A Bayesian Approach to Designing Microstructures and Processing Pathways for Tailored Material Properties

Adam P. Generale[1] **Conlain Kelly**[2] **Grayson H. Harrington**[2] **Andreas E. Robertson**[1]
Michael O. Buzzy[2] **Surya R. Kalidindi**[1,2†]

[1] George W. Woodruff School of Mechanical Engineering
[2] School of Computational Science and Engineering
† *Corresponding author*
Georgia Institute of Technology, Atlanta GA 30332
{agenerale3,ckelly84,grayson.harrington,arobertson38,mbuzzy3}@gatech.edu
surya.kalidindi@me.gatech.com

Abstract

Inverse problems are central to material design. While numerous studies have focused on designing microstructures by inverting structure–property linkages for various material systems, such efforts stop short of providing realizable paths to manufacture such structures. Accomplishing the dual task of designing a microstructure *and* a feasible manufacturing pathway to achieve a target property requires inverting the complete process–structure–property linkage. However, this inversion is complicated by a variety of challenges such as inherent microstructure stochasticity, high-dimensionality, and ill-conditioning of the inversion. In this work, we propose a Bayesian framework leveraging a lightweight flow-based generative approach for the stochastic inversion of the complete process–structure–property linkage. This inversion identifies a solution distribution in the processing parameter space; utilizing these processing conditions realizes materials with the target property sets. Our modular framework readily incorporates the output of stochastic forward models as conditioning variables for a flow-based generative model, thereby learning the complete joint distribution over processing parameters and properties. We demonstrate its application to the multi-objective task of designing processing routes of heterogeneous materials given target sets of bulk elastic moduli and thermal conductivities.

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

1 Introduction

The problems of material discovery, development, and deployment have historically been posed in the form of *inverse problems* [1] – mathematical challenges which repeatedly appear across almost every discipline [2–9]. In the context of materials, the primary challenges involve mapping between the spaces of material processing parameters Ω_1, microstructures Ω_2, and resulting functional properties Ω_3. This introduces the notion of process–structure (PS) and structure–property forward linkages, which we denote as $\mathcal{G} : \Omega_1 \to \Omega_2$ and $\mathcal{F} : \Omega_2 \to \Omega_3$. However, the underlying goal of inverting the relationship between processing and property, remains an open problem. This composite relation $\mathcal{H} = \mathcal{F} \circ \mathcal{G} : \Omega_1 \to \Omega_3$ is commonly referred to as a process–structure–property (PSP) linkage.

It is impossible to disentangle material microstructures from how they deform, and consequently their bulk properties and performance. Unfortunately, consideration of microstructures introduces significant complexities to materials-specific inverse problems [1, 10]. First, material microstructures are complex hierarchical stochastic systems spanning a wide range of length-scales [11, 12]. Across these length-scales, the evolution of local material states during processing steps, as well as their bulk response to external stimuli, are highly nonlinear and stochastic in nature [13–20]. Second, microstructure representation is inherently irregular and high-dimensional [21], such that any PS, $\mathcal{G}(\cdot)$, or SP linkage, $\mathcal{F}(\cdot)$, would necessarily need to be capable of modeling or approximating this behavior across high-dimensional domains. Taken as a whole, these factors result in an ill-posed and unstable composite inversion problem. In the presence of such stochastic forward models, the problem is more formally known as a *stochastic inverse problem* [22].

As a direct result of these challenges, current works have primarily focused on inverting individual SP linkages [23–30]. Even these works have achieved significant advancements over traditional approaches [15, 31, 32] through their use of distributional deep learning [33–38]. By learning a mapping of individual samples from the complex, irregular microstructure space to a topologically-well-structured, continuous latent space, traversal of this latent space can lead to the identification of unique and undiscovered materials. However, these current approaches rarely account for microstructure stochasticity. Further, learned latent spaces lack any particular domain knowledge or physics-supported topology (e.g., they do not naturally account for underlying symmetries) [23, 24, 26–30, 39]. Altogether, these limitations add unnecessary complexity into the generative model in order to distinguish between statistically-identical microstructures. Most critically, these approaches only partially solve the problem at hand; while knowledge of property-conditioned microstructures is certainly valuable, it does not provide any processing route for manufacturing such target microstructures of interest. This secondary inversion is usually handled via a domain-knowledge-dependent iterative search [14, 15].

In this work, we leverage several learning methods to efficiently solve the full PSP inverse problem. Specifically, the inverse problem is posed and solved in a Bayesian setting, which naturally accounts for stochasticity in the composite forward model and degeneracy in the inversion. The solution is generated using amortized variational inference [40]. In order to disentangle the microstructure space Ω_2, we employ a statistical latent representation of microstructure, namely, 2-point spatial correlations [41, 42]. Stochastic forward models are then constructed linking this representation to both processing parameters ($\Omega_1 \to \Omega_2$) and effective properties ($\Omega_2 \to \Omega_3$). The inverse linkage ($\Omega_1 \leftarrow \Omega_3$) itself is modeled using a conditional continuous normalizing flow, which is trained with its conditioning variable taken as the property output of the chained stochastic forward models. Once trained, this model can sample from the approximate posterior distribution of property-conditioned processing parameters $p(\boldsymbol{x}|\boldsymbol{y})$, where $\boldsymbol{y}$ is a set of target properties and $\boldsymbol{x}$ is a set of processing or manufacturing parameters. We conclude by applying the framework to the design of heterogeneous microstructures produced by phase-field simulation [43] and provide statistical analysis of our method's overall performance and reliability.

2 Background

Phase-Field Modeling: Phase-field simulations are commonly applied to model a number of manufacturing processes involving evolving interfaces (such as solidification/melting, spinodal decomposition, grain growth, recrystallization, and crack propagation [44, 45]). In particular, the Cahn-Hilliard equation is frequently used to describe spinodal decomposition, a spontaneous

thermodynamic-instability-induced phase separation [46]; this partial differential equation models a diffusion-driven process with a diffusivity constant $\mathcal{D}$, driving the evolution of composition variations, c, over characteristic length scales dictated by a gradient energy coefficient, γ. In our problem the spatially-dependent composition will take the role of material microstructure.

$$\frac{\delta c}{\delta t} = \mathcal{D}\nabla^2\left(c^3 - c - \gamma\nabla^2 c\right) \tag{1}$$

Gaussian Processes: Throughout this work, the forward mappings (PS and SP linkages) are assumed to be stochastic, thus requiring forward uncertainty propagation. We approach this task through the use of Gaussian processes ($\mathcal{GP}s$), an expressive and flexible construct, defining a distribution over a family of functions by describing their functional output covariance across the input domain. This covariance necessarily dictates the resulting function characteristics, and is constructed with a kernel function $k(\cdot,\cdot')$. Provided observations $\{\boldsymbol{x}_n, \boldsymbol{y}_n\}_{n=1}^N$ from an underlying generating process, a conditioned distribution over the space of functions $\{\boldsymbol{x} \rightarrow \boldsymbol{y}\}$ can be established, estimating the trend and the uncertainty of the data. For a more in depth discussion we defer to [47, 48].

2-Point Spatial Correlations: This work uses 2-point spatial correlations [13, 14, 41] to construct descriptive microstructure features. This representation builds on the idea that the microstructure itself is a stochastic function, where individual observed microstructure instances are samples from the governing stochastic microstructure function [41,42]. This theoretical treatment accounts for the inherent stochasticity displayed by material microstructures, and allows for the underlying stochastic function to be linked to homogenized properties. It also provides a convenient mechanism to account for underlying symmetries (such as translation-equivarinace and periodicity).

In practice, 2-point spatial correlations can be computed as a convolution of the sampled discrete microstructure function m_s^α, where α is the material local state at voxel s. The resulting 2-point spatial correlations are then defined by the operation

$$f_r^{\alpha\beta} = \frac{1}{S}\sum_{s=1}^{S} m_s^\alpha m_{s+r}^\beta \tag{2}$$

where S defines the set of all voxels in the microstructural domain. These represent lower-order terms in a moment expansion of the true microstructure random process; for a number of materials systems, this term is dominant and captures most of the variation in bulk material properties. Dimensionality reduction techniques can then be effectively applied to the ensemble of 2-point spatial correlations to provide robust, information-dense features for the construction of PS and SP linkages [16, 49–55].

Conditional Flow Matching: Continuous normalizing flows (CNFs) [37] are an expressive and efficient method for density estimation which construct an invertible mapping between the standard normal distribution and a given data distribution. The mapping itself is specified via the dynamics of a neural ordinary differential equation (ODE) [37]. While these methods have historically exhibited difficulties during training [56, 57], *Conditional Flow Matching* (CFM) [38] has emerged as a powerful *simulation-free* technique for stabilizing the training of CNFs. This objective enables training of the CNF *without* requiring internal calls to an ODE solver during the training phase, instead matching the vector field of the conditional transport map.

3 Framework

In this section we propose a Bayesian Inversion [58] framework for solving the process-property stochastic inverse problem. Figure 1 visually summarizes the overarching strategy. To begin, we define notation specific to this particular problem. Inputs of the forward model are denote by the vector of processing parameters $\boldsymbol{\theta} \in \mathbb{R}^D$, the target effective property vector is denoted by $\boldsymbol{k} \in \mathbb{R}^P$,

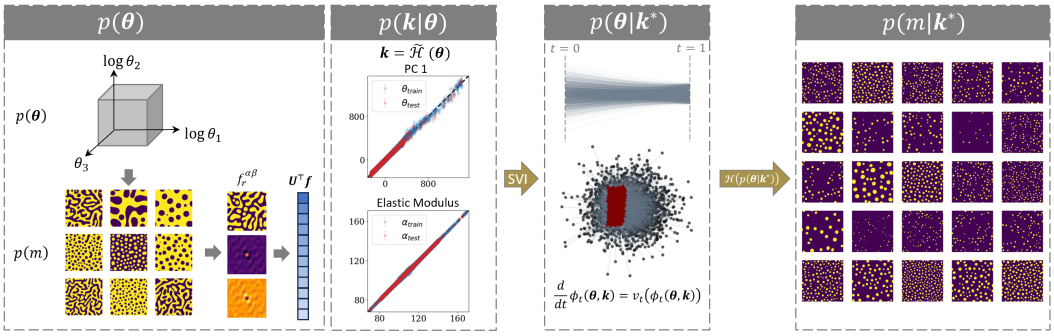

Figure 1. Visual representation of proposed framework. The prior distribution in the first two dimensions is invertibly transformed to the uniform distribution, which can then be sampled to produce unique microstructure instantiations. This prior over processing parameters results in an equivalent prior in microstructure space when phase-field simulations are performed. This latent microstructure prior, alongside processing through quantification with 2-point spatial correlations and PCA embedding is also shown. The likelihood is then defined as a composition of probabilistic PS and SP linkages, which is subsequently used to perform amortized stochastic variational inference (SVI). Lastly, the image of posterior samples under the (true) PS forward model is displayed, demonstrating microstructure specificity to match target properties. $\tilde{\mathcal{H}}$ denotes the composite surrogate forward model, $\boldsymbol{k}$ the property set of interest, $\boldsymbol{\alpha}$ the statistical low-dimensional PC representation of microstructure, and m microstructure.

and intermediate (2D, in this work) microstructures are written $m \in \mathbb{R}^{L\times L}$. Given an initial dataset of processing parameter, microstructure, and property triples, $\{\boldsymbol{\theta}_n, m_n, \boldsymbol{k}_n\}_{n=1}^N$, we create tractable estimations for the components of Bayes' theorem necessary for computing the desired posterior (e.g., $p(\boldsymbol{\theta})$, $p(\boldsymbol{k}|\boldsymbol{\theta})$)[1].

We note that a naïve approach to identifying the mappings $\mathcal{G} : \boldsymbol{\theta} \to m$ and $\mathcal{F} : m \to \boldsymbol{k}$ would necessarily traverse the high-dimensional stochastic microstructure space. This challenge would only be amplified during the inversion. We address these issues directly by quantifying microstructures through 2-point spatial correlations, and then embedding them into a lower-dimensional manifold through Principal Component Analysis (PCA) [53, 59–61], retaining R principal components (PC) scores (denoted $\boldsymbol{\alpha}$) as descriptors. Importantly, this physics-informed R-dimensional representation has effectively demonstrated a robust ability of stratifying associated properties in the first handful of terms [49–55], enabling us to move through a R-dimensional microstructure space as higher-order terms result in minimal shifts in resulting properties. This distillation of global microstructure information greatly simplifies the overarching inversion, as identifying a distribution over the exceedingly high-dimensional microstructure space is wholly unnecessary. Rather, we can traverse this compact representation of the microstructure random process. Briefly, we note that this approach could potentially break down when the target properties are the result of localized processes, such as fracture and fatigue resistance [62–65].

In summary, the proposed modular framework utilizes surrogate forward models $\tilde{\mathcal{G}}$ and $\tilde{\mathcal{F}}$ to train a flow-based generative model to estimate the posterior mapping between desired properties and processing parameters. This cooperative reformulation avoids the problem of inverting this mapping directly from just data using solely a flow-based generative model, and instead provides a platform for which to construct smaller forward models as a continuous source of training data for the CNF.

4 Experiments & Discussion

In order to validate the proposed framework, we apply it to two different materials systems undergoing spinodal decomposition, with in depth analysis performed for two distinct test cases per system, and subsequently, a more general validation procedure applied to the entire property space. These two cases consist of (processing parameters, property) pairs denoted as $\{\boldsymbol{\theta}^*, \boldsymbol{k}^*\}$,

[1] A reference summary of notation used throughout this work is presented in Appendix D.

Table 1. Summary of constituent property assignment for low- and high-contrast ratio cases, displaying the isotropic elastic modulus (E), Poisson's ratio (ν), and thermal conductivity (κ) of the constituents considered.

	Al-Si			High-Contrast		
	E (GPa)	ν	κ (W/mK)	E (GPa)	ν	κ (W/mK)
Constituent 0	70	0.33	247	1	0.3	100
Constituent 1	169	0.28	400	100	0.3	1

which we selected from a dataset of spinodal decomposition phase-field simulations. The phase-field simulations take a 3-dimensional processing parameter input $\boldsymbol{\theta}$, dictating disparate heterogeneous microstructure evolution pathways. The property set $\boldsymbol{k}$ considered includes both effective anisotropic elastic moduli and thermal conductivity.

The two materials systems considered were chosen to evaluate model performance with varying contrast between phases, as well as its ability to balance competing property targets. The low-contrast case corresponds to the Al-Si alloy material system, which is physically realizable. Conversely, the high-contrast case is a fictitious 2-element alloy in which the elastic and thermal properties trade off (and vary heavily) between phases. It is important to consider multiple contrast ratios because it has been well documented that PSP linkages become harder to construct with higher contrast ratios [66]. A summary of the constituent material parameters can be found in Table 1. These stress tests were selected to best scrutinize the proposed approach across sparse regions of the dataset and increasingly-complex material responses.

The dataset contains 10,000 two-phase microstructures of size 256×256 voxels, each associated with processing parameters θ_1 and θ_2, sampled according to the log-uniform $\log(\theta) \sim \mathcal{U}(\log(0.1), \log(100))$, and $\theta_3 \sim \mathcal{U}(-0.7, 0.7)$. These processing parameters correspond to the mobility parameters of the two constituents (θ_1, θ_2), and initial relative concentrations (θ_3). Homogenized properties were extracted for each microstructure in the dataset through performing finite element simulations with Abaqus/Standard [67]. Figure 2a visually summarizes the resulting property set for these microstructures with Al-Si constituent property assignment and Figure 2b with the high-contrast assignment. Figure 2c displays the microstructure ensemble in the first 4 PC dimensions. The selected test cases are highlighted in each projection as p_1 and p_2.

Figure 2c illustrates projections of the microstructure ensemble in the first 4 PCs. Each microstructure is statistically represented by a subset of its 2-point spatial correlations: we computed $\{f_r^{00}, f_r^{01}\}$ using Eq. (2), with the material local state indexed by $\{0, 1\}$. While this set of spatial correlations provides a robust statistical description of internal microstructure arrangement, it also doubles the dimensionality relative to the initial microstructure. As such, Principal Component Analysis (PCA) was performed to reduce the dataset's dimensionality, using a concatenated flattened representation of $\{f_r^{00}, f_r^{01}\}$. Scaling factors were applied to the set as $\{f_r^{00}, 26.389 f_r^{01}\}$ to balance contributions between auto- and cross-correlations so that neither set dominated the transformation. Overall, Figure 2 highlights the motivation for selecting these particular test cases, with each selected to balance coverage of the property space while simultaneously identifying sparse regions in PC space.

4.1 Al-Si Alloy

After training, samples from the approximate posterior $p(\boldsymbol{\theta}|\boldsymbol{k}^*)$ were drawn from the CNF model for both of the test cases, displayed in Figure 3. The marginal distributions in the log-space of the processing parameters can be seen in Figure 3, where the processing parameter set $\boldsymbol{\theta}^*$ associated with the conditioning property set $\boldsymbol{k}^*$ is demarcated with a vertical line. Inspection of the posteriors highlights that in the first case study, larger deviations in θ_1 and θ_2 can be tolerated, but extreme specificity is required in θ_3, which, notably, was found to be very near to zero. This observation aligns with the fundamental dynamics of the Cahn-Hilliard phase-field model. When the order parameter θ_3 is near zero, signifying an almost equal fraction of both phases in the initial mixture, the system experiences a heightened driving force for phase separation rendering it less sensitive to mobility parameter discrepancies in the early stages of the microstructure evolution. As a result,

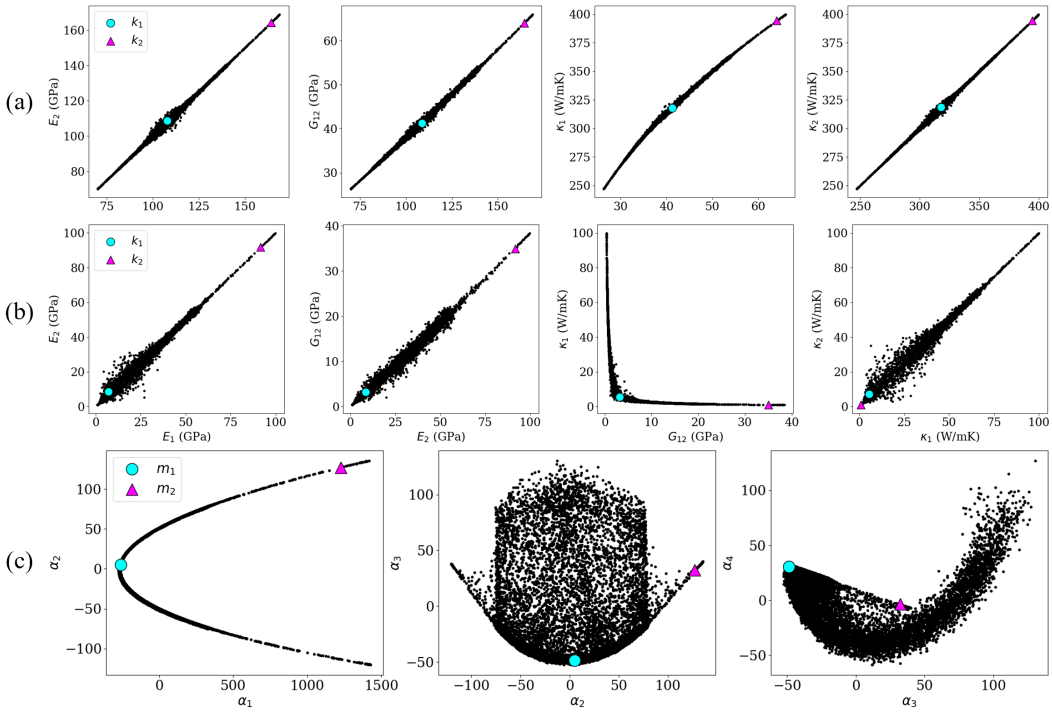

Figure 2. Experimental dataset considered in this work with select microstructures identified as test cases. (a) Property set of the low-contrast Al-Si microstructure ensemble, (b) property set of the high-contrast microstructure ensemble, and (c) low-dimensional microstructure representation in the first 4 PC dimensions.

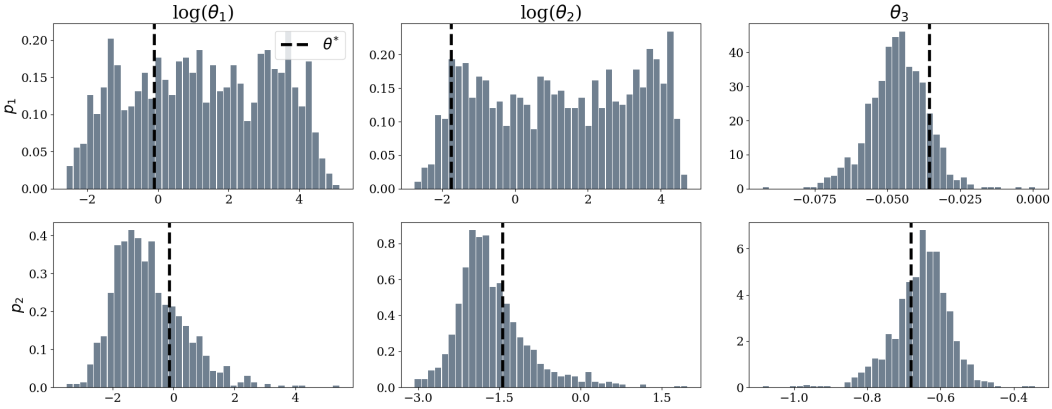

Figure 3. Marginal posterior distributions for the two test cases p_1 and p_2 with the low-contrast constituent property assignment. The first two dimensions of $\boldsymbol{\theta}$ are displayed in the log-space. Associated processing parameter sets $\boldsymbol{\theta}^*$ to the target property set $\boldsymbol{k}^*$ are displayed as a black dashed line.

even significant variations in the mobility parameters, θ_1 and θ_2, have a diminished impact on the microstructure evolution. In contrast, the second case study posterior demonstrates significantly more sharpening across the processing parameter domain, indicating that in sparser regions of the microstructure and property spaces, a greater degree of specificity across all processing parameters is necessary to meet the target property set. In each of the marginal posterior plots, the framework was able to recover the latent $\boldsymbol{\theta}^*$ associated with the target property set.

Next, we interrogate how well the posterior predictions match the desired target property set. To address this we draw 200 valid samples from the posterior $p(\boldsymbol{\theta}|\boldsymbol{k}^*)$), and apply both the surrogate

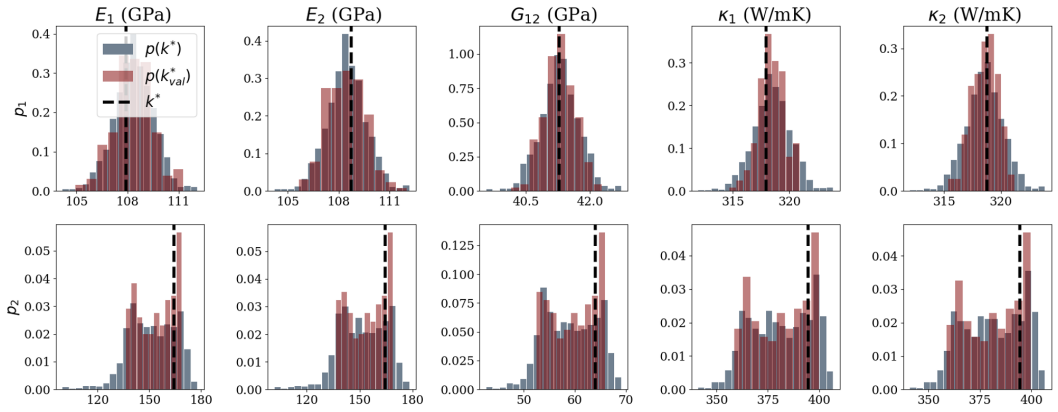

Figure 4. Marginal distributions of $\mathcal{H}(p(\boldsymbol{\theta}|\boldsymbol{k}^*))$ (ground truth forward model) for the two low-contrast test cases p_1 and p_2 shown alongside the validation set of 200 processing parameters sampled from the posteriors. The validation set was evaluated through phase-field and finite element simulations.

and "ground truth" composite forward models ($\tilde{\mathcal{H}}, \mathcal{H}$ respectively) on each sample. Posterior samples outside of the validity bounds of the phase-field model were rejected. The results of this can be seen in Figure 4, where the property target falls well within the predicted distribution (for both models). Additionally, the distributions of resultant properties match relatively well between surrogate and true forward models.

4.2 High-Contrast Alloy

The low-contrast Al-Si test cases demonstrated that the framework is capable of adequately identifying the processing parameter sets necessary to meet various target property sets. Due to the higher contrast of the microstructures' constituent properties and increasing nonlinearity of their responses in this case study, we expect less variability in the distributions over processing parameters and latent microstructure statistics. Such intuition is reflected in the resulting posteriors $p(\boldsymbol{\theta}|\boldsymbol{k}^*)$, displayed in Figure 5. p_1 demonstrates a departure from the low-contrast log-uniform, with a preference towards lower mobility values (θ_1 and θ_2), while the posterior p_2 remained roughly equivalent to the prior case study due to its extremal location in the dataset.

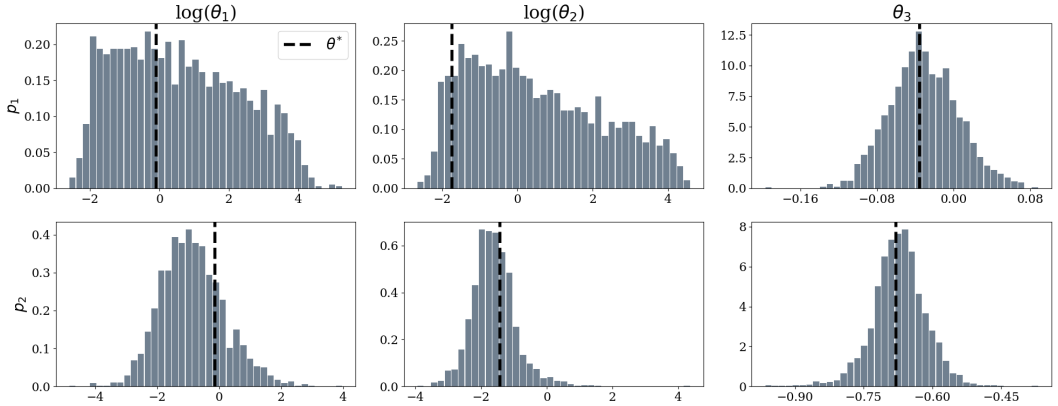

Figure 5. Marginal posterior distributions for the two test cases p_1 and p_2 with the high-contrast constituent property assignment. The first two dimensions of $\boldsymbol{\theta}$ are displayed in the log-space. Associated processing parameter sets $\boldsymbol{\theta}^*$ to the target property set $\boldsymbol{k}^*$ are displayed as a black dashed line.

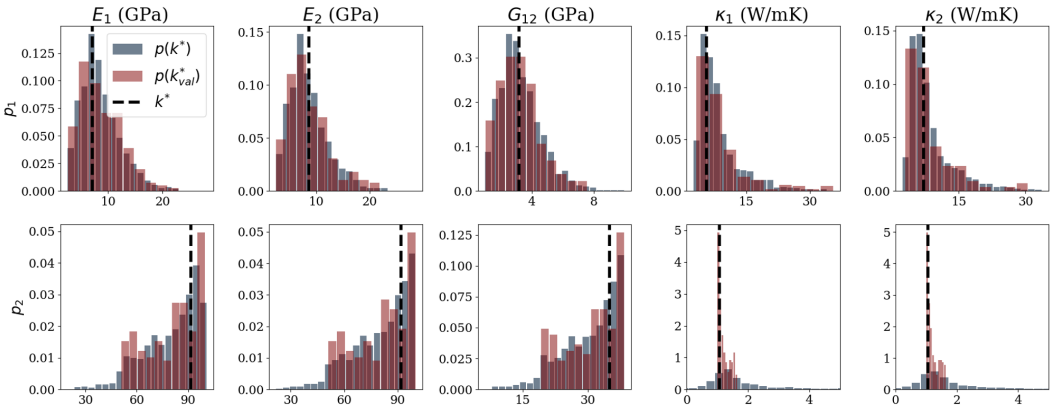

Figure 6. Marginal distributions of $\mathcal{H}(p(\boldsymbol{\theta}|\boldsymbol{k}^*))$ for the two high-contrast test cases p_1 and p_2 shown alongside the validation set of 200 processing parameters sampled from the posteriors. The validation set was evaluated through phase-field and finite element simulations.

In a similar fashion to the prior case study, we turn towards inspecting the resulting property set predictions of the processing parameter posterior. We apply the same methodology used for Figure 4 for the high-contrast case and present the results in Figure 6. Again, the target property vector is recovered by both the surrogate and exact forward models. Interestingly, the surrogate model over-estimates the variance in predictions k_1 and k_2 compared to the true model. We believe this is caused by a degradation in performance of the SP linkage $\tilde{\mathcal{F}}$ for thermal conductivities near the homogeneous microstructure lower-property limit, as seen in Appendix A. Importantly, even with this degradation in forward property prediction, the framework is capable of identifying promising processing parameter sets in validation.

4.3 Validation Across Property Space

While utilizing these two selected test cases provided the opportunity for validation through the exact forward process, it is limited by the computational cost of evaluating this process. Consequently, the limitation to certain target property sets precludes a complete understanding of the frameworks performance across the entire property space. In order to further explore framework performance, we now turn towards identifying posterior solutions for every property set in the available dataset. In this validation, we restrict our consideration to property values in the hold-out testing dataset from training the forward models. 128 samples are drawn from $p(\boldsymbol{\theta}|\boldsymbol{k}^*)$ for every property value in the dataset, and then each passed through the learned sparse variational multi-output Gaussian process forward models $\tilde{\mathcal{H}}(p(\boldsymbol{\theta}|\boldsymbol{k}^*))$, such that for each target property set there now exists a distribution of forward propagated samples from the corresponding solution posterior. This is practically useful as it provides a method for directly and rapidly comparing the capabilities of the framework, while providing a methodology for inspecting uncertainty in the identified solutions. Figure 7 displays the results of this evaluation for both the low- and high-contrast applications where Figure 7(a) and Figure 7(b) show parity plots colored by the estimated density, highlighting the accurate predictions across the majority of the property space, particularly in the most dense regions. The shape of the estimated cumulative distribution functions for each are overlaid in Figure 7(c) only reinforcing this point. Degraded performance can be observed as we move towards lower-density regions of the property space. Quite naturally, we also correctly observe an increase in uncertainty in such extremal property sets – an early indication uncertainty in the framework is well calibrated.

4.4 Calibration

While the previous case studies indicate that our model captures the correct overall trends, the question remains of how well the uncertainty in the estimated distribution matches that of the true distribution; that is: is our model well-calibrated? The validity of the identified joint distributions $p(\boldsymbol{\theta}, \boldsymbol{k}) = p(\boldsymbol{k}|\boldsymbol{\theta})p(\boldsymbol{\theta})$ is explored through recent work by Säilynoja *et al.* [68]. This simulation-based method evaluates the quality of a Bayesian analysis across the entire joint $p(\boldsymbol{\theta}, \boldsymbol{k})$, sampling

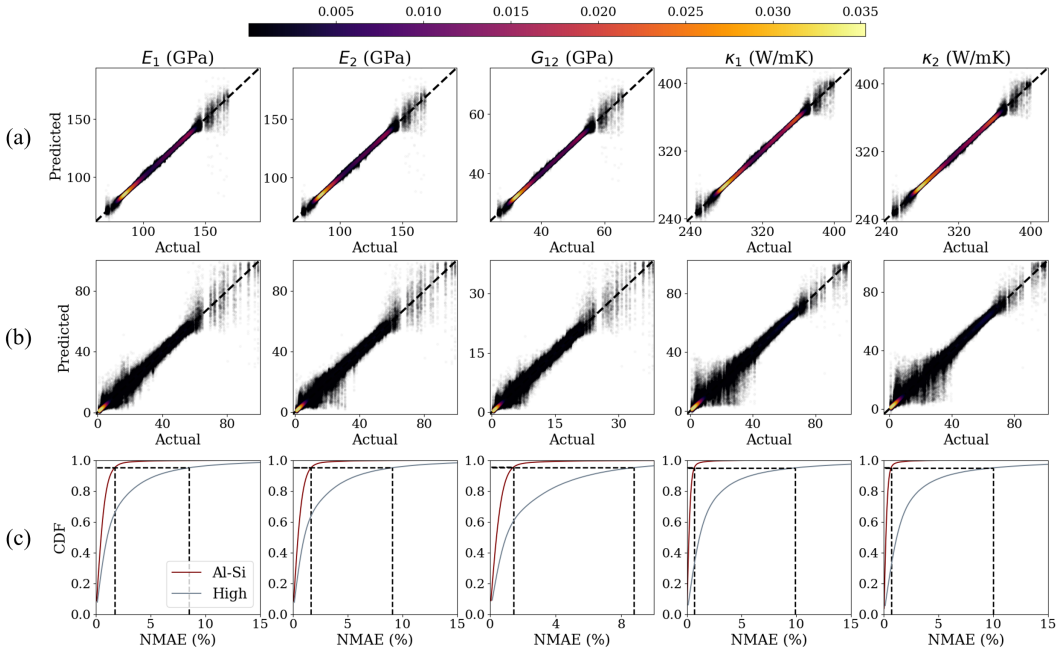

Figure 7. Forward mapping of inverse solutions for all property sets in the available dataset. (a) Parity plot of all target property sets against 128 points propagated from the posterior solution, (b) estimated density of the parity plot, (c) estimated cumulative distribution of Normalized Mean Absolute Error (NMAE) over all propagated property sets with the 95$^{\text{th}}$ percentile highlighted.

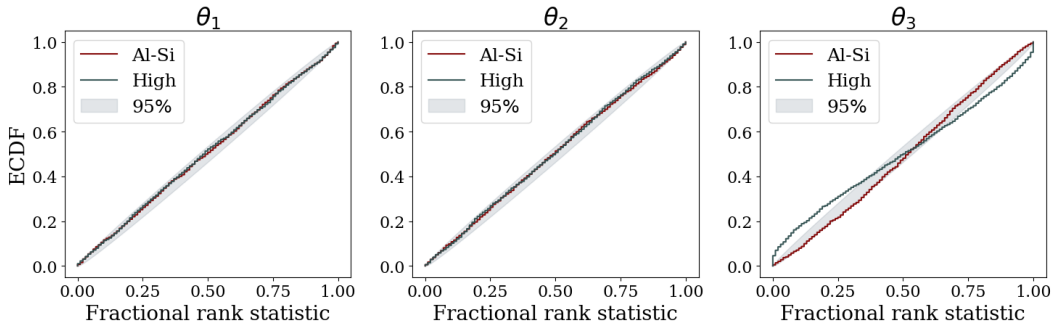

Figure 8. Empirical cumulative distribution functions to asses uniformity of the samples for both low- and high-contrast cases.

a ground truth from the prior $\tilde{\boldsymbol{\theta}} \sim p(\boldsymbol{\theta})$ and data from the likelihood $\tilde{\boldsymbol{k}} \sim p(\boldsymbol{k}|\boldsymbol{\theta})$. Evaluations with these simulated pairs can then be used in evaluating the discrepancy between the data averaged posterior and the prior through the following expression.

$$p(\boldsymbol{\theta}) = \int p(\boldsymbol{\theta}|\tilde{\boldsymbol{k}})p(\tilde{\boldsymbol{k}}|\tilde{\boldsymbol{\theta}})p(\tilde{\boldsymbol{\theta}})d\tilde{\boldsymbol{\theta}}d\tilde{\boldsymbol{k}}. \tag{3}$$

The empirical cumulative distribution function (ECDF) graphical test for uniformity involves the probability integral transform in comparing the data averaged posterior to the prior, providing an intuitive methodology for evaluating the quality of a Bayesian analysis. A linear result is desirable. The results of this procedure for both joints from the low- and high-contrast cases can be seen in Figure 8, as well as 95% confidence intervals (based on a binomial distribution). Briefly, this figure indicates that the model is well-calibrated in the low-contrast case, with discrepancies in the fractional rank statistic for θ_3 in the high-contrast posterior.

5 Conclusions

We present a framework for addressing inverse stochastic materials design problems. The framework provides a tool set for the designer to identify *novel* structures and identify distributions of processing routes to *realize* such structures. We exhibit demonstrations of its utility on increasing challenging test cases, and demonstrate that it provides well calibrated posterior predictions of processing parameters. Due to its stochastic nature, this framework also provides a mechanism for Bayesian model calibration, selection, and uncertainty quantification.

6 Author Contributions

AG conceived the project and designed the models with assistance from CK, AR, and MB. GH generated the spinodal decomposition data, and AG the finite-element results. SK ensured funding support. AG wrote the initial manuscript, CK revised, and all authors edited the final version.

7 Acknowledgments

This research acknowledges funding by various sources. Adam P. Generale: Pratt & Whitney; Conlain Kelly: NSF 2027105, NSF Graduate Research Fellowship DGE-1650044, and ONR N00014-18-1-2879; Grayson Harrington: NSF Graduate Research Fellowship DGE-2039655 and Sandia LDRD Contract Agreement #2177713; Andreas E. Robertson: ONR N00014-18-1-2879; M. Buzzy: NSF DMREF-2119640. Surya R. Kalidindi acknowledges funding from all previously listed funding sources. This work utilized computing resources and services provided by the Partnership for an Advanced Computing Environment (PACE) and the Hive computing cluster at the Georgia Institute of Technology, Atlanta, Georgia, USA. A.E. Robertson also acknowledges the continuing support of the Jack Kent Cooke Foundation, and A.P. Generale that of the Alfred P. Sloan Foundation. Finally, the authors wish to thank Rafael Orozco for invaluable conversations, particularly regarding model calibration.

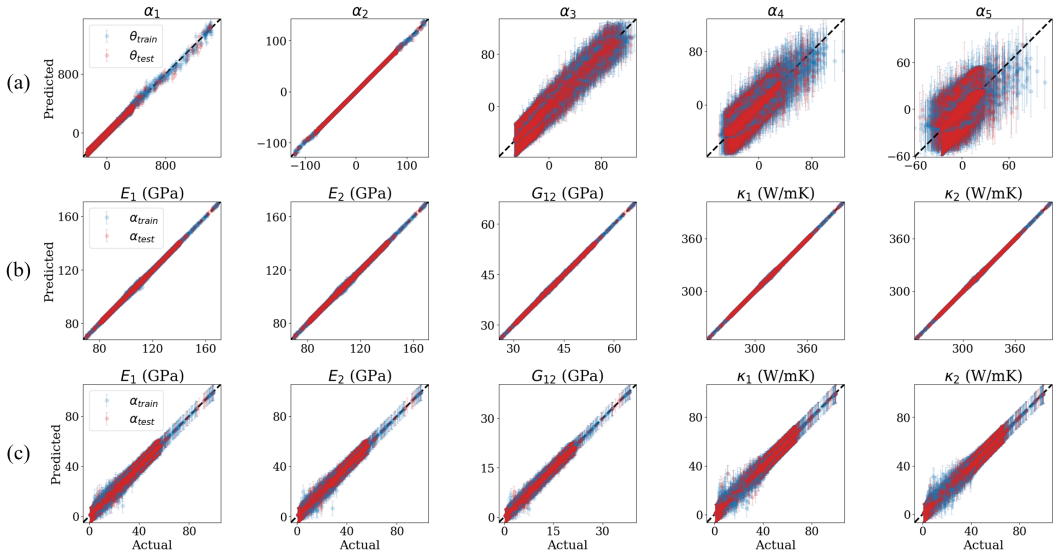

Figure A.1. Parity plot of SV-MOGP forward models performance on the available microstructure ensemble. Predicted $\pm 2.0\sigma$ confidence intervals are displayed alongside mean predictions. (a) PS linkage for first 5 PC scores, (b) SP linkage for Al-Si case, and (c) SP linkage for high-contrast ratio case.

A Appendix: Forward Models

Process-structure (PS) and structure-property (SP) linkages for the Al-Si case were constructed with a sparse variational multi-output Gaussian process (SV-MOGP) [69, 70], where the multi-output covariance matrix was constructed with the Linear Model of Coregionalization (LMC) [71, 72] and a spectral mixture kernel [73]. The intermediate microstructure representation consisted of 5 PC scores due to signal decay and limitations of the PS linkage to accurately predict past this point. The PS linkage was constructed with 800 inducing points, 5 latent processes, and 6 mixtures in the SM kernel, while the SP linkage was defined through 160 inducing points, 5 latent processes, and 6 mixtures in the SM kernel. Each model was trained with 32 bit-precision for 4,000 epochs with a minibatch size of 1024. Optimization of model parameters was performed with the Adam optimizer and cosine annealing from an initial learning rate of 1e-2, ending at 1e-8.

Due to the increasingly complex material response in the high-contrast case, we found improved performance for the SP linkage through the creation of a deep Gaussian process [74, 75]. The model consists of 3 layers, each constructed with a Matérn 5/2 kernel and 400 inducing points. Similarly, the model was trained with the Adam optimizer and the same floating point precision, learning rate schedule, and minibatch size for 2,000 epochs.

The available experimental dataset was partitioned into a 80/20 train/test split. Parity plots of each model's performance can be seen in Figure A.1, with corresponding error metrics in Table A.1. The additional challenge of learning the complex nonlinear mapping between structure and property in the high-contrast ratio case is clearly evident.

B Appendix: CNF Architecture

The architecture underpinning the CNF in both the low- and high-contrast cases was identified as the result of an ablation study across 6 various architectures. Relatively simple architectures were considered in this work, although we have no doubt that further refinement may be possible. Due to the low-dimensional microstructure space considered, the architectures considered are significantly more light-weight than similar models considered in other works [23, 24, 26–29], albeit no other models exist which enable inversion of the complete PSP linkage. The networks and training settings considered are displayed in Figure B.1. All models were trained with the Adam optimizer [76] and 32 bit-precision. Specifics regarding the architecture selected can be found in Table B.1. Training

Table A.1. Summary error metrics (Mean Absolute Error (MAE) and Normalized Mean Absolute Error (NMAE)) of SV-MOGP forward model performance with reported mean and standard deviation.

		α_1	α_2	α_3	α_4	α_5
MAE	Train	3.396±8.458	0.377±0.515	3.376±4.401	4.133±5.727	4.348±6.119
	Test	4.055±12.113	0.411±0.661	3.573±4.636	4.302±6.211	4.826±7.087
NMAE (%)	Train	1.746±4.350	0.853±1.164	8.879±11.577	18.230±25.258	31.317±44.078
	Test	2.086±5.809	0.929±1.164	9.400±11.577	18.975±27.394	36.447±51.054
		E_1 (GPa)	E_2 (GPa)	G_{12} (GPa)	k_1 (W/mK)	k_2 (W/mK)
MAE	Train	0.267±0.341	0.272±0.355	0.094±0.103	0.233±0.310	0.234±0.100
	Test	0.272±0.345	0.262±0.336	0.094±0.101	0.231±0.299	0.232±0.304
NMAE (%)	Train	0.247±0.315	0.252±0.328	0.228±0.249	0.074±0.098	0.074±0.100
	Test	0.252±0.319	0.242±0.311	0.227±0.246	0.073±0.095	0.073±0.096
MAE	Train	0.579±0.957	0.566±0.915	0.198±0.298	0.578±1.099	0.611±1.217
	Test	0.744±1.381	0.777±1.481	0.264±0.464	0.896±2.279	0.838±2.071
NMAE (%)	Train	2.990±4.945	2.923±4.726	2.760±4.159	2.241±4.423	2.348±4.713
	Test	3.844±7.134	4.010±7.646	3.677±6.475	3.477±8.844	3.246±8.023

Table B.1. Summary of CNF model architectures considered.

Hyperparameter	Arch. 1	Arch. 2	Arch. 3	Arch. 4	Arch. 5	Arch. 6
Depth	3	4	6	3	4	6
Width	128	128	128	256	256	256

was performed for 1000 epochs with a batch size of 1024 (9766 iterations), cosine annealing of the learning rate from $1e-3$ to $1e-8$, an EMA decay rate of 0.9999, and $\sigma_{\text{min}} = 1e-3$. All probability flow paths were solved for at $t = 1$ with the adaptive step size `dopri5` solver and tolerances $\texttt{atol} = \texttt{rtol} = 1e-5$.

Model performance was evaluated for the low-contrast dataset for the purposes of architecture selection. Posteriors were sampled for each conditioning property set available in the dataset, and the NMAE evaluated across all samples. The threshold for the 95$^{\text{th}}$ percentile was also evaluated, with the results of both across the architectures considered presented in Figure C.3. This process is identical to the process laid out in constructing Figure 7. Architecture #5 was selected for all case studies presented in this work, as it was the most light-weight architecture displaying minimal values for the error metrics tracked in Table B.2 and Figure C.3.

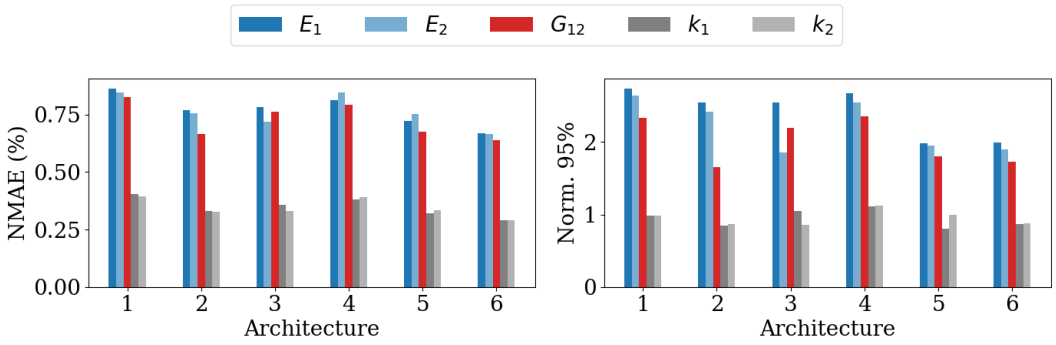

Figure B.1. Results of ablation study. (a) NMAE determined through the propagation of 128 posterior samples, conditioned across all available property sets. Propagated posterior means are evaluated directly against the condioning property set. (b) Normalized 95$^{\text{th}}$ percentile of properties, as determined through the estimated CDF from all propagated posterior samples.

Table B.2. Model architectures and mean error metrics considered across the property space.

		Arch. 1	Arch. 2	Arch. 3	Arch. 4	Arch. 5	Arch. 6
E_1 (GPa)	MAE	0.937±0.969	0.822±0.742	0.850±0.843	0.896±0.939	0.782±0.813	0.724±0.740
	NMAE (%)	0.864±0.778	0.768±0.616	0.783±0.654	0.814±0.695	0.723±0.633	0.668±0.564
	95$^{\text{th}}$	2.959	2.751	2.749	2.887	2.144	2.150
	Norm. 95$^{\text{th}}$	2.738	2.546	2.544	2.672	1.984	1.990
E_2 (GPa)	MAE	0.917±0.958	0.810±0.734	0.777±0.832	0.929±0.929	0.809±0.802	0.718±0.731
	NMAE (%)	0.846±0.770	0.757±0.609	0.718±0.647	0.846±0.691	0.751±0.627	0.666±0.559
	95$^{\text{th}}$	2.860	2.616	2.009	2.754	2.103	2.047
	Norm. 95$^{\text{th}}$	2.647	2.422	1.859	2.549	1.946	1.895
G_{12} (GPa)	MAE	0.345±367	0.273±0.288	0.319±0.330	0.332±0.359	0.279±0.310	0.265±0.285
	NMAE (%)	0.826±0.768	0.665±0.621	0.762±0.660	0.792±0.693	0.674±0.632	0.639±0.565
	95$^{\text{th}}$	0.963	0.683	0.904	0.970	0.743	0.711
	Norm. 95$^{\text{th}}$	2.337	1.658	2.193	2.353	1.803	1.726
k_1 (W/mK)	MAE	1.269±1.204	1.036±0.987	1.134±1.114	1.201±1.149	1.007±0.997	0.920±0.908
	NMAE (%)	0.404±0.388	0.332±0.329	0.359±0.347	0.380±0.352	0.321±0.324	0.292±0.290
	95$^{\text{th}}$	3.122	2.685	3.326	3.514	2.543	2.732
	Norm. 95$^{\text{th}}$	0.988	0.850	1.053	1.112	0.905	0.864
k_2 (W/mK)	MAE	1.244±1.206	1.106±0.988	1.048±1.108	1.241±1.151	1.042±0.997	0.914±0.907
	NMAE (%)	0.396±0.389	0.326±0.329	0.332±0.345	0.393±0.353	0.333±0.324	0.291±0.290
	95$^{\text{th}}$	3.121	2.746	2.695	3.558	3.136	2.766
	Norm. 95$^{\text{th}}$	0.988	0.869	0.853	1.126	0.992	0.875

C Appendix: Resulting Microstructures

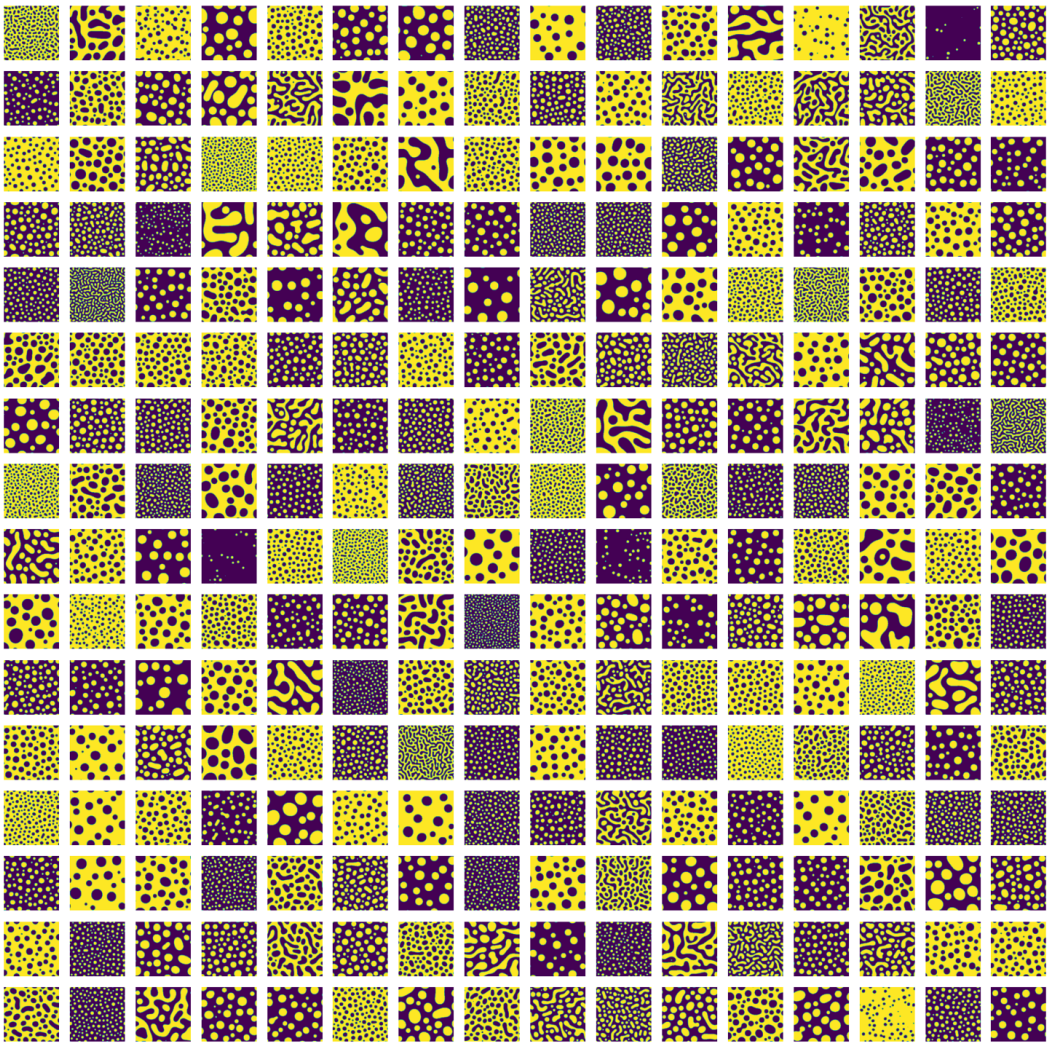

Figure C.1. Heterogeneous microstructures drawn at random from the available dataset.

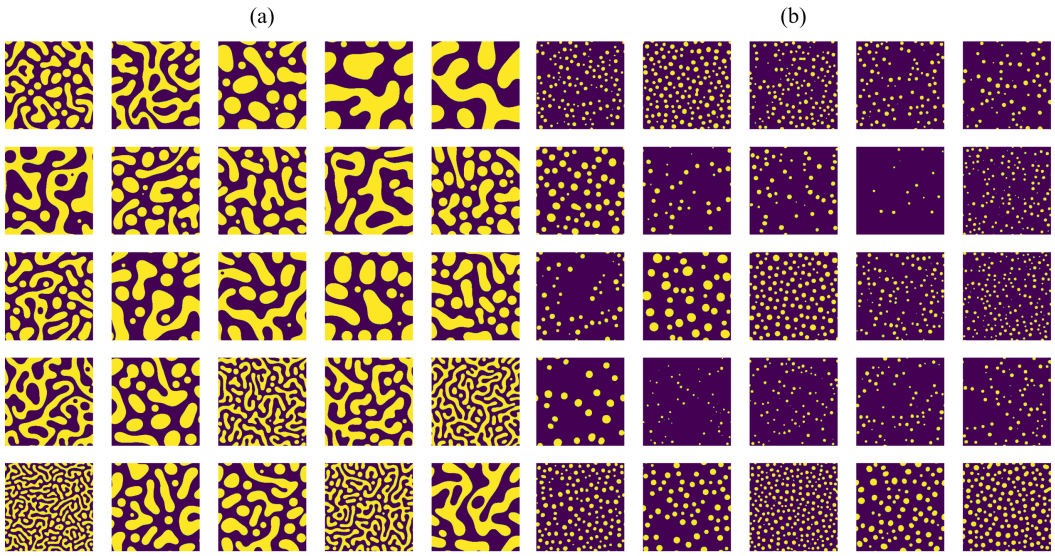

Figure C.2. Heterogeneous microstructures identified through transforming 25 posterior samples through the phase-field model forward process, $p(m|\boldsymbol{k}^*)$. (a) Al-Si low-contrast ratio microstructures for the first test case, and (b) the second test case.

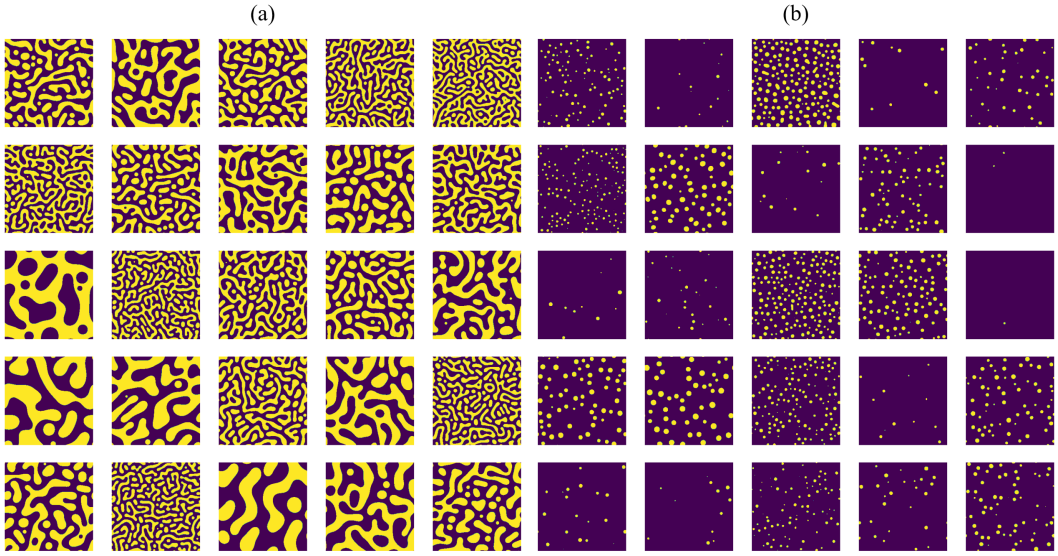

Figure C.3. Heterogeneous microstructures identified through transforming 25 posterior samples through the phase-field model forward process, $p(m|\boldsymbol{k}^*)$. (a) High-contrast ratio microstructures for the first test case, and (b) the second test case.

D Appendix: Notation

Table D.1. Definitions of common variables used in this work

Variable	Definition
Ω_1	Process Space
Ω_2	Structure Space
Ω_3	Property Space
$\mathcal{F}$	Process-Structure forward mapping (tilde denotes surrogate model)
$\mathcal{G}$	Structure-Property forward mapping (tilde denotes surrogate model)
$\mathcal{H}$	Process-Property forward mapping (tilde denotes surrogate model)
m_s^α	Discrete microstructure function for state α in voxel s
$f_r^{\alpha\beta}$	2-point spatial correlations between phases α and β with offset r
$\boldsymbol{\theta}$	Vector of processing parameters
θ_1, θ_2	First two processing parameters for spinodal decomposition (mobilities for each phase)
θ_3	Third processing parameter for spinodal decomposition (initial relative concentration)
$\boldsymbol{k}$	Vector of material properties
E	Young's Modulus
ν	Poisson Ratio
κ	Thermal conductivity
$p_n(\boldsymbol{\theta}\|\boldsymbol{k}^*)$	Posterior distribution of process parameters given a vector of target properties $\boldsymbol{k}^*$

References

[1] G. Olson, Computational Design of Hierarchically Structured Materials, Science 277 (1997) 1237–1242. doi:10.1126/science.277.5330.1237.

[2] E. Kalnay, Atmospheric Modeling, Data Assimilation and Predictability, iSBN: 9780511802270 Publisher: Cambridge University Press (Nov. 2002). doi:10.1017/CBO9780511802270.
URL https://www.cambridge.org/highereducation/books/atmospheric-modeling-data-assimilation-and-predictability/C5FD207439132836E85027754CE9BC1A

[3] A. Abdul-Aziz, L. J. Ghosn, G. Y. Baaklini, R. Bhatt, Combined NDE/finite element technique to study the effects of matrix porosity on the behavior of ceramic matrix composites, San Diego, CA, 2003, p. 144. doi:10.1117/12.484775.
URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.484775

[4] A. F. Bennett, Inverse Modeling of the Ocean and Atmosphere, Cambridge University Press, Cambridge, 2002. doi:10.1017/CBO9780511535895.
URL https://www.cambridge.org/core/books/inverse-modeling-of-the-ocean-and-atmosphere/5CFE830AEB412531F6419D17DB620B90

[5] J. Chen, D. Kipping, PROBABILISTIC FORECASTING OF THE MASSES AND RADII OF OTHER WORLDS, The Astrophysical Journal 834 (1) (2016) 17. doi:10.3847/1538-4357/834/1/17.
URL https://iopscience.iop.org/article/10.3847/1538-4357/834/1/17

[6] C. Lubich, From Quantum to Classical Molecular Dynamics: Reduced Models and Numerical Analysis, iSBN: 9783037190678 9783037195673 (Sep. 2008). doi:10.4171/067.
URL https://ems.press/books/zlam/55

[7] J. Adler, O. Oktem, Learned Primal-Dual Reconstruction, IEEE transactions on medical imaging 37 (6) (2018) 1322–1332. doi:10.1109/TMI.2018.2799231.

[8] E. Candes, M. Wakin, An Introduction To Compressive Sampling, IEEE Signal Processing Magazine 25 (2) (2008) 21–30. doi:10.1109/MSP.2007.914731.
URL http://ieeexplore.ieee.org/document/4472240/

[9] D. Donoho, Compressed sensing, IEEE Transactions on Information Theory 52 (4) (2006) 1289–1306. doi:10.1109/TIT.2006.871582.
URL http://ieeexplore.ieee.org/document/1614066/

[10] D. L. McDowell, Microstructure-Sensitive Computational Structure-Property Relations in Materials Design, in: D. Shin, J. Saal (Eds.), Computational Materials System Design, Springer International Publishing, Cham, 2018, pp. 1–25. doi:10.1007/978-3-319-68280-8_1.
URL https://doi.org/10.1007/978-3-319-68280-8_1

[11] D. L. McDowell, A perspective on trends in multiscale plasticity, International Journal of Plasticity 26 (9) (2010) 1280–1309, special Issue In Honor of David L. McDowell. doi:https://doi.org/10.1016/j.ijplas.2010.02.008.
URL https://www.sciencedirect.com/science/article/pii/S0749641910000306

[12] D. J. Luscher, D. L. McDowell, C. A. Bronkhorst, A second gradient theoretical framework for hierarchical multiscale modeling of materials, International Journal of Plasticity 26 (8) (2010) 1248–1275, special Issue In Honor of Lallit Anand. doi:https://doi.org/10.1016/j.ijplas.2010.05.006.
URL https://www.sciencedirect.com/science/article/pii/S0749641910000719

[13] S. R. Kalidindi, Hierarchical Materials Informatics: Novel Analytics for Materials Data, 1st Edition, Butterworth-Heinemann, Amsterdam, 2015.

[14] B. L. Adams, S. R. Kalidindi, D. T. Fullwood, Microstructure-Sensitive Design for Performance Optimization, in: Microstructure Sensitive Design for Performance Optimization, Elsevier, 2013, p. i. doi:10.1016/B978-0-12-396989-7.01001-7.
URL https://linkinghub.elsevier.com/retrieve/pii/B9780123969897010017

[15] D. T. Fullwood, S. R. Niezgoda, B. L. Adams, S. R. Kalidindi, Microstructure sensitive design for performance optimization, Progress in Materials Science 55 (6) (2010) 477–562. doi:10.1016/j.pmatsci.2009.08.002.
URL http://www.sciencedirect.com/science/article/pii/S0079642509000760

[16] G. H. Harrington, C. Kelly, V. Attari, R. Arroyave, S. R. Kalidindi, Application of a Chained-ANN for Learning the Process–Structure Mapping in Mg2SixSn1-x Spinodal Decomposition, Integrating Materials and Manufacturing Innovation 11 (3) (2022) 433–449. doi:10.1007/s40192-022-00274-3.
URL https://doi.org/10.1007/s40192-022-00274-3

[17] P. Caron, T. Khan, Evolution of Ni-based superalloys for single crystal gas turbine blade applications, Aerospace Science and Technology 3 (8) (1999) 513–523. doi:10.1016/S1270-9638(99)00108-X.
URL https://linkinghub.elsevier.com/retrieve/pii/S127096389900108X

[18] M. R. Palacín, Recent advances in rechargeable battery materials: a chemist's perspective, Chemical Society Reviews 38 (9) (2009) 2565–2575, publisher: The Royal Society of Chemistry. doi:10.1039/B820555H.
URL https://pubs.rsc.org/en/content/articlelanding/2009/cs/b820555h

[19] D. S. Li, H. Garmestani, S. Ahzi, Processing path optimization to achieve desired texture in polycrystalline materials, Acta Materialia 55 (2) (2007) 647–654. doi:10.1016/j.actamat.2006.04.041.
URL https://www.sciencedirect.com/science/article/pii/S1359645406006458

[20] J. B. Shaffer, M. Knezevic, S. R. Kalidindi, Building texture evolution networks for deformation processing of polycrystalline fcc metals using spectral approaches: Applications to process design for targeted performance, International Journal of Plasticity 26 (8) (2010) 1183–1194. doi:10.1016/j.ijplas.2010.03.010.

URL https://www.sciencedirect.com/science/article/pii/S0749641910000586

[21] S. R. Niezgoda, D. T. Fullwood, S. R. Kalidindi, Delineation of the space of 2-point correlations in a composite material system, Acta Materialia 56 (18) (2008) 5285–5292. doi:10.1016/j.actamat.2008.07.005.
URL http://www.sciencedirect.com/science/article/pii/S1359645408004886

[22] T. Butler, J. Jakeman, T. Wildey, Combining Push-Forward Measures and Bayes' Rule to Construct Consistent Solutions to Stochastic Inverse Problems, SIAM Journal on Scientific Computing 40 (2) (2018) A984–A1011. doi:10.1137/16M1087229.
URL https://epubs.siam.org/doi/10.1137/16M1087229

[23] J. Jung, J. I. Yoon, H. K. Park, H. Jo, H. S. Kim, Microstructure design using machine learning generated low dimensional and continuous design space, Materialia 11 (2020) 100690. doi:10.1016/j.mtla.2020.100690.
URL https://linkinghub.elsevier.com/retrieve/pii/S2589152920301071

[24] Z. Pei, K. A. Rozman, Ö. N. Doğan, Y. Wen, N. Gao, E. A. Holm, J. A. Hawk, D. E. Alman, M. C. Gao, Machine-Learning Microstructure for Inverse Material Design, Advanced Science 8 (23) (2021) 2101207. doi:10.1002/advs.202101207.
URL https://onlinelibrary.wiley.com/doi/10.1002/advs.202101207

[25] V. Fung, J. Zhang, G. Hu, P. Ganesh, B. G. Sumpter, Inverse design of two-dimensional materials with invertible neural networks, npj Computational Materials 7 (1) (2021) 200. doi:10.1038/s41524-021-00670-x.
URL https://www.nature.com/articles/s41524-021-00670-x

[26] R. K. Tan, N. L. Zhang, W. Ye, A deep learning–based method for the design of microstructural materials, Structural and Multidisciplinary Optimization 61 (4) (2020) 1417–1438. doi:10.1007/s00158-019-02424-2.
URL http://link.springer.com/10.1007/s00158-019-02424-2

[27] H. J. Lim, K.-H. Lee, G. J. Yun, Microstructure Design of Multifunctional Particulate Composite Materials using Conditional Diffusion Models, arXiv:2301.09051 [cond-mat, physics:physics] (Jan. 2023). doi:10.48550/arXiv.2301.09051.
URL http://arxiv.org/abs/2301.09051

[28] Y. Kim, H. K. Park, J. Jung, P. Asghari-Rad, S. Lee, J. Y. Kim, H. G. Jung, H. S. Kim, Exploration of optimal microstructure and mechanical properties in continuous microstructure space using a variational autoencoder, Materials & Design 202 (2021) 109544. doi:10.1016/j.matdes.2021.109544.
URL https://www.sciencedirect.com/science/article/pii/S0264127521000976

[29] L. Xu, N. Hoffman, Z. Wang, H. Xu, Harnessing structural stochasticity in the computational discovery and design of microstructures, Materials & Design 223 (2022) 111223. doi:10.1016/j.matdes.2022.111223.
URL https://www.sciencedirect.com/science/article/pii/S0264127522008450

[30] A. Deshwal, J. Doppa, Combining Latent Space and Structured Kernels for Bayesian Optimization over Combinatorial Spaces, in: Advances in Neural Information Processing Systems, Vol. 34, Curran Associates, Inc., 2021, pp. 8185–8200.
URL https://proceedings.neurips.cc/paper/2021/hash/44e76e99b5e194377e955b13fb12f630-Abstract.html

[31] A. Solomou, G. Zhao, S. Boluki, J. K. Joy, X. Qian, I. Karaman, R. Arróyave, D. C. Lagoudas, Multi-objective Bayesian materials discovery: Application on the discovery of precipitation strengthened NiTi shape memory alloys through micromechanical modeling, Materials & Design 160 (2018) 810–827. doi:10.1016/j.matdes.2018.10.014.
URL https://www.sciencedirect.com/science/article/pii/S026412751830769X

[32] A. Talapatra, S. Boluki, T. Duong, X. Qian, E. Dougherty, R. Arróyave, Autonomous efficient experiment design for materials discovery with Bayesian model averaging, Physical Review Materials 2 (11) (2018) 113803, publisher: American Physical Society. doi:10.1103/PhysRevMaterials.2.113803.
URL https://link.aps.org/doi/10.1103/PhysRevMaterials.2.113803

[33] D. P. Kingma, M. Welling, Auto-Encoding Variational Bayes, arXiv:1312.6114 [cs, stat] (2013).
URL http://arxiv.org/abs/1312.6114

[34] J. Ho, A. Jain, P. Abbeel, Denoising Diffusion Probabilistic Models, arXiv:2006.11239 [cs, stat] (Dec. 2020). doi:10.48550/arXiv.2006.11239.
URL http://arxiv.org/abs/2006.11239

[35] L. Ardizzone, J. Kruse, S. Wirkert, D. Rahner, E. W. Pellegrini, R. S. Klessen, L. Maier-Hein, C. Rother, U. Köthe, Analyzing Inverse Problems with Invertible Neural Networks, arXiv:1808.04730 [cs, stat] (Feb. 2019). doi:10.48550/arXiv.1808.04730.
URL http://arxiv.org/abs/1808.04730

[36] D. J. Rezende, S. Mohamed, Variational Inference with Normalizing Flows, arXiv:1505.05770 [cs, stat] (Jun. 2016).
URL http://arxiv.org/abs/1505.05770

[37] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. K. Duvenaud, Neural Ordinary Differential Equations, in: Advances in Neural Information Processing Systems, Vol. 31, Curran Associates, Inc., 2018.
URL https://proceedings.neurips.cc/paper_files/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html

[38] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, M. Le, Flow Matching for Generative Modeling.

[39] S. Sridhara, A. Chandrasekhar, K. Suresh, A generalized framework for microstructural optimization using neural networks, Materials & Design 223 (2022) 111213. doi:10.1016/j.matdes.2022.111213.
URL https://www.sciencedirect.com/science/article/pii/S0264127522008358

[40] C. Zhang, J. Bütepage, H. Kjellström, S. Mandt, Advances in Variational Inference, IEEE Transactions on Pattern Analysis and Machine Intelligence 41 (8) (2019) 2008–2026, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. doi:10.1109/TPAMI.2018.2889774.

[41] S. Torquato, Random Heterogeneous Materials: Microstructure and Macroscopic Properties, Interdisciplinary Applied Mathematics, Springer-Verlag, New York, 2002. doi:10.1007/978-1-4757-6355-3.
URL https://www.springer.com/gp/book/9780387951676

[42] E. Kröner, Statistical Continuum Mechanics, Vol. 92 of CISM International Centre for Mechanical Sciences, Springer Vienna, Vienna, 1971. doi:10.1007/978-3-7091-2862-6.
URL http://link.springer.com/10.1007/978-3-7091-2862-6

[43] R. P. M. Dingreville, J. A. Stewart, E. Y. Chen, J. M. Monti, Benchmark problems for the mesoscale multiphysics phase field simulator (memphis), Tech. rep., Sandia National Lab.(SNL-NM), Albuquerque, NM (United States) (2020).

[44] I. Steinbach, Phase-field models in materials science, Modelling and simulation in materials science and engineering 17 (7) (2009) 073001.

[45] C. Miehe, M. Hofacker, F. Welschinger, A phase field model for rate-independent crack propagation: Robust algorithmic implementation based on operator splits, Computer Methods in Applied Mechanics and Engineering 199 (45-48) (2010) 2765–2778.

[46] J. W. Cahn, J. E. Hilliard, Free energy of a nonuniform system. i. interfacial free energy, The Journal of chemical physics 28 (2) (1958) 258–267.

[47] C. E. Rasmussen, C. K. I. Williams, Gaussian processes for machine learning, Adaptive computation and machine learning, MIT Press, Cambridge, Mass, 2006, oCLC: ocm61285753.

[48] M. Seeger, Gaussian processes for machine learning, International journal of neural systems 14 (02) (2004) 69–106.

[49] A. Gupta, A. Cecen, S. Goyal, A. K. Singh, S. R. Kalidindi, Structure–property linkages using a data science approach: Application to a non-metallic inclusion/steel composite system, Acta Materialia 91 (2015) 239–254. doi:10.1016/j.actamat.2015.02.045.
URL http://www.sciencedirect.com/science/article/pii/S1359645415001603

[50] M. I. Latypov, L. S. Toth, S. R. Kalidindi, Materials knowledge system for nonlinear composites, Computer Methods in Applied Mechanics and Engineering 346 (2019) 180–196. doi:10.1016/j.cma.2018.11.034.
URL http://www.sciencedirect.com/science/article/pii/S0045782518305930

[51] Y. C. Yabansu, P. Altschuh, J. Hötzer, M. Selzer, B. Nestler, S. R. Kalidindi, A digital workflow for learning the reduced-order structure-property linkages for permeability of porous membranes, Acta Materialia 195 (2020) 668–680. doi:10.1016/j.actamat.2020.06.003.
URL http://www.sciencedirect.com/science/article/pii/S1359645420304274

[52] A. Marshall, S. R. Kalidindi, Autonomous Development of a Machine-Learning Model for the Plastic Response of Two-Phase Composites from Micromechanical Finite Element Models, JOM 73 (7) (2021) 2085–2095. doi:10.1007/s11837-021-04696-w.
URL https://doi.org/10.1007/s11837-021-04696-w

[53] N. H. Paulson, M. W. Priddy, D. L. McDowell, S. R. Kalidindi, Reduced-order structure-property linkages for polycrystalline microstructures based on 2-point statistics, Acta Materialia 129 (2017) 428–438. doi:10.1016/j.actamat.2017.03.009.
URL http://www.sciencedirect.com/science/article/pii/S135964541730188X

[54] A. P. Generale, S. R. Kalidindi, Reduced-order models for microstructure-sensitive effective thermal conductivity of woven ceramic matrix composites with residual porosity, Composite Structures 274 (2021) 114399. doi:10.1016/j.compstruct.2021.114399.
URL https://linkinghub.elsevier.com/retrieve/pii/S0263822321008618

[55] S. R. Kalidindi, Feature engineering of material structure for AI-based materials knowledge systems, Journal of Applied Physics 128 (4) (2020) 041103, publisher: American Institute of Physics. doi:10.1063/5.0011258.
URL https://aip.scitation.org/doi/10.1063/5.0011258

[56] W. Grathwohl, R. T. Q. Chen, J. Bettencourt, I. Sutskever, D. Duvenaud, FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models, arXiv:1810.01367 [cs, stat] (Oct. 2018). doi:10.48550/arXiv.1810.01367.
URL http://arxiv.org/abs/1810.01367

[57] D. Onken, S. Wu Fung, X. Li, L. Ruthotto, OT-Flow: Fast and Accurate Continuous Normalizing Flows via Optimal Transport, Proceedings of the AAAI Conference on Artificial Intelligence 35 (10) (2021) 9223–9232. doi:10.1609/aaai.v35i10.17113.
URL https://ojs.aaai.org/index.php/AAAI/article/view/17113

[58] A. M. Stuart, Inverse problems: A Bayesian perspective, Acta Numerica 19 (2010) 451–559. doi:10.1017/S0962492910000061.
URL https://www.cambridge.org/core/product/identifier/S0962492910000061/type/journal_article

[59] A. Khosravani, A. Cecen, S. R. Kalidindi, Development of high throughput assays for establishing process-structure-property linkages in multiphase polycrystalline metals: Application to dual-phase steels, Acta Materialia 123 (2017) 55–69. doi:10.1016/j.actamat.2016.10.033.
URL https://www.sciencedirect.com/science/article/pii/S135964541630800X

[60] M. I. Latypov, S. R. Kalidindi, Data-driven reduced order models for effective yield strength and partitioning of strain in multiphase materials, Journal of Computational Physics 346 (2017) 242–261. doi:10.1016/j.jcp.2017.06.013.
URL https://www.sciencedirect.com/science/article/pii/S0021999117304588

[61] A. Iskakov, Y. C. Yabansu, S. Rajagopalan, A. Kapustina, S. R. Kalidindi, Application of spherical indentation and the materials knowledge system framework to establishing microstructure-yield strength linkages from carbon steel scoops excised from high-temperature exposed components, Acta Materialia 144 (2018) 758–767. doi:10.1016/j.actamat.2017.11.024.
URL https://www.sciencedirect.com/science/article/pii/S1359645417309679

[62] D. L. McDowell, F. P. E. Dunne, Microstructure-sensitive computational modeling of fatigue crack formation, International Journal of Fatigue 32 (9) (2010) 1521–1542. doi:10.1016/j.ijfatigue.2010.01.003.
URL https://www.sciencedirect.com/science/article/pii/S0142112310000162

[63] A. Muth, A. Venkatraman, R. John, A. Pilchak, S. R. Kalidindi, D. L. McDowell, Neighborhood spatial correlations and machine learning classification of fatigue hot-spots in Ti–6Al–4V, Mechanics of Materials 182 (2023) 104679. doi:10.1016/j.mechmat.2023.104679.
URL https://www.sciencedirect.com/science/article/pii/S0167663623001254

[64] P. Meyer, A. M. Waas, FEM predictions of damage in continous fiber ceramic matrix composites under transverse tension using the crack band method, Acta Materialia 102 (2016) 292–303. doi:10.1016/j.actamat.2015.09.002.
URL https://www.sciencedirect.com/science/article/pii/S1359645415006680

[65] S. Murakami, Continuum Damage Mechanics, Vol. 185 of Solid Mechanics and Its Applications, Springer Netherlands, Dordrecht, 2012. doi:10.1007/978-94-007-2666-6.
URL http://link.springer.com/10.1007/978-94-007-2666-6

[66] T. Fast, S. R. Kalidindi, Formulation and calibration of higher-order elastic localization relationships using the mks approach, Acta Mater. 59 (2011) 4595–4605. doi:10.1016/j.actamat.2011.04.005.

[67] S. Dassault, Abaqus/CAE User's Manual (2019) 1174.

[68] T. Säilynoja, P.-C. Bürkner, A. Vehtari, Graphical Test for Discrete Uniformity and its Applications in Goodness of Fit Evaluation and Multiple Sample Comparison, Statistics and Computing 32 (2) (2022) 32, arXiv:2103.10522 [stat]. doi:10.1007/s11222-022-10090-6.
URL http://arxiv.org/abs/2103.10522

[69] Alexander G. de G. Matthews, James Hensman, Richard E. Turner, Zoubin Ghahramani, On sparse variational methods and the Kullback-Leibler divergence between stochastic processes, in: Proceedings of Machine Learning Research, Artificial Intelligence and Statistics, 2016, pp. 231–239.

[70] M. Titsias, Variational Learning of Inducing Variables in Sparse Gaussian Processes, in: Proceedings of the Twelth International Conference on Artificial Intelligence and Statistics, PMLR, 2009, pp. 567–574, iSSN: 1938-7228.
URL https://proceedings.mlr.press/v5/titsias09a.html

[71] M. A. Alvarez, L. Rosasco, N. D. Lawrence, Kernels for Vector-Valued Functions: a Review, number: arXiv:1106.6251 arXiv:1106.6251 [cs, math, stat] (Apr. 2012).
URL http://arxiv.org/abs/1106.6251

[72] A. G. Journel, C. J. Huijbregts, Mining Geostatistics, Blackburn Press, 2003, google-Books-ID: Id1GAAAAYAAJ.

[73] A. G. Wilson, R. P. Adams, Gaussian Process Kernels for Pattern Discovery and Extrapolation, arXiv:1302.4245 [cs, stat] (Dec. 2013). doi:10.48550/arXiv.1302.4245.
URL http://arxiv.org/abs/1302.4245

[74] A. Damianou, N. D. Lawrence, Deep Gaussian Processes, in: Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics, PMLR, 2013, pp. 207–215, iSSN: 1938-7228.
URL https://proceedings.mlr.press/v31/damianou13a.html

[75] H. Salimbeni, M. Deisenroth, Doubly Stochastic Variational Inference for Deep Gaussian Processes, arXiv:1705.08933 [stat] (Nov. 2017). doi:10.48550/arXiv.1705.08933.
URL http://arxiv.org/abs/1705.08933

[76] D. P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, arXiv:1412.6980 [cs] (Jan. 2017). doi:10.48550/arXiv.1412.6980.
URL http://arxiv.org/abs/1412.6980