review of microstructural computer models used to simulate grain growth and recrystallisation in aluminium alloys. *Journal of Light Metals*, 2(3):125–135, 2002.
- [151] Surya R Kalidindi. Incorporation of deformation twinning in crystal plasticity models. *Journal of the Mechanics and Physics of Solids*, 46(2):267–290, 1998.
- [152] H Wang, PD Wu, J Wang, and CN Tomé. A crystal plasticity model for hexagonal close packed (hcp) crystals including twinning and de-twinning mechanisms. *International Journal of Plasticity*, 49:36–52, 2013.
- [153] LP Evers, WAM Brekelmans, and MGD Geers. Scale dependent crystal plasticity framework with dislocation density and grain boundary effects. *International Journal of solids and structures*, 41(18-19):5209–5230, 2004.
- [154] S Mercier, A Molinari, S Berbenni, and M Berveiller. Comparison of different homogenization approaches for elastic–viscoplastic materials. *Modelling and Simulation in Materials Science and Engineering*, 20(2):024004, 2012.
- [155] John Douglas Eshelby. The determination of the elastic field of an ellipsoidal inclusion, and related problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 241(1226):376–396, 1957.
- [156] RODNEY Hill. Continuum micro-mechanics of elastoplastic polycrystals. *Journal of the Mechanics and Physics of Solids*, 13(2):89–101, 1965.
- [157] RA Lebensohn, CN Tomé, and P Ponte Castaneda. Self-consistent modelling of the mechanical behavior of viscoplastic polycrystals incorporating intragranular field fluctuations. *Philosophical Magazine*, 87(28):4287–4322, 2007.
- [158] E Voce. The relationship between stress and strain for homogeneous deformation. *Journal of the Institute of Metals*, 74:537–562, 1948.
- [159] E Voce. A practical strain hardening function. *Metallurgia*, 51:219–226, 1955.
- [160] UF Kocks. Laws for work-hardening and low-temperature creep. *Journal of engineering materials and technology*, 98(1):76–85, 1976.
- [161] Ricardo A Lebensohn. N-site modeling of a 3d viscoplastic polycrystal using fast fourier transform. *Acta Materialia*, 49(14):2723–2737, 2001.
- [162] D Kiener, W Grosinger, G Dehm, and R Pippan. A further step towards an

- understanding of size-dependent crystal plasticity: In situ tension experiments of miniaturized single-crystal copper samples. *Acta Materialia*, 56(3):580–592, 2008.
- [163] E8-99 ASTM et al. Standard test methods for tension testing of metallic materials. *Annual book of ASTM standards*. ASTM, 2001.
 - [164] A.T. Dinsdale. Sgte data for pure elements. *Calphad*, 15(4):317–425, 1991.
 - [165] Luiz TF Eleno and Cláudio G Schön. Ternary interaction parameters in calphad solution models. *Brazilian Journal of Physics*, 44(2-3):208–214, 2014.
 - [166] Ales Kroupa. Modelling of phase diagrams and thermodynamic properties using calphad method—development of thermodynamic databases. *Computational Materials Science*, 66:3–13, 2013.
 - [167] Francis A Shunk, Max Hansen, and Kurt Anderko. *Constitution of binary alloys*. Number 2. McGraw-Hill, 1969.
 - [168] Jan-Olof Andersson, Thomas Helander, Lars Höglund, Pingfang Shi, and Bo Sundman. Thermo-calc & dictra, computational tools for materials science. *Calphad*, 26(2):273–312, 2002.
 - [169] AT Dinsdale. Sgte unary database, 2001.
 - [170] RH Davies, AT Dinsdale, JA Gisby, JAJ Robinson, and S M Martin. Mtdata-thermodynamic and phase equilibrium software from the national physical laboratory. *Calphad*, 26(2):229–271, 2002.
 - [171] W Cao, YA Chang, J Zhu, S Chen, and WA Oates. Thermodynamic modeling of the cu–ag–au system using the cluster/site approximation. *Intermetallics*, 15(11):1438–1446, 2007.
 - [172] XC He, H Wang, HS Liu, and ZP Jin. Thermodynamic description of the cu–ag–zr system. *Calphad*, 30(4):367–374, 2006.

- [173] S-L Chen, S Daniel, F Zhang, YA Chang, X-Y Yan, F-Y Xie, R Schmid-Fetzer, and WA Oates. The pandat software package and its applications. *Calphad*, 26(2):175–188, 2002.
- [174] Denielle E Ricciardi, Oksana A Chkrebtii, and Stephen R Niezgoda. Uncertainty quantification for parameter estimation and response prediction. *Integrating Materials and Manufacturing Innovation*, 8(3):273–293, 2019.
- [175] Bertil Matérn. Spatial variation, volume 36 of. *Lecture Notes in Statistics*, 1960.
- [176] Ottar N Bjørnstad and Wilhelm Falck. Nonparametric spatial covariance functions: estimation and testing. *Environmental and Ecological Statistics*, 8(1):53–70, 2001.
- [177] Scott Fortmann-Roe. Understanding the bias-variance tradeoff. I: URL: <http://scott.fortmann-roe.com/docs/BiasVariance.html> (hämtad 2017-09-29), 2012.
- [178] H Wang, PD Wu, CN Tomé, and Yonggang Huang. A finite strain elastic–viscoplastic self-consistent model for polycrystalline materials. *Journal of the Mechanics and Physics of Solids*, 58(4):594–612, 2010.

APPENDIX A

A.1 MCMC SAMPLER IMPLEMENTATION

Algorithm 1 Algorithm

- 1: Input: initialize the vector of unknown model parameters $\Phi^{(0)} = (\theta^{[1](0)}, \dots, \theta^{[S](0)}, \theta^{(0)})$ and partition it into $\mathbf{B}$ blocks denoted by $\Phi^{1(0)}, \dots, \Phi^{B(0)}$, respectively. Denote by $\Phi^{-b(0)}$ the vector $\Phi^{(0)}$ with the b th block removed. Also initialize parameters $\Lambda^{(0)}$ and $\delta^{2(0)}$.
 - 2: **for** $m = 1:M$ **do**
 - 3: Set $\Phi^{(m)} = \Phi^{(m-1)}$.
 - 4: **for** $b = 1:B$ **do**
 - 5: Sample $\Phi^{b(p)}$ from the symmetric proposal distribution with density $p(\Phi^b | \Phi^{b(m-1)})$.
 - 6: Generate a sample u from the uniform distribution on the interval $(0,1)$.
 - 7: **if** $\ln(u) < \ln \pi(\Phi^{b(p)}, \Phi^{-b(m)}, \delta^{2(m-1)}, \Lambda^{(m-1)} | y^{[1]}, \dots, y^{[S]}) - \ln \pi(\Phi^{(m-1)}, \delta^{2(m-1)}, \Lambda^{(m-1)} | y^{[1]}, \dots, y^{[S]})$ **then**
 - 8: Set $\Phi^{b(m)} = \Phi^{b(p)}$.
 - 9: **end if**
 - 10: **end for**
 - 11: Sample $\Lambda^{(m)}$ from full conditional distribution with density $f(\Lambda | y^{[1]}, \dots, y^{[S]}, \Phi^{(m)})$ (Eq. (3.14)).
 - 12: Sample $\delta^{2(m)}$ from full conditional distribution with density $f(\delta^2 | y^{[1]}, \dots, y^{[S]}, \Phi^{(m)})$ (Eq. (3.15)).
 - 13: Sample $\theta^{[new](m)}$ from the random effect distribution with density $f(\theta^{[new]} | \theta^{(m)}, \Lambda^{(m)})$ (Eq. (3.5)).
 - 14: Sample $y^{[new](m)}$ from the likelihood $f(y^{[new]} | \theta^{[new]}, \Psi^{(m)})$ (Eq. (3.6)).
 - 15: **end for**
 - 16: Output: $\{(\Phi^{(1)}, \Lambda^{(1)}, \delta^{2(1)}), \dots, (\Phi^{(M)}, \Lambda^{(M)}, \delta^{2(M)})\}$ (realizations of MC targeting the posterior in Eq. (3.7)); $\{y^{[new](1)}, \dots, y^{[new](M)}\}$ (realizations of the MC targeting the posterior predictive density in Eq. (3.8)).
-

A.2 DERIVATION OF FULL-CONDITIONALS

Derivation of full conditional distributions over δ^2 , Ψ and Λ . Note that the ℓ_2 normdefined in eq. (3.13) is utilized here.

$$\begin{aligned}
 f(\delta^2 | y^{[1]}, \dots, y^{[s]}, \theta^{[1]}, \dots, \theta^{[s]}, \theta) &\propto \prod_{s=1}^S MVN_N(y^{[s]}; m(\theta^{[s]}), \delta^{-2} \mathbb{I}_N) \text{Gamma}(\delta^2; a, b), \\
 a, b &> 0 \\
 &\propto \det(\delta^2 \mathbb{I}_N)^{\frac{S}{2}} \exp\left\{-\frac{1}{2} \sum_{s=1}^S \|y^{[s]} - m(\theta^{[s]})\|_2^2 (\delta^2 \mathbb{I}_N)\right\} (\delta^2)^{a-1} \exp\{-b\delta^2\}, \\
 a, b &> 0 \\
 &= (\delta^2)^{\frac{SN}{2} + a - 1} \exp\left\{-\delta^2 \left(b + \frac{1}{2} \sum_{s=1}^S \|y^{[s]} - m(\theta^{[s]})\|_2^2\right)\right\}, \\
 a, b &> 0 \\
 &\propto \text{Gamma}\left(\delta^2; \frac{SN}{2} + a, b + \frac{1}{2} \sum_{s=1}^S \|y^{[s]} - m(\theta^{[s]})\|_2^2\right) \\
 a, b &> 0
 \end{aligned} \tag{A.1}$$

$$\begin{aligned}
 f(\Psi^2 | y^{[1]}, \dots, y^{[s]}, \theta^{[1]}, \dots, \theta^{[s]}, \theta) &\propto \prod_{s=1}^S MVN_N(y^{[s]}; m(\theta^{[s]}), \Psi^{-1}) \text{Wishart}_N(\Psi; V_o, v_o), \\
 v_o &\geq N + 1 \\
 &\propto \det(\Psi)^{\frac{S}{2}} \exp\left\{-\frac{1}{2} \sum_{s=1}^S \|y^{[s]} - m(\theta^{[s]})\|_2^2 (\Psi)\right\} \det(\Psi)^{(v_o - N - 1)/2} \exp\left\{-\text{Tr}\left(\frac{V_o^{-1} \Psi}{2}\right)\right\}, \\
 v_o &\geq N + 1 \\
 &\propto \det(\Psi)^{(v_o - N - 1 + S)/2} \exp\left\{-\frac{1}{2} \text{Tr}\left(\Psi \left[V_o^{-1} + \sum_{s=1}^S \|y^{[s]} - m(\theta^{[s]})\|_2^2\right]\right)\right\}, \\
 v_o &\geq N + 1 \\
 &\propto \text{Wishart}_N\left(\Psi; \left[V_o^{-1} + \sum_{s=1}^S \|y^{[s]} - m(\theta^{[s]})\|_2^2\right]^{-1}, v_o + S\right), \\
 v_o &\geq N + 1
 \end{aligned} \tag{A.2}$$

$$\begin{aligned}
 f(\Lambda | y^{[1]}, \dots, y^{[s]}, \theta^{[1]}, \dots, \theta^{[s]}, \theta) &\propto \prod_{s=1}^S MVN_D(\theta^{[s]}; \theta, \Lambda^{-1}) \text{Wishart}_D(\Lambda; V_o, v_o), \\
 v_o &\geq D + 1
 \end{aligned}$$

$$\begin{aligned}
& \propto \det(\Lambda)^{\frac{S}{2}} \exp \left\{ -\frac{1}{2} \sum_{s=1}^S \|\theta^{[s]} - \theta\|_2^2(\Lambda) \right\} \det(\Lambda)^{(v_o - D - 1)/2} \exp \left\{ -\text{Tr} \left(\frac{V_o^{-1} \Lambda}{2} \right) \right\}, \\
& v_o \geq D + 1 \\
& \propto \det(\Lambda)^{(v_o - D - 1 + S)/2} \exp \left\{ -\frac{1}{2} \text{Tr} \left(\Lambda \left[V_o^{-1} + \sum_{s=1}^S \|\theta^{[s]} - \theta\|_2^2 \right] \right) \right\}, \\
& v_o \geq D + 1 \\
& \propto \text{Wishart}_D \left(\Lambda; \left[V_o^{-1} + \sum_{s=1}^S \|\theta^{[s]} - \theta\|_2^2 \right]^{-1}, v_o + S \right), \\
& v_o \geq D + 1
\end{aligned} \tag{A.3}$$

A.3 PARAMETER CORRELATION

See Tables A.1, A.2, A.3, and A.4.

Table A 1: Element-wise posterior correlation for a random effect from case study I

	$\tau_0^{[1]}$	$\tau_1^{[1]}$	$\xi_0^{[1]}$	$\xi_1^{[1]}$
$\tau_0^{[1]}$	1.0			
$\tau_1^{[1]}$	.38	1.0		
$\xi_0^{[1]}$	-.83	-.74	1.0	
$\xi_1^{[1]}$	-.38	-1.0	.74	1.0

Table A 2: Element-wise posterior correlation for the overall effect from case study I

	τ_0	τ_1	ξ_0	ξ_1
τ_0	1.0			
τ_1	.75	1.0		
ξ_0	.96	.82	1.0	
ξ_1	.59	-.020	.51	1.0

Table A 3: Element-wise posterior correlation for the random effect from case study II

	$L_0^{[1]}$	$L_1^{[1]}$
$L_0^{[1]}$	1.0	
$L_1^{[1]}$	.33	1.0

Table A 4: Element-wise posterior correlation for the overall effect from case study II

	L_0	L_1
L_0	1.0	
L_1	.97	1.0

A.4 ADDITIONAL FIGURES

Marginal posterior distributions from case studies I and II. Figures A.1 and A.2 compare the distributions between the overall (dashed lines) and random effects (solid lines). Figure A 3 shows the marginal histogram for the error precision plotted with its prior density distribution (red line) for case study I.






Figure A 1: Scaled marginal posterior densities from case study I for random effects (solid lines) and the overall effect (dashed lines) for each Voce hardening parameter.



(a) L_0



(b) L_1

Figure A 2: Scaled marginal histograms for Redlich-Kister parameters at 1000K. The parameters are L_0 in (a) and L_1 in (b) and include all random effects (with 'o' markers) as well as the overall effect parameters (dashed lines).




Figure A 3: Marginal histogram for the error precision plotted with its prior density distribution (red line) for case study I.

APPENDIX B

B.1 MCMC SAMPLER IMPLEMENTATION

Algorithm 2 Algorithm

- 1: Input: initialize the vector of unknown model parameters $\Phi^{(0)} = (\theta^{[1](0)}, \dots, \theta^{[S](0)}, \theta^{(0)})$ and partition it into B blocks denoted by $\Phi^{1(0)}, \dots, \Phi^{B(0)}$, respectively. Denote by $\Phi^{-b(0)}$ the vector $\Phi^{(0)}$ with the b th block removed. Also initialize parameters $R^{(0)}, (t^2)^{(0)}, \Delta^{(0)}$ and $\delta^{2(0)}$.
 - 2: Calculate the covariance structure Γ by (4.27).
 - 3: **for** $m = 1:M$ **do**
 - 4: Set $\Phi^{(m)} = \Phi^{(m-1)}$.
 - 5: **for** $b = 1:B$ **do**
 - 6: Sample $\Phi^{b(p)}$ from the symmetric proposal distribution with density $p(\Phi^b | \Phi^{b(m-1)})$.
 - 7: Generate a sample u from the uniform distribution on the interval $(0, 1)$.
 - 8: **if** $\ln(u) < \ln \pi(\Phi^{b(p)}, \Phi^{-b(m)}, \delta^{2(m-1)}, \Lambda^{(m-1)} | y^{[1]}, \dots, y^{[S]}) - \ln \pi(\Phi^{(m-1)}, \delta^{2(m-1)}, \Lambda^{(m-1)} | y^{[1]}, \dots, y^{[S]})$ **then**
 - 9: Set $\Phi^{b(m)} = \Phi^{b(p)}$.
 - 10: **end if**
 - 11: **end for**
 - 12: Sample $R^{(m)}$ from full conditional distribution with density $f(R | \{y^{[s]}\}_{s=1}^S, \Phi^{(m)})$.
 - 13: Sample $(\delta^2)^{(m)}$ from full conditional distribution with density $f(\delta^2 | \Delta, \{y^{[s]}\}_{s=1}^S, \Phi^{(m)})$.
 - 14: Sample $\Delta^{(m)}$ from full conditional distribution with density $f(\Delta | \Psi, \{y^{[s]}\}_{s=1}^S, \Phi^{(m)})$.
 - 15: Sample $\theta^{[new](m)}$ from the random effect distribution with density $f(\theta^{[new]} | \theta^{(m)}, \Lambda^{(m)})$.
 - 16: Sample $y^{[new](m)}$ from the likelihood $f(y^{[new]} | \theta^{[new]}, \Psi^{(m)}, \Delta^{(m)})$.
 - 17: **end for**
 - 18: Output: $(\Phi^{(1)}, \Lambda^{(1)}, \delta^{2(1)}, \Delta^{(1)}), \dots, (\Phi^{(M)}, \Lambda^{(M)}, \delta^{2(M)}, \Delta^{(M)})$ (realizations of MC targeting the posterior in Eq. (4.12)); $y^{[new](1)}, \dots, y^{[new](M)}$ (realizations of the MC targeting the posterior predictive density in Eq. (4.14)).
-

B.2 DERIVATION OF FULL-CONDITIONALS

Derivation of full conditional distributions over δ^2 , Ψ , R and Δ .

$$\begin{aligned}
 f\left(\delta^2 \mid \{y^{[s]}\}_{s=1}^S, \{\theta^{[s]}\}_{s=1}^S, \theta, \Delta\right) &\propto \prod_{s=1}^S MVN_N\left(y^{[s]} - m\left(\theta^{[s]}\right); \Delta, \delta^{-2} \mathbb{I}_N\right) \text{Gamma}\left(\delta^2; a, b\right), \\
 a, b &> 0 \\
 &\propto \det\left(\delta^2 \mathbb{I}_N\right)^{\frac{S}{2}} \exp\left\{-\frac{1}{2} \sum_{s=1}^S \left\|y^{[s]} - m\left(\theta^{[s]}\right) - \Delta\right\|_2^2 (\delta^2 \mathbb{I}_N)\right\} (\delta^2)^{a-1} \exp\{-b\delta^2\}, \\
 a, b &> 0 \\
 &= (\delta^2)^{\left(\frac{SN}{2} + a - 1\right)} \exp\left\{-\delta^2 \left(b + \frac{1}{2} \sum_{s=1}^S \left\|y^{[s]} - m\left(\theta^{[s]}\right) - \Delta\right\|_2^2\right)\right\}, \\
 a, b &> 0 \\
 &\propto \text{Gamma}\left(\delta^2; \frac{SN}{2} + a, b + \frac{1}{2} \sum_{s=1}^S \left\|y^{[s]} - m\left(\theta^{[s]}\right) - \Delta\right\|_2^2\right), \\
 a, b &> 0
 \end{aligned} \tag{B.1}$$

$$\begin{aligned}
 f(\Psi \mid \{y^{[s]}\}_{s=1}^S, \{\theta^{[s]}\}_{s=1}^S, \theta, \Delta) &\propto \prod_{s=1}^S MVN_N\left(y^{[s]} - m\left(\theta^{[s]}\right); \Delta, \Psi^{-1}\right) \text{Wishart}_N(\Psi; V_0, v_0), \\
 v_0 &\geq N + 1 \\
 &\propto \det(\Psi)^{\frac{S}{2}} \exp\left\{-\frac{1}{2} \sum_{s=1}^S \left\|y^{[s]} - m\left(\theta^{[s]}\right) - \Delta\right\|_2^2 (\Psi)\right\} \det(\Psi)^{(v_0 - N - 1)/2} \exp\left\{-\text{Tr}\left(\frac{V_0^{-1} \Psi}{2}\right)\right\}, \\
 v_0 &\geq N + 1 \\
 &\propto \det(\Psi)^{(v_0 - N - 1 + S)/2} \exp\left\{-\frac{1}{2} \text{Tr}\left(\Psi \left[V_0^{-1} + \sum_{s=1}^S \left\|y^{[s]} - m\left(\theta^{[s]}\right) - \Delta\right\|_2^2\right]\right)\right\}, \\
 v_0 &\geq N + 1 \\
 &\propto \text{Wishart}_N\left(\Psi; \left[V_0^{-1} + \sum_{s=1}^S \left\|y^{[s]} - m\left(\theta^{[s]}\right) - \Delta\right\|_2^2\right]^{-1}, v_0 + S\right), \\
 v_0 &\geq N + 1
 \end{aligned} \tag{B.2}$$

$$\begin{aligned}
f(R \mid \{y^{[s]}\}_{s=1}^S, \{\theta^{[s]}\}_{s=1}^S, \theta) &\propto \left\{ \prod_{s=1}^S MVN_D \left(\theta^{[s]}; \theta, \Lambda^{-1} \right) \right\} Wishart_D(R; R_o, r_o), \\
r_o &\geq D+1 \\
&\propto \det(\Lambda)^{\frac{S}{2}} \exp \left\{ -\frac{1}{2} \sum_{s=1}^S \left\| \theta^{[s]} - \theta \right\|_2^2 (\Lambda) \right\} \det(R)^{(r_o-D-1)/2} \exp \left\{ -\text{Tr} \left(\frac{R_o^{-1} R}{2} \right) \right\}, \\
r_o &\geq D+1 \\
&\propto \det(R)^{\frac{1}{2}(S+r_o-D-1)} \exp \left\{ -\frac{1}{2} \text{Tr} \left(R \left[R_o^{-1} + \text{diag}(t) \left(\sum_{s=1}^S \left\| \theta^{[s]} - \theta \right\|_2^2 \right) \text{diag}(t) \right] \right) \right\}, \\
r_o &\geq D+1 \\
&\propto Wishart_D \left(R; \left[R_o^{-1} + \text{diag}(t) \left(\sum_{s=1}^S \left\| \theta^{[s]} - \theta \right\|_2^2 \right) \text{diag}(t) \right]^{-1}, r_o + S \right), \\
r_o &\geq D+1
\end{aligned} \tag{B.3}$$

$$f(\Delta \mid \Psi, \{y^{[s]}\}_{s=1}^S, \{\theta^{[s]}\}_{s=1}^S, \theta) \propto \prod_{s=1}^S MVN_N \left(y^{[s]} - m(\theta^{[s]}); \Delta, \Psi^{-1} \right) MVN_N(\Delta; 0, \Gamma)$$

which is just a normal-normal hierarchy with full-conditional,

$$\Delta \mid \{y^{[s]}\}_{s=1}^S, \{\theta^{[s]}\}_{s=1}^S, \Psi \sim MVN_N(\Delta; m, V)$$

$$f(\Delta \mid \Psi, \{y^{[s]}\}_{s=1}^S, \{\theta^{[s]}\}_{s=1}^S, \theta) \propto MVN_N(\Delta; m, V)$$

where,

$$m = (\Gamma^{-1} + N\Psi)^{-1} N\Psi \left(\overline{y^{[s]} - m(\theta^{[s]})} \right)$$

$$V = (\Gamma^{-1} + N\Psi)^{-1}.$$

(B.4)

B.3 PARAMETER CORRELATION

See Tables B.1 and B.2.

Table B 1: Element-wise posterior correlation for a representative random effect from case study I

	$\tau_0^{[1]}$	$\tau_1^{[1]}$	$\xi_0^{[1]}$	$\xi_1^{[1]}$
$\tau_0^{[1]}$	1.0			
$\tau_1^{[1]}$	.030	1.0		
$\xi_0^{[1]}$	-.91	-.31	1.0	
$\xi_1^{[1]}$	.010	-1.0	.25	1.0

Table B 2: Element-wise posterior correlation for the overall effect from case study I

	τ_0	τ_1	ξ_0	ξ_1
τ_0	1.0			
τ_1	-.63	1.0		
ξ_0	.82	-.56	1.0	
ξ_1	.75	-.94	.67	1.0

B.4 ADDITIONAL FIGURES

Marginal posterior distributions: Figure B.1 compares the distributions between the overall (dashed lines) and random effects (solid lines).




Figure B 1: Scaled marginal posterior densities from case study I for random effects (solid lines) and the overall effect (dashed lines) for each Voce hardening parameter

APPENDIX C

The following contains the MCMC algorithm written in MATLAB which was applied to the study. Details of the code can be found in a previous section.

C.1 MCMC MAIN

```
1 % MCMC
2 %   simulation using an adaptive MH-MCMC
3 %   algorithm to target the posterior distribution
4 %   of interest with Gibbs steps where appropriate
5 %
6 % Example Usage
7 %   Parameter estimation for the VPSC model (Tome and Lebensohn)
8 %   taking into account various sources of uncertainty stemming
9 %   from noisy and/or inconsistent data (both epistemic and
10 %   aleatoric) and model-form error
11 %
12 % Unknown Parameters
13 %   theta^[s]: VPSC model parameters for the random effects
14 %       MH (blocks 1:S)
15 %   theta: VPSC model parameters for overall effect
16 %       MH (block S+1)
17 %   Lambda: Random effects precision
18 %   R: Decomposed correlation of Lambda
19 %       Gibbs
20 %   t2: Decomposed variance of Lambda
21 %       MH (block S+2)
22 %   delta: Error precision
23 %       Gibbs
24 %   Delta: Discrepancy (Model-form error)
25 %       Gibbs
26 %
27 % Output
```

```

28 % k = 'num iter' samples targeting the posterior distribution of
29 % all relevant parameters as well as posterior and posterior
30 % predictive samples
31 %
32 %
33 %Copyright (c) 2020, Denielle Ricciardi
34 %All rights reserved.
35
36
37 %% USER DEFENITIONS
38
39 num iter = 1E5;
40 thin = 1; %factor chain is thinned by
41 num par = 4; %physical model unknown parameters
42 store ppd = 500; %freq of storing posterior and
    posterior ...predictive samples
43 start fix = 1; %start adaptive proposal covariance at this iteration
44 stop fix = 25000; %end adaptive proposal covariance at this ...
    iteration
45 covariance check = 2000; %frequency of adapting porposal covariance
46 target accept = [0.2,0.5]; %target acceptance rate
47
48 %load data
49 obs = 'import data'; %observations
50 S = size(obs,2); %number of random effects
51
52 cov mat = cell(1,S + 2); %Initialize proposal covaraince
53 for i = 1:size(cov mat,2)
54     cov mat i{ } = eye(num par)*10; %initiate
55 end
56
57 %% CREATE BLOCKS
58
59 %One block per RE + 1 for OE + 1 for t^2
60 v = 1:S+2; % + 2 (1 for overall and 1 for t2)
61 blocks = repelem(v,num par);%define blocks
62 theta chain index = 1:1:length(blocks); %7*voce
63 u = unique(blocks);
64 num blocks = length(u);
65 block index = cell(1,num blocks);
66 ml par = cell(1,num blocks);
67 for nn = 1:num blocks
68     block index nn{ } = theta chain index(blocks == u(nn));
69 end
70
71 %% LOAD DATA AND SET CODE DEFENITIONS
72
73 sigma obs = zeros(size(obs.SVM,1),S); %observed stress
74 S = size(sigma obs,2);
75 D = 4;
76 N = length(sigma obs);

```

```

77
78 for i = 1:S
79     sigma obs(:,i) = obs.SVM(:,i);
80 end
81
82 strain inc = obs.EVM; %strain increments for observed stresses
83
84 %Proposal Distribution
85 proposal sample = @(mu,proposal variance) ...
86     mvnrnd(mu,proposal variance);
87
88 %% INITIALIZE ARRAYS
89
90 theta chain = zeros(num iter/thin+1,block index S+1 (D));
91 Lambda chain = zeros(D,D,num iter/thin+1);
92 delta chain = zeros(num iter/thin+1,1);
93 Delta chain = zeros(N,num iter/thin+1);
94
95 %parameters in Lambda decomposition
96 t2 chain = zeros(D,num iter/thin+1);
97 R chain = zeros(D,D,num iter/thin+1);
98
99 %array holding covariance check number of samples for full adjustment
100 %lines (242-265) - for theta^[s], theta, and t2
101
102 temp theta chain = zeros(covariance check,block index{S+2}(D));
103
104 log post = zeros(num iter/thin+1,1);
105 acceptance = zeros(num iter,S+2);
106 ppd draw = zeros(N,num iter/store ppd);
107 post draw = zeros(N,num iter/store ppd);
108
109 %% INITIALIZE MARKOV CHAIN
110
111 %hyperparameters for error precision - Gamma
112 a delta = 1; %SHAPE parameter
113 b delta = .1; %RATE parameter
114
115 % hyperparameters for R - Wishart
116 r not = D + 2;
117 R not = 1/r not*eye(D);
118
119 % hyperparameters for t - Gamma
120 a t = 1; %SHAPE
121 b t = 0.1; %RATE
122
123 % initialize chains
124 theta chain(1,:) = repmat([55 55 700 150],1,S+1);
125 R chain(:, :, 1) = r not.*R not;
126 t2 chain(:, 1) = ones(D,1)*.1;
127 Lambda chain(:, :, 1) = diag(sqrt(t2 chain(:, 1)))*...
128     R chain(:, :, 1)*diag(sqrt(t2 chain(:, 1)));

```

```

127 delta chain(1,:) = 10;
128 Delta chain(:,1) = diag(ones(N));
129
130 current theta = theta chain(1,:);
131 current Lambda = Lambda chain(:, :, 1);
132 current delta = delta chain(1,1);
133 current t2_ = t2 chain(:,1);
134 current R = R chain(:, :, 1);
135 current Delta = Delta chain(:,1);
136
137 temp chain(1,1:block index{S+1}(D)) = theta chain(1,:);
138 temp chain(1,block index{S+2}) = t2 chain(:,1)';
139
140 %% INITIALIZE AND DEFINE COVARIANCE STRUCTURE FOR DISCREPANCY
141
142 discretization = diff(strain inc);
143 w 1 = .35*discretization(1);
144 w 2 = 50*discretization(2);
145 a = 0; %left limit of domain
146 nug 1 = .1;
147 nug 2 = 1E-9;
148 signal variance 1 = 1E14;
149 signal variance 2 = 1E-4;
150 limit = .0025; %transition point between elastic and plastic
151
152 %flags
153 show figs = 1;
154 analyze = 1;
155
156 [Gamma,e trans,cov 1,cov 2] = ...
157     gen cov(signal variance 1,signal variance 2,...
158     w 1,w 2,nug 1,nug 2,N,strain inc,limit,show figs,analyze);
159
160 log det Gamma = double(vpa(log(det(vpa(Gamma,1000))),1000));
161
162 % Calculate Inverse of Covariance structure
163 digitsOld = digits;
164 digits(16)
165 syms A
166 A = vpa(Gamma);
167
168 disp(['Condition of matrix is ',num2str(cond(double(A)))])
169
170 inv Gamma = vpa(inv(Gamma));
171
172 if issymmetric(double(inv Gamma))
173     warning('Inverse is not symmetric')
174 end
175
176 if isposdef(double(inv Gamma))
177     warning('The inverse is not positive definite')

```

```

178 end
179
180 disp([' || I - inv(Gamma)*Gamma || = ',...
181       num2str(double(norm(eye(N) - inv Gamma*Gamma,1))))]
182 dubIG = double(inv Gamma);
183
184 %% INITIAL LOG LIKELIHOOD AND LOG POSTERIOR EVALUATION
185
186 sigma vpsc = zeros(N,S);
187
188 for i = 1:S
189     par = theta chain(1,block index i);
190     sigma vpsc(:,i) = VPSC(par(1),par(2),par(3),par(4),strain inc);
191 end
192
193 [initial log likelihood,vpsc prop] = ...
194     log likelihood(S,N,D,theta chain(1,:),Lambda chain(:,1),...
195     delta chain(1),Delta chain(:,1),sigma obs,strain inc,...
196     sigma vpsc,1,block index);
197
198 sigma vpsc = vpsc prop;
199
200 initial log posterior = log posterior(S,N,D,block index,...
201     initial log likelihood,theta chain(1,:),t2 chain(:,1),...
202     R chain(:,1),delta chain(1),Delta chain(:,1),log det Gamma,...
203     dubIG,a delta,b delta,R not,r not,a t,b t);
204
205 log post(1,1) = initial log posterior;
206 current log post = log post(1,1);
207
208 %% MCMC LOOP
209
210 for k = 1:num iter
211     if mod(k,covariance check) == 0
212         if k ≤ stop fix && k ≥ start fix
213             accept rate = ...
214                 (1 + sum(acceptance(k-covariance check + ...
215                     1:k,:))./covariance check;
216             for aa = 1:num blocks
217                 if accept rate(aa) < target accept(1)
218                     accept rate(aa) > target accept(2)
219                     C = cov(temp chain(:,block index{aa}));
220                     cov mat{aa}(logical(eye(D))) = ...
221                         diag(C).*accept rate(aa)./(target accept(1)...
222                             + range(target accept)/2);
223                     cov mat aa{ } = ...
224                         full cov mat adj(temp chain(:,blocks ==
225                             ...aa),...
226                             D,cov mat aa{ });
227                 end
228             end
229         end
230     end

```

```

228     end
229     temp chain = zeros(covariance check,block index S+{ D});
230 end
231
232 proposed theta = current theta;
233 for b = 1:length(unique(blocks)) - 1
234     Kosher = 0;
235     while Kosher ==0
236         proposed theta(block index{b}) = ...
237             proposal sample(current theta(block index{b}),...
238                 cov mat{b});
239         if proposed theta(b*D) ≥ proposed theta(b*D-1)
240             continue
241         end
242         if proposed theta(b*D) < 0 |proposed theta(b*D-3)
243             < ...0 ...
244             |proposed theta(b*D-2) < 0
245             continue
246         end
247         Kosher = 1;
248     end
249     try
250
251         u = unifrnd(0,1);
252
253         [proposed log likelihood,vpsc prop] = ...
254             log likelihood(S,N,D,proposed theta,current Lambda,...
255             current delta,current Delta,sigma obs,strain inc,...
256             sigma vpsc,b,block index);
257
258         proposed log posterior = ...
259             log posterior(S,N,D,block index,...
260             proposed log likelihood,proposed theta,...
261             current t2,current R,current delta,...
262             current Delta,log det Gamma,dubIG,...
263             a delta,b delta,R not,r not,a t,b t);
264
265         if log(u) < proposed log posterior-current log post %accept
266             current theta(block index b { }
267                 = ...proposed theta(block { }
268                 index b );
269             current log post = proposed log posterior;
270             acceptance(k,b) = 1;
271             sigma vpsc = vpsc prop;
272         else % reject
273             acceptance(k,b) = 0;
274             proposed theta(block index b { } = ...
275                 current theta(block index b { }

```

```

276         catch
277             sprintf('Error in MH step iteration (%i,%i)',k,b)
278             kk = k+1;
279             sprintf(' n\Continuing to next iteration with...
280                 parameter(%i,%i) = parameter(%i,%i)',kk,b,k,b)
281             proposed theta(block index b { } = ...
282                 current theta(block index b { })
283             acceptance(k,b) = 0;
284             continue
285         end
286     end
287
288     %sample Psi (delta), Lambda and Delta from
289     %full conditional (gibbs sampler)
290     %sampled directly from posterior - always accepted
291     %sample t2 in MH step
292
293     %% SAMPLE T2 - THE VARIANCE FOR THE RANDOM EFFECTS PRECISION MATRIX
294     b = S+2;
295     pos t2_ = 0; %t2 must be positive
296     while pos t2_ == 0
297         proposed t2_ = proposal sample(current t2,cov mat b { });
298         if proposed t2_ > 0
299             pos t2_ = 1;
300         end
301     end
302     proposed lambda = ...
303         diag(sqrt(proposed t2_))*current R*diag(sqrt(proposed t2_));
304     u = unifrnd(0,1);
305
306     try
307
308         [proposed log likelihood, ] = ...
309             log likelihood(S,N,D,current theta,proposed lambda,...
310                 current delta,current Delta,sigma obs,strain inc,...
311                 sigma vpssc,b,block index);
312
313         proposed log posterior = ...
314             log posterior(S,N,D,block index,...
315                 proposed log likelihood,current theta,...
316                 proposed t2,current R,current delta,...
317                 current Delta,log det Gamma,dubIG,a delta,b delta,...
318                 R not,r not,a t,b t);
319
320         if log(u) < proposed log posterior-current log post
321             current t2_ = proposed t2;
322             current log post = proposed log posterior;
323             acceptance(k,b) = 1;
324         else
325             acceptance(k,b) = 0;
326         end

```

```

327
328     catch
329         sprintf('Error in MH step iteration (%i,%i)',k,b)
330         kk = k+1;
331         sprintf(' n \Continuing to next iteration
332             with ...parameter(%i,%i)...
333             = parameter(%i,%i)',kk,b,k,b)
334     acceptance(k,b) = 0;
335 end
336 %% SAMPLE R - THE CORRELATION FOR THE RANDOM EFFECTS PRECISION MATRIX
337 try
338     sq diff pars = 0;
339     for i = 1:S
340         sq diff pars = sq diff pars + diag(sqrt(current t2))...
341             *(current theta(block index{i})-...
342                 current theta(block index{S+1}))'...
343             *(current theta(block index{i})-...
344                 current theta(block index{S+1}))...
345             *diag(sqrt(current t2));
346     end
347     current R =...
348         wishrnd((R not eye(D) + sq diff pars) eye(D),r not + S);
349     catch
350         disp(['Numerical error in Gibbs step for R at...
351             iteration',num2str(k)])
352     end
353     current Lambda = diag(sqrt(current t2))*...
354         current R*diag(sqrt(current t2));
355
356 %% SAMPLE PSI (delta) - THE ERROR PRECISION
357
358     sq diff = (sigma obs - sigma vpsc - current Delta).^2;
359     current delta = gamrnd(S*N/2 + a delta,...
360         1/(b delta + 1/2*(sum(sum(sq diff)))));
361
362 %% SAMPLE Delta - THE DISCREPANCY
363
364     err pr mat = diag(repmat(current delta,1,N));
365     sample mean = mean((sigma obs - sigma vpsc),2);
366     fc var = (dubIG + N.*err pr mat)\eye(N);
367     fc mean = fc var*(N.*err pr mat)*sample mean;
368     current Delta = mvnrnd(fc mean',fc var)';
369
370
371 %% CALCULATE THE LOG LIKELIHOOD AFTER GIBBS STEPS
372     [proposed log likelihood,vpsc prop] = ...
373         log likelihood(S,N,D,current theta,current Lambda,...
374             current delta,current Delta,sigma obs,strain inc,...
375             sigma vpsc,b,block index);
376

```

```

377 %% CALCULATE THE LOG POSTERIOR AFTER GIBBS STEPS
378     proposed log posterior = ...
379         log posterior(S,N,D,block index,...
380         proposed log likelihood,current theta,...
381         current t2,current R,current delta,...
382         current Delta,log det Gamma,dubIG,a delta,...
383         b delta,R not,r not,a t,b t);
384
385     current log post = proposed log posterior;
386
387     if mod(k,thin) == 0
388         theta chain(k/thin + 1,:) = current theta;
389         t2 chain(:,k/thin+1) = current t2;
390         R chain(:,k/thin+1) = current R;-
391         Lambda chain(:,k/thin + 1) = current Lambda;
392         delta chain(k/thin + 1,1) = current delta;
393         Delta chain(:,k/thin+1) = current Delta;
394         log post(k/thin+1,1) = current log post;
395     end
396
397     temp index = rem(k,covariance check);
398     temp chain(temp index+1,1:block index{S+1}(D)) = current theta;
399     temp chain(temp index+1,block index{S+2}) = current t2';
400
401 %% DRAW FROM POSTERIOR AND POSTERIOR PREDICTIVE DISTRIBUTIONS
402
403     if mod(k,store ppd) == 0
404         try
405             [ppd draw, post draw] = ...
406                 post pred draw(N,S,D,current theta,current Lambda,...
407                 current delta,current Delta,strain inc,sigma obs,...
408                 ppd draw,post draw,k,store ppd,figure path);
409         catch
410             end
411     end
412 end

```

C.2 LOG LIKELIHOOD

```

1  %Copyright (c) 2020, Denielle Ricciardi
2  %All rights reserved.
3
4  function [log_likelihood,prop vpsc] = ...
5      log_likelihood(S,N,D,curr_theta,curr_Lambda,curr_delta,...
6      curr_Delta,sigma_obs,strain_inc,sigma_vpsc,b,block_index)
7

```

```

8 err_pr_mat = diag(repmat(curr_delta,1,N));
9
10 log_like_between = 0;
11 log_like_within = 0;
12
13 prop_vpssc = sigma_vpssc;
14
15 if b ≤ 6
16     par = curr_theta(block_index{b});
17     prop_vpssc(:,b) = VPSC(par(1),par(2),par(3),par(4),strain_inc);
18 end
19
20 Z1 = -D*log(2*pi) + log(det(curr_Lambda));
21 Z2 = -N*log(2*pi) + sum(log(diag(err_pr_mat)));
22
23 for i = 1:S
24     log_like_between = ...
25         log_like_between + (1/2)*Z1 - (1/2)*...
26         (curr_theta(block_index{i}) - ...
27             curr_theta(block_index{S+1})) ...
28         *curr_Lambda*...
29         (curr_theta(block_index{i}) - curr_theta(block_index{S+1}))';
29     log_like_within = ...
30         log_like_within + (1/2)*Z2 - (1/2)*...
31         (sigma_obs(:,i) - prop_vpssc(:,i) - curr_Delta)'...
32         *err_pr_mat*(sigma_obs(:,i) - prop_vpssc(:,i) - curr_Delta);
33 end
34
35 log_likelihood = log_like_between + log_like_within;
36 end

```

C.2 LOG POSTERIOR

```

1 %Copyright (c) 2020, Denielle Ricciardi
2 %All rights reserved.
3
4 function log_posterior = ...
5     log_posterior(S,N,D,block_index,log_like,...
6     curr_theta,curr_t2,curr_R,curr_delta,...
7     curr_Delta,log_det_Gamma,inv_Gamma,a_delta,...
8     b_delta,R_not,r_not,a_t,b_t)
9
10 hyp1 = [.98,.98,1800,600];
11 hyp2 = [.014,.014,1E6,1E6];
12 overall_log_priors = ...
13     {@log-gampdf,@log-gampdf,@log-normal,@log-normal};

```

```

14
15 overall_par = curr_theta(block_index{S+1});
16 overall_log_prior = 0;
17 for dd = 1:D
18     overall_log_prior = ...
19         overall_log_prior + ...
20         overall_log_priors{dd}(overall_par(dd),hyp1(dd),hyp2(dd));
21 end
22
23 Z = -N*log(2*pi) - log_det_Gamma;
24 Delta_log_prior = (1/2)*Z - (1/2)*curr_Delta'*inv_Gamma*curr_Delta;
25
26 t2_log_prior = 0;
27 for d = 1:D
28     t2_log_prior = t2_log_prior + log_gampdf(curr_t2(d),a_t,b_t);
29 end
30
31 R_log_prior = log_wishart_pdf(R_not,r_not,curr_R,D);
32 delta_log_prior = log_gampdf(curr_delta(end),a_delta,b_delta);
33
34 log_posterior = double(log_like + overall_log_prior + ...
35     Delta_log_prior + t2_log_prior + R_log_prior + delta_log_prior);
36 end

```

C.4 POSTERIOR AND POSTERIOR PREDICTIVE DRAWS

```

1 %Copyright (c) 2020, Denielle Ricciardi
2 %All rights reserved.
3
4 function [new_traj, post_draw] = ...
5     post_pred_draw(N,S,D,curr_theta,curr_Lambda,...
6     curr_delta,curr_Delta,strain_inc,FFT_stress,...
7     new_traj,post_draw,k,store,figure_path)
8
9 %% Sample from Posterior Predictive Distribution
10
11 %sample random effect from likelihood
12 Kosher = 0;
13 while Kosher ==0
14     rand_eff = mvnrnd(curr_theta(S*D + 1:end),curr_Lambda\eye(D));
15     if rand_eff(4) >= rand_eff(3)
16         continue
17     end
18     if sum(rand_eff <= 0) /= 0
19         continue
20     end

```

```

21         Koshier = 1;
22     end
23
24     %sample new experiment from likelihood
25     rand eff vpsc = VPSC(rand eff(1),rand eff(2),...
26         rand eff(3),rand eff(4),strain inc);
27     err pr mat = diag(repmat(curr delta,1,N));
28     new traj(:,k/store) = mvnrnd(rand eff vpsc + ...
29         curr Delta,err pr mat\eye(N));
30
31     figure
32     set(gcf,'Visible','off')
33     h1 = plot(strain inc,FFT stress(:,:),'r--','Linewidth',.5);
34     hold on
35
36     for s = 1:k/store
37         xflip = [strain inc' fliplr(strain inc)];
38         yflip = [new traj(:,s)' fliplr(new traj(:,s))];
39         p = patch(xflip,yflip,'b','EdgeAlpha',.2,...
40             'FaceColor','none','Linesmoothing','on',...
41             'Linewidth',1);
42     end
43
44     xlim([-0.01 .091])
45     ylim([0 200])
46     title('Draws from Posterior Predictive Distribution')
47     xlabel('strain')
48     ylabel('stress (MPa)')
49
50     %% Sample from Posterior Distribution
51     figure
52     set(gcf,'Visible','off')
53     h1 = plot(strain inc,FFT stress(:,:),'r--','Linewidth',.5);
54     hold on
55
56     par = curr theta(S*D+1:end);
57     pdraw = VPSC(par(1),par(2),par(3),par(4),strain inc);
58
59     post draw(:,k/store) = pdraw + curr Delta;
60     for s = 1:k/store
61         xflip = [strain inc' fliplr(strain inc)];
62         yflip = [post draw(:,s)' fliplr(post draw(:,s))];
63         p = patch(xflip,yflip,'b','EdgeAlpha',.2,...
64             'FaceColor','none','Linesmoothing','on','Linewidth',1); ...
65             %edge alpha originally .1
66     end
67     xlim([-0.01 .091])
68     ylim([0 200])
69     title('Draws from Posterior Distribution')
70     xlabel('strain')
71     ylabel('stress (MPa)')

```

C.5 ADAPTIVE PROPOSAL COVARIANCE

```

1 function cov_mat_new = full_cov_mat_adj(par,D,cov_mat_orig)
2     def_by = 1;
3
4     chaincorr = corr(par);
5     chaincorr(isnan(chaincorr)) = 0;
6
7     new_stepcor_chain = zeros(D,D);
8     temp_stepvar_chain = zeros(D,D);
9
10    for rr = 1:D
11        for cc = 1:D
12            if (rr==cc)
13                new_stepcor_chain(rr,cc) = 1;
14            elseif (abs(chaincorr(rr,cc)) < 0.6)
15                new_stepcor_chain(rr,cc) = 0;
16            else
17                new_stepcor_chain(rr,cc) = sign(chaincorr(rr,cc)) ...
18                    * (abs(chaincorr(rr,cc)) - 0.4);
19            end
20        end
21    end
22
23    pvars = diag(cov_mat_orig)/def_by;
24
25    for rr = 1:D
26        for cc = 1:D
27            temp_stepvar_chain(rr,cc) = new_stepcor_chain(rr,cc) ...
28                * sqrt(pvars(rr)*pvars(cc));
29        end
30    end
31
32    tempvar = diag(temp_stepvar_chain);
33
34    for dd = 1:D
35        if tempvar(dd) > 0
36            temp_stepvar_chain(dd,dd) = cov_mat_orig(dd,dd);
37        end
38    end
39
40    cov_mat_new = temp_stepvar_chain;
41 end

```

LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

AFRL	Air Force Research Laboratory
CP	Crystal Plasticity
CPFE	Crystal Plasticity Finite Elements
CRSS	Critical Resolved Shear Stress
EVP-FFT	Elasto-ViscoPlastic Fast Fourier Transform
GP	Gaussian Process
HEM	Homogeneous Effective Medium
HPI	Highest Posterior Intervals
ICME	Integrated Computational Materials Engineering
MAP	Maximum A Posteriori
MCMC	Markov Chain Monte Carlo
MH	Metropolis-Hastings
PDF	Probability Density Function
PUQ	PRISM Uncertainty Quantification
RXCA	Materials State Awareness Branch, Structural Materials Division, Materials and Manufacturing Directorate
SR	Statistically Representative
SGTE	Scientific Group Thermodata Europe
UP	Uncertainty Propagation
UM	Uncertainty Mitigation
UQ	Uncertainty Quantification
V&V	Verification and Validation
VPSC	Visco-Plastic Self-Consistent
WPAFB	Wright Patterson Air Force Base