

PROBABILISTIC REDUCED-ORDER MODELING FOR STOCHASTIC PARTIAL DIFFERENTIAL EQUATIONS

Constantin Grigo¹ and Phaedon-Stelios Koutsourelakis¹

¹Continuum Mechanics Group,
Department of Mechanical Engineering,
Technical University of Munich,
D-85748 Garching b. München

Keywords: Reduced-order modeling, generative Bayesian model, SPDE, effective material properties

Abstract. *We discuss a Bayesian formulation to coarse-graining (CG) of PDEs where the coefficients (e.g. material parameters) exhibit random, fine scale variability. The direct solution to such problems requires grids that are small enough to resolve this fine scale variability which unavoidably requires the repeated solution of very large systems of algebraic equations.*

We establish a physically inspired, data-driven coarse-grained model which learns a low-dimensional set of microstructural features that are predictive of the fine-grained model (FG) response. Once learned, those features provide a sharp distribution over the coarse scale effective coefficients of the PDE that are most suitable for prediction of the fine scale model output.

This ultimately allows to replace the computationally expensive FG by a generative probabilistic model based on evaluating the much cheaper CG several times. Sparsity enforcing priors further increase predictive efficiency and reveal microstructural features that are important in predicting the FG response. Moreover, the model yields probabilistic rather than single-point predictions, which enables the quantification of the unavoidable epistemic uncertainty that is present due to the information loss that occurs during the coarse-graining process.

1 INTRODUCTION

Many engineering design problems such as flow in porous media and mechanical properties of composite materials require simulations that are capable of resolving the microstructure of the underlying medium. If the material components under consideration exhibit fine-scale heterogeneity, popular discretization schemes (e.g. finite elements) yield very large systems of algebraic equations. Pertinent solution strategies at best (e.g. multigrid methods) scale linearly with the dimension of the unknown state vector. Despite the ongoing improvements in computer hardware, repeated solutions of such problems, as is required in the context of uncertainty quantification (UQ), poses insurmountable difficulties.

It is obvious that viable strategies for such problems, as well as a host of other deterministic problems where repeated evaluations are needed such as inverse, control/design problems etc, should focus on methods that exhibit *sublinear* complexity with respect to the dimension of the original problem. In the context of UQ a popular and general such strategy involves the use of surrogate models or emulators which attempt to learn the input-output map implied by the fine-grained model. Such models, e.g. Gaussian Processes [26], polynomial chaos expansions [8], neural nets [2] and many more, are trained on a finite set of fine-grained model runs. Nevertheless, their performance is seriously impeded by the curse of dimensionality, i.e. they usually become inaccurate for input dimensions larger than a few tens or hundreds, or equivalently, the number of FG runs required to achieve an acceptable level of accuracy grows exponentially fast with the input dimension.

Alternative strategies for high-dimensional problems make use of multi-fidelity models [10, 16] as inexpensive predictors of the FG output. As shown in [12], lower-fidelity models whose output deviates significantly from that of the FG can still yield accurate estimates with significant computational savings, as long as the outputs of the models exhibit statistical dependence. In the case of PDEs where finite elements are employed as the FG, multi-fidelity solvers can be simply obtained by using coarser discretizations in space or time. While linear and nonlinear dimensionality reduction techniques are suitable for dealing with high-dimensional inputs [14], it is known which of the microstructural features are actually predictive of FG outputs [22].

The model proposed in the present paper attempts to address this question. By using a two-component Bayesian network, we are able to predict fine-grained model outputs based on only a finite number of training data runs and a repeated solution of a much coarser model. Uncertainties can be easily quantified as our model leads to probabilistic rather than point-like predictions.

2 THE FINE-GRAINED MODEL

Let (Ω, \mathcal{F}, P) be a probability space. Let \mathcal{H} be the Hilbert space of functions defined over the domain \mathcal{D} over which the physical problem is defined. We consider problems in the context of heterogeneous media which exhibit properties given by a random process $\lambda(\mathbf{x}, \xi)$ defined over the product space $\mathcal{D} \times \Omega$. The corresponding stochastic PDE may be written as

$$\mathcal{L}(\mathbf{x}, \lambda(\mathbf{x}, \xi))u(\mathbf{x}, \lambda(\mathbf{x}, \xi)) = f(\mathbf{x}), \quad +\text{B.C.} \quad (1)$$

where \mathcal{L} is a stochastic differential operator and $\mathbf{x} \in \mathcal{D}$, $\xi \in \Omega$ are elements of the physical domain and the sample space, respectively. Discretization of the random process

$\lambda(\mathbf{x}, \xi) \xrightarrow{\text{discretize}} \boldsymbol{\lambda}_f \in \mathbb{R}^{n_f}$ as well as the governing equation leads to a system of n_f (potentially

nonlinear) algebraic equations, which can be written in residual form as

$$\mathbf{r}_f(\mathbf{U}_f; \boldsymbol{\lambda}_f) = \mathbf{0}, \quad (2)$$

where $\mathbf{U}_f(\boldsymbol{\lambda}_f) \in \mathbb{R}^{n_f}$ is the n_f -dimensional discretized solution vector for a given $\boldsymbol{\lambda}_f$ and $\mathbf{r}_f : \mathbb{R}^{n_f} \times \mathbb{R}^{n_{\boldsymbol{\lambda}_f}} \rightarrow \mathbb{R}^{n_f}$ the discretized residual vector. It is the model described by equation (2) which is denoted as the fine-grained model (FG) henceforth.

3 A GENERATIVE BAYESIAN SURROGATE MODEL

Let

$$p(\boldsymbol{\lambda}_f) = \int \delta(\boldsymbol{\lambda}_f - \boldsymbol{\lambda}_f(\xi)) p(\xi) d\xi \quad (3)$$

be the density of $\boldsymbol{\lambda}_f$. The density of the fine-scale response \mathbf{U}_f is then given by

$$p(\mathbf{U}_f) = \int p(\mathbf{U}_f | \boldsymbol{\lambda}_f) p(\boldsymbol{\lambda}_f) d\boldsymbol{\lambda}_f, \quad (4)$$

where the conditional density $p(\mathbf{U}_f | \boldsymbol{\lambda}_f)$ degenerates to a $\delta(\mathbf{U}_f - \mathbf{U}_f(\boldsymbol{\lambda}_f))$ when the only uncertainties in the problem are due to $\boldsymbol{\lambda}_f$.

The objective of this paper is to approximate this input-output map implied by $\mathbf{U}_f(\boldsymbol{\lambda}_f)$, or equivalently in terms of probability densities, the conditional distribution $p(\mathbf{U}_f | \boldsymbol{\lambda}_f)$. The latter case can also account for problems where additional sources of uncertainty are present and the input-output map is stochastic. To that end, we introduce a coarse-grained model (CG) leading to an approximate distribution $\bar{p}(\mathbf{U}_f | \boldsymbol{\lambda}_f)$ which will be trained on a limited number of FG solutions $\mathcal{D}_N = \left\{ \boldsymbol{\lambda}_f^{(i)}, \mathbf{U}_f^{(i)}(\boldsymbol{\lambda}_f^{(i)}) \right\}_{i=1}^N$.

Our approximate model $\bar{p}(\mathbf{U}_f | \boldsymbol{\lambda}_f)$ employs a set of latent [3], reduced, collective variables which we denote by $\boldsymbol{\lambda}_c \in \mathbb{R}^{n_{\boldsymbol{\lambda}_c}}$ for reasons that will be apparent in the sequel, such that

$$\begin{aligned} \bar{p}(\mathbf{U}_f | \boldsymbol{\lambda}_f) &= \int \bar{p}(\mathbf{U}_f, \boldsymbol{\lambda}_c | \boldsymbol{\lambda}_f) d\boldsymbol{\lambda}_c \\ &= \int \underbrace{\bar{p}(\mathbf{U}_f | \boldsymbol{\lambda}_c)}_{\text{decoder}} \underbrace{\bar{p}(\boldsymbol{\lambda}_c | \boldsymbol{\lambda}_f)}_{\text{encoder}} d\boldsymbol{\lambda}_c. \end{aligned} \quad (5)$$

As it can be understood from the equation above, the latent variables $\boldsymbol{\lambda}_c$ act as a probabilistic filter (encoder) on the FG input $\boldsymbol{\lambda}_f$, retaining the features necessary for predicting the FG output \mathbf{U}_f . In order for $\bar{p}(\mathbf{U}_f | \boldsymbol{\lambda}_f)$ to approximate well $p(\mathbf{U}_f | \boldsymbol{\lambda}_f)$, the latent variables $\boldsymbol{\lambda}_c$ should not simply be the outcome of a dimensionality reduction on $\boldsymbol{\lambda}_f$. Even if $\boldsymbol{\lambda}_f$ is amenable to such a dimensionality reduction, it is not necessary that the $\boldsymbol{\lambda}_c$ found would be predictive of \mathbf{U}_f . Posed differently, it is not important that $\boldsymbol{\lambda}_c$ provides a high-fidelity encoding of $\boldsymbol{\lambda}_f$ but it suffices that it is capable of providing a good prediction of the corresponding $\mathbf{U}_f(\boldsymbol{\lambda}_f)$.

The aforementioned desiderata do not unambiguously define the form of the encoding/ decoding densities in (5) nor the type/dimension of the latent variables $\boldsymbol{\lambda}_c$. In order to retain some of the physical and mathematical structure of the FG, we propose employing a coarsened, discretized version of the original continuous equation (1). In residual form this can again be written as

$$\mathbf{r}_c(\mathbf{U}_c; \boldsymbol{\lambda}_c) = \mathbf{0}, \quad (6)$$




Figure 1: A two-step Bayesian network/generative model defining $\bar{p}(U_f|\lambda_f, \theta_c, \theta_{cf})$.

where $U_c \in \mathbb{R}^{n_c}$ is the n_c -dimensional ($n_c \ll n_f$) discretized solution and $\mathbf{r}_c : \mathbb{R}^{n_f} \times \mathbb{R}^{n_{\lambda_c}} \rightarrow \mathbb{R}^{n_c}$ is the discretized residual vector. Due to the significant discrepancy in dimensions $n_c \ll n_f$, the cost of solving the CG in (6) is negligible compared to the FG in (2).

It is clear that λ_c plays the role of effective/equivalent properties but it is not obvious (except for some limiting cases where homogenization results can be invoked [23]) how these should depend on the fine-scale input λ_f nor how the solution $U_c(\lambda_c)$ of the CG should relate to $U_f(\lambda_f)$ in (2). Furthermore, it is important to recognize a priori that the use of the reduced variables λ_c in combination with the coarse model in (6) would in general imply some information loss. The latter should introduce an additional source of uncertainty in the predictions of the fine-scale output U_f [25], even in the limit of infinite training data. For that purpose and in agreement with (5), we propose a generative probabilistic model composed of the following two densities:

- A probabilistic mapping from λ_f to λ_c , which determines the effective properties λ_c given λ_f . We write this as $p_c(\lambda_c|\lambda_f, \theta_c)$ where θ_c denotes a set of model parameters,
- A coarse-to-fine map $p_{cf}(U_f|U_c, \theta_{cf})$, which is the PDF of the FG output U_f given the output U_c of the CG. It is parametrized by θ_{cf} .

This model is illustrated graphically in Figure 1. The density p_c encodes λ_f into λ_c and the coarse-to-fine map p_{cf} plays the role of a decoder, i.e. given the CG output U_c , it predicts U_f .

Using the abbreviated notation $\theta = [\theta_c, \theta_{cf}]$, from (5) we obtain

$$\bar{p}(U_f|\lambda_f, \theta) = \int p_{cf}(U_f|U_c(\lambda_c), \theta_{cf})p_c(\lambda_c|\lambda_f, \theta_c)d\lambda_c, \quad (7)$$

where $U_c(\lambda_c)$ is the solution vector to equation (6). The previous discussion suggests the following generative process for drawing samples from $\bar{p}(U_f|\lambda_f, \theta)$ i.e. predicting the FG output U_f given a FG input λ_f ,

- draw a sample $\lambda_c \sim p_c(\lambda_c|\lambda_f, \theta_c)$,
- solve the CG to obtain $U_c(\lambda_c)$,
- draw a sample $U_f \sim p_{cf}(U_f|U_c(\lambda_c), \theta_{cf})$.

3.1 Model training

In order to train the model described above, it is a reasonable strategy to minimize the Kullback-Leibler divergence [5] between the target density $p(U_f, \lambda_f) = \delta(U_f - U_f(\lambda_f))$

and $\bar{p}(\mathbf{U}_f|\boldsymbol{\lambda}_f, \boldsymbol{\theta})$. As these are conditional distributions, the KL divergence would depend on $\boldsymbol{\lambda}_f$. In order to calibrate the model for the $\boldsymbol{\lambda}_f$ values that are of significance, we operate on the augmented densities $p(\mathbf{U}_f, \boldsymbol{\lambda}_f) = p(\mathbf{U}_f|\boldsymbol{\lambda}_f)p(\boldsymbol{\lambda}_f)$ and $\bar{p}(\mathbf{U}_f, \boldsymbol{\lambda}_f|\boldsymbol{\theta}) = \bar{p}(\mathbf{U}_f|\boldsymbol{\lambda}_f, \boldsymbol{\theta})p(\boldsymbol{\lambda}_f)$, where $p(\boldsymbol{\lambda}_f)$ is defined by (3). In particular, we propose minimizing with respect to $\boldsymbol{\theta}$

$$\begin{aligned}
 \text{KL}(p(\mathbf{U}_f, \boldsymbol{\lambda}_f)||\bar{p}(\mathbf{U}_f, \boldsymbol{\lambda}_f|\boldsymbol{\theta})) &= \text{KL}(p(\mathbf{U}_f, \boldsymbol{\lambda}_f)||\bar{p}(\mathbf{U}_f|\boldsymbol{\lambda}_f, \boldsymbol{\theta})p(\boldsymbol{\lambda}_f)) \\
 &= \int p(\mathbf{U}_f, \boldsymbol{\lambda}_f) \log \left(\frac{p(\mathbf{U}_f, \boldsymbol{\lambda}_f)}{\bar{p}(\mathbf{U}_f|\boldsymbol{\lambda}_f, \boldsymbol{\theta})p(\boldsymbol{\lambda}_f)} \right) d\mathbf{U}_f d\boldsymbol{\lambda}_f \\
 &= - \int p(\mathbf{U}_f, \boldsymbol{\lambda}_f) \log \bar{p}(\mathbf{U}_f|\boldsymbol{\lambda}_f, \boldsymbol{\theta}) d\mathbf{U}_f d\boldsymbol{\lambda}_f + H(p(\mathbf{U}_f, \boldsymbol{\lambda}_f)) \\
 &\approx -\frac{1}{N} \sum_{i=1}^N \log \bar{p}(\mathbf{U}_f^{(i)}|\boldsymbol{\lambda}_f^{(i)}, \boldsymbol{\theta}) + H(p(\mathbf{U}_f, \boldsymbol{\lambda}_f)),
 \end{aligned} \tag{8}$$

where N is the number of training samples drawn from $p(\mathbf{U}_f, \boldsymbol{\lambda}_f)$, i.e.

$$\boldsymbol{\lambda}_f^{(i)} \sim p(\boldsymbol{\lambda}_f), \quad \mathbf{U}_f^{(i)} = \mathbf{U}_f(\boldsymbol{\lambda}_f^{(i)}), \tag{9}$$

and $H(p(\mathbf{U}_f, \boldsymbol{\lambda}_f))$ is the entropy of $p(\mathbf{U}_f, \boldsymbol{\lambda}_f)$ that is nevertheless independent of the model parameters $\boldsymbol{\theta}$. It is obvious from the final expression in (8) that $\sum_{i=1}^N \log \bar{p}(\mathbf{U}_f^{(i)}|\boldsymbol{\lambda}_f^{(i)}, \boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf})$ is a log-likelihood function of the data \mathcal{D}_N which we denote by $L(\mathcal{D}_N|\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf})$. In a fully Bayesian setting, this can be complemented with a prior $p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf})$ leading to the posterior

$$p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}|\mathcal{D}_N) \propto e^{L(\mathcal{D}_N|\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf})} p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}). \tag{10}$$

It is up to the analyst if predictions using equation (7) are carried out using point estimates of $\boldsymbol{\theta}$ (e.g. maximum likelihood (MLE) or maximum a posteriori (MAP)) or if a fully Bayesian approach is followed by averaging over the posterior $p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}|\mathcal{D}_N)$. The latter has the added advantage of quantifying the uncertainty introduced due the finite training data N . We pursue the former in the following as it is computationally more efficient.

3.1.1 Maximizing the posterior

Our objective is to find $\boldsymbol{\theta}^* = [\boldsymbol{\theta}_c^*, \boldsymbol{\theta}_{cf}^*]$ which maximizes the posterior given in equation (10), i.e.

$$\begin{aligned}
 [\boldsymbol{\theta}_c^*, \boldsymbol{\theta}_{cf}^*] &= \arg \max_{\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}} e^{L(\mathcal{D}_N|\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf})} p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}) \\
 &= \arg \max_{\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}} (L(\mathcal{D}_N|\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}) + \log p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf})) \\
 &= \arg \max_{\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}} \left(\sum_{i=1}^N \log \bar{p}(\mathbf{U}_f^{(i)}|\boldsymbol{\lambda}_f^{(i)}, \boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}) + \log p(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}) \right).
 \end{aligned} \tag{11}$$

The main difficulty in this optimization problem arises from the log-likelihood term which involves the log of an analytically intractable integral with respect to $\boldsymbol{\lambda}_c$ since

$$\begin{aligned}
 L_i(\boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}) &= \log \bar{p}(\mathbf{U}_f^{(i)}|\boldsymbol{\lambda}_f^{(i)}, \boldsymbol{\theta}_c, \boldsymbol{\theta}_{cf}) \\
 &= \log \int p_{cf}(\mathbf{U}_f^{(i)}|\mathbf{U}_c(\boldsymbol{\lambda}_c^{(i)}), \boldsymbol{\theta}_{cf}) p_c(\boldsymbol{\lambda}_c^{(i)}|\boldsymbol{\lambda}_f^{(i)}, \boldsymbol{\theta}_c) d\boldsymbol{\lambda}_c^{(i)},
 \end{aligned} \tag{12}$$

where we note that an independent copy of $\lambda_c^{(i)}$ pertains to each data point i . Due to this integration, typical deterministic optimization algorithms are not applicable.

However, as λ_c appears as a latent variable, we may resort to the well-known Expectation-Maximization algorithm [6]. Using Jensen's inequality, we establish a lower bound on every single term L_i of the sum in the log-likelihood $L(\mathcal{D}_N|\theta_c, \theta_{cf})$ by employing an arbitrary density $q_i(\lambda_c^{(i)})$ (different for each sample i) as

$$\begin{aligned}
 L_i(\theta_c, \theta_{cf}) &= \log \bar{p}(\mathbf{U}_f^{(i)}|\lambda_f^{(i)}, \theta_c, \theta_{cf}) \\
 &= \log \int p_{cf}(\mathbf{U}_f^{(i)}|U_c(\lambda_c^{(i)}), \theta_{cf}) p_c(\lambda_c^{(i)}|\lambda_f^{(i)}, \theta_c) d\lambda_c^{(i)} \\
 &= \log \int q_i(\lambda_c^{(i)}) \frac{p_{cf}(\mathbf{U}_f^{(i)}|U_c(\lambda_c^{(i)}), \theta_{cf}) p_c(\lambda_c^{(i)}|\lambda_f^{(i)}, \theta_c)}{q_i(\lambda_c^{(i)})} d\lambda_c^{(i)} \\
 &\geq \int q_i(\lambda_c^{(i)}) \log \left(\frac{p_{cf}(\mathbf{U}_f^{(i)}|U_c(\lambda_c^{(i)}), \theta_{cf}) p_c(\lambda_c^{(i)}|\lambda_f^{(i)}, \theta_c)}{q_i(\lambda_c^{(i)})} \right) d\lambda_c^{(i)} \text{ (Jensen)} \\
 &= \mathcal{F}_i(q_i; \theta_c, \theta_{cf}),
 \end{aligned} \tag{13}$$

Hence, the log-posterior in (10) can be lower bounded by

$$\begin{aligned}
 \log p(\theta_c, \theta_{cf}|\mathcal{D}_N) &= \log L(\mathcal{D}_N|\theta_c, \theta_{cf}) + \log p(\theta_c, \theta_{cf}) \\
 &= \sum_{i=1}^N \log L_i(\theta_c, \theta_{cf}) + \log p(\theta_c, \theta_{cf}) \\
 &\geq \sum_{i=1}^N \mathcal{F}_i(q_i; \theta_c, \theta_{cf}) + \log p(\theta_c, \theta_{cf}) \\
 &= \mathcal{F}(\{q_i\}_{i=1}^N, \theta_c, \theta_{cf}) + \log p(\theta_c, \theta_{cf}).
 \end{aligned} \tag{14}$$

The introduction of the auxiliary densities q_i suggests the following recursive procedure [4] for maximizing the log-posterior:

E-step: Given some $\theta_c^{(t)}, \theta_{cf}^{(t)}$ in iteration t , find the auxiliary densities $q_i^{(t+1)}$ that maximize \mathcal{F} ,

M-step: Given $q_i^{(t+1)}$, find the parameters $\theta_c^{(t+1)}, \theta_{cf}^{(t+1)}$ that maximize \mathcal{F} .

It can be readily shown that the optimal q_i is given by

$$q_i(\lambda_c^{(i)}) \propto p_{cf}(\mathbf{U}_f^{(i)}|U_c(\lambda_c^{(i)}), \theta_{cf}) p_c(\lambda_c^{(i)}|\lambda_f^{(i)}, \theta_c) \tag{15}$$

with which the inequality in (13) becomes an equality. In fact, both E- and M-steps can be relaxed to find suboptimal $q_i^{(t)}, \theta_c^{(t)}, \theta_{cf}^{(t)}$, which enables the application of approximate schemes such as e.g. Variational Inference (VI) [24, 9]. For the M-step, we may resort to any (stochastic) optimization algorithm [17] or, on occasion, closed-form updates might also be feasible.

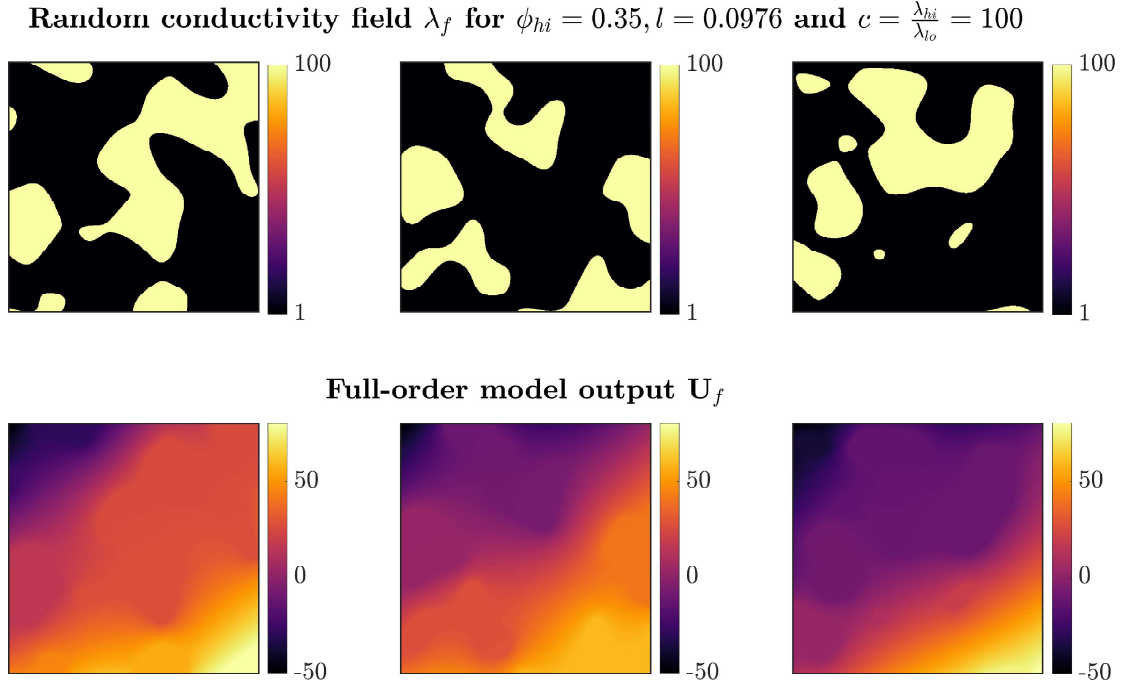


Figure 2: Random microstructure samples and corresponding fine-grained model outputs.

4 SAMPLE PROBLEM: 2D STATIONARY HEAT EQUATION

As a sample problem, we consider the 2D stationary heat equation on the unit square $[0, 1]^2$

$$\nabla_{\mathbf{x}}(-\lambda(\mathbf{x}, \xi(\mathbf{x}))\nabla_{\mathbf{x}}U(\mathbf{x}, \xi(\mathbf{x}))) = 0 \quad (16)$$

where $U(\mathbf{x}, \xi(\mathbf{x}))$ represents the temperature field. For the boundary conditions, we fix the temperature U in the upper left corner (see Figure 2) to -50 and prescribe the heat flux $\mathbf{Q}(\mathbf{x}) = \begin{pmatrix} 150-30y \\ 100-30x \end{pmatrix}$ on the remaining domain boundary $\partial\mathcal{D}$.

We consider a binary random medium whose conductivity $\lambda(\mathbf{x})$ can take the values λ_{hi} and λ_{lo} . To define such a field we consider transformations of a zero-mean Gaussian process $\xi(\mathbf{x})$ of the form [18, 13]

$$\lambda(\mathbf{x}, \xi(\mathbf{x})) = \begin{cases} \lambda_{hi}, & \text{if } \xi(\mathbf{x}) > c, \\ \lambda_{lo}, & \text{otherwise} \end{cases} \quad (17)$$

where the thresholding constant c is selected so as to achieve the target volume fraction ϕ_{hi} (or equivalently $\phi_{lo} = 1 - \phi_{hi}$) of the material with conductivity λ_{hi} (or equivalently of λ_{lo}). In the following we use an isotropic squared-exponential covariance function for $\xi(\mathbf{x})$ of the form

$$\text{cov}(\mathbf{x}_i, \mathbf{x}_j) = \exp\left\{-\frac{|\mathbf{x}_i - \mathbf{x}_j|^2}{l^2}\right\}. \quad (18)$$

In the following studies, we used values $l \approx 0.01$.

Due to the small correlation length l , we discretize the SPDE in Equation (16) using 256×256 standard quadrilateral finite elements, leading to a linear system of $N_{\text{eq}} = \dim(\mathbf{U}_f) - 1 =$

$257 \times 257 - 1 = 66048$ algebraic equations. We choose the discretization mesh of the random process $\lambda(\mathbf{x}, \xi(\mathbf{x}))$ to coincide with the finite element discretization mesh of the SPDE, i.e. $\dim(\boldsymbol{\lambda}_f) = 256 \times 256 = 65536$.

Samples $\mathcal{D}_N = \{\lambda_f^{(i)}, \mathbf{U}_f^{(i)}\}_{i=1}^N$ are readily obtained by simulating realizations of the discretized Gaussian field, transforming them according to (17) and solving the discretized SPDE. Three such samples can be seen in Figure 2.

4.1 Model specifications

We define a coarse model employing n_{λ_c} quadrilateral elements, the conductivities of which are given by the vector $\boldsymbol{\lambda}_c$. Since these need to be strictly positive, we operate instead on \mathbf{z}_c defined as

$$\mathbf{z}_c = \log \boldsymbol{\lambda}_c. \quad (19)$$

For each element $k = 1, \dots, n_{\lambda_c}$ of the coarse model/mesh, we assume that

$$z_{c,k} = \sum_{j=1}^{N_{\text{features}}} \theta_{c,j} \chi_j(\boldsymbol{\lambda}_{f,k}) + \sigma_k Z_k, \quad Z_k \sim \mathcal{N}(0, 1), \quad (20)$$

where $\boldsymbol{\lambda}_{f,k}$ is the subset of the vector $\boldsymbol{\lambda}_f$ that belongs to coarse element k and χ_j some feature functions of the underlying $\boldsymbol{\lambda}_{f,k}$ that we specify below. In a more compact form, we can write

$$p_c(\mathbf{z}_c | \boldsymbol{\lambda}_f, \boldsymbol{\theta}_c, \boldsymbol{\sigma}) = \mathcal{N}(\mathbf{z}_c | \boldsymbol{\Phi}(\boldsymbol{\lambda}_f) \boldsymbol{\theta}_c, \text{diag}(\boldsymbol{\sigma}^2)), \quad (21)$$

where $\boldsymbol{\Phi}(\boldsymbol{\lambda}_f)$ is an $n_{\lambda_c} \times N_{\text{features}}$ design matrix with $\Phi_{kj}(\boldsymbol{\lambda}_f) = \chi_j(\boldsymbol{\lambda}_{f,k})$ and $\text{diag}(\boldsymbol{\sigma}^2)$ is a diagonal covariance matrix with components σ_k^2 .

For the coarse-to-fine map p_{cf} , we postulate the relation

$$p_{cf}(\mathbf{U}_f | \mathbf{U}_c(\mathbf{z}_c), \boldsymbol{\theta}_{cf}) = \mathcal{N}(\mathbf{U}_f | \mathbf{W} \mathbf{U}_c(\mathbf{z}_c), \mathbf{S}), \quad (22)$$

where $\boldsymbol{\theta}_{cf} = (\mathbf{W}, \mathbf{S})$ are parameters to be learned from the data. The matrix \mathbf{W} is of dimension $n_f \times n_c$ and \mathbf{S} is a positive definite covariance matrix of size $n_f \times n_f$. To reduce the large amount of free parameters in the model, we fix \mathbf{W} to express coarse model's shape functions. Furthermore, we assume that $\mathbf{S} = \text{diag}(s)$ where s is the n_f -dimensional vector of diagonal entries of the diagonal matrix \mathbf{S} . The aforementioned expression implies, on average, a linear relation between the fine and coarse model outputs, which is supplemented by the residual Gaussian noise implied by \mathbf{S} . The latter expresses the uncertainty in predicting the FG output when only the CG solution is available. We note that the relation in Equation (20) (i.e. $\boldsymbol{\theta}_c$ and $\boldsymbol{\sigma}_k^2$) will be adjusted during training so that the model implied in Equation (22) represents the data as good as possible.

From equation (15), we have that

$$\begin{aligned} \log q_i^{(t+1)}(\mathbf{z}_c^{(i)}) &\propto \frac{1}{2} \sum_{k=1}^{N_{\text{el}}^c} \log \left((\sigma_k^{(t)})^{-2} \right) - \frac{1}{2} \sum_{k=1}^{N_{\text{el}}^c} (\sigma_k^{(t)})^{-2} \left(\mathbf{z}_c^{(i)} - \boldsymbol{\Phi}(\boldsymbol{\lambda}_f^{(i)}) \boldsymbol{\theta}_c^{(t)} \right)_k^2 \\ &\quad - \frac{1}{2} \sum_{j=1}^{N_{\text{el}}^f} \log s_j^{(t)} - \frac{1}{2} \sum_{j=1}^{N_{\text{el}}^f} s_j^{(t)} \left(\mathbf{T}_f^{(i)} - \mathbf{W}^{(t)} \mathbf{T}_c(\mathbf{z}_c^{(i)}) \right)_j^2, \end{aligned} \quad (23)$$

where with $(\cdot)_i$, we mean the i -th component of the vector in brackets. For use in the M-step, we compute the gradients

$$\frac{\partial \mathcal{F}}{\partial \sigma_k^{-2}} = \frac{N}{2} \sigma_k^2 - \frac{1}{2} \sum_{i=1}^N \left\langle \left(\mathbf{z}_c^{(i)} - \Phi(\boldsymbol{\lambda}_f^{(i)}) \boldsymbol{\theta}_c^{(t)} \right)_k \right\rangle_{q_i}, \quad (24)$$

$$\nabla_{\boldsymbol{\theta}_c} \mathcal{F} = \sum_{i=1}^N \left(\Phi^T(\mathbf{z}_f^{(i)}) \boldsymbol{\Sigma}^{-1} \langle \mathbf{z}_c^{(i)} \rangle_{q_i^{(t)}} - \Phi^T(\boldsymbol{\lambda}_f^{(i)}) \boldsymbol{\Sigma}^{-1} \Phi(\boldsymbol{\lambda}_f^{(i)}) \boldsymbol{\theta}_c \right) \quad (25)$$

where $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_{n_{\boldsymbol{\lambda}_c}}^2)$. We can readily compute the roots to find the optimal $\boldsymbol{\Sigma}, \boldsymbol{\theta}_c$. Furthermore, from $\frac{\partial \mathcal{F}}{\partial s_j} = 0$ we get

$$s_j^{(t+1)} = \frac{1}{N} \sum_{i=1}^N \left\langle \left(\mathbf{T}_f^{(i)} - \mathbf{W} \mathbf{T}_c(\mathbf{z}_c^{(i)}) \right)_j \right\rangle_{q_i^{(t)}}. \quad (26)$$

4.2 Feature functions χ

The framework advocated allows for any form and any number of feature functions χ_j in (20). Naturally such a selection can be guided by physical insight [11]. In practice, the feature functions we are using can be roughly classified into two different groups:

4.2.1 Effective-medium approximations

We consider existing effective-medium approximation formula that can be found in the literature [23] as ingredients for construction of feature functions χ_j in equation (20). The majority of commonly used such features only retain low-order topological information. The following approximations to the effective property λ_{eff} can be used as building blocks for feature functions χ in the sense that we can transform them nonlinearly such that $\chi(\boldsymbol{\lambda}_f) = f(\lambda_{\text{eff}}(\boldsymbol{\lambda}_f))$. In particular, as we are modeling $\mathbf{z}_c = \log \boldsymbol{\lambda}_c$, we include feature functions of type $\chi(\boldsymbol{\lambda}_f) = \log(\lambda_{\text{eff}}(\boldsymbol{\lambda}_f))$.

Maxwell-Garnett approximation (MGA) The Maxwell-Garnett approximation is assumed to be valid if the microstructure consists of a matrix phase λ_{mat} and spherical inclusions λ_{inc} that are small and dilute enough such that interactions between them can be neglected. Moreover, the heat flux far away from any inclusion is assumed to be constant. Under such conditions, the effective conductivity can be approximated in 2D as

$$\lambda_{\text{eff}} = \lambda_{\text{mat}} \frac{\lambda_{\text{mat}} + \lambda_{\text{inc}} + \phi_{\text{inc}}(\lambda_{\text{inc}} - \lambda_{\text{mat}})}{\lambda_{\text{mat}} + \lambda_{\text{inc}} - \phi_{\text{inc}}(\lambda_{\text{inc}} - \lambda_{\text{mat}})}, \quad (27)$$

where the volume fraction $\phi_i(\boldsymbol{\lambda}_f)$ is given by the fraction of phase i elements in the binary vector $\boldsymbol{\lambda}_f$.

Self-consistent approximation (SCA) The self-consistent approximation (or Bruggeman formula) was originally developed for effective electrical properties of random microstructures. It also considers non-interacting spherical inclusions and follows from the assumption that perturbations of the electric field due to the inclusions average to 0. In 2D, the formula reads

$$\lambda_{\text{eff}} = \frac{\alpha + \sqrt{\alpha^2 + 4\lambda_{\text{mat}}\lambda_{\text{inc}}}}{2}, \quad \alpha = \lambda_{\text{mat}}(2\phi_{\text{mat}} - 1) + \lambda_{\text{inc}}(2\phi_{\text{inc}} - 1). \quad (28)$$




Figure 3: Left: Sample microstructure for $l = 0.078$ and $\phi_{hi} = 0.7$. The blue line encompasses the convex area of the encircled low conducting phase blob, the two red lines are its maximum extent in x - and y -direction. The green lines are paths along which we count pixels and compute generalized means in the pixel-cross and straight path mean functions. Right: distance transform (distance to nearest black pixel) of the microstructure.

Note that the SCA exhibits phase inversion symmetry.

Differential effective medium (DEM) From a first-order expansion in volume fraction of the effective conductivity in the dilute limit of spherical inclusions, one can deduce the differential equation [23]

$$(1 - \phi_{inc}) \frac{d}{d\phi_{inc}} \lambda_{eff}(\phi_{inc}) = 2\lambda_{eff}(\phi_{inc}) \frac{\lambda_{inc} - \lambda_{eff}(\phi_{inc})}{\lambda_{inc} + \lambda_{eff}(\phi_{inc})}, \quad (29)$$

which can be integrated to

$$\left(\frac{\lambda_{inc} - \lambda_{eff}}{\lambda_{inc} - \lambda_{mat}} \right) \sqrt{\frac{\lambda_{mat}}{\lambda_{eff}}} = 1 - \phi_{inc}. \quad (30)$$

We solve for λ_{eff} and use it as a feature function χ .

4.2.2 Morphology-describing features

Apart from the effective-medium approximations which only take into account the phase conductivities and volume fractions, we wish to have feature functions χ_j that more thoroughly describe the morphology of the underlying microstructure. Popular members of this class of microstructural features are the two-point correlation, the lineal-path function, the pore-size density or the specific surface to mention only a few of them [23].

We are however free to use any function $\chi_j : (\mathbb{R}^+)^{\dim(\lambda_{f,k})} \mapsto \mathbb{R}$ as a feature, no matter from which field or consideration it may originate. We thus make use of existing image processing features [15] as well as novel topology-describing functions. Some important examples are

Convex area of connected phase blobs This feature identifies distinct, connected “blobs” of only high or low conducting phase pixels and computes the area of the convex hull to each blob. We then can e.g. use the mean or maximum value thereof as a feature function.

Blob extent From an identified phase blob, we can compute its maximum extension in x - and y -directions. One can for instance take the mean or maximum of maximum extension among identified blobs as a feature.

Distance transformation functions The distance transform of a binary image assigns a number to each pixel i that is the distance from that pixel i to the nearest nonzero pixel in the binary image, see right part of Figure 3. One can use either phase to correspond to nonzero in the binary image as well as different distance metrics. As a feature, one can e.g. take the mean or maximum of the distance transformed image.

Pixel-cross function This feature counts the number of high or low conducting pixels one has to cross going on a straight line from boundary to boundary in x - or y -direction. One can again take e.g. mean, maximum or minimum values as the feature function outputs.

Straight path mean function A further refinement of the latter function is to take generalized means instead of numbers of crossed pixels along straight lines from boundary to boundary. In particular, we use harmonic, geometric and arithmetic means as features.

4.3 Sparsity priors

It is clear from the previous discussion that the number of feature functions is practically limitless. The more such χ_j one introduces, the more unknown parameters $\theta_{c,j}$ must be learned. From the modeling point of view, while ML estimates can always be found, it is desirable to have as clear of a distinction as possible between relevant and irrelevant features that could provide further insight as well being able to do so with the fewest possible training data available. For that purpose we advocate the use of sparsity-enforcing priors in θ_c [1, 7, 19]. From a statistical perspective, this is also motivated by the *bias-variance-tradeoff*. Model prediction accuracy is adversely affected by two factors: One is noise in the training data (*variance*), the other is due to overly simple model assumptions (*bias*). Maximum-likelihood estimates of model parameters tend to have low bias but high variance, i.e. they accurately predict the training data but generalize poorly. To address this issue, a common Bayesian approach to control model complexity is the use of priors, which is the equivalent to regularization in frequentist formulations. A particularly appealing family of prior distributions is the Laplacian (or LASSO regression, [21]), as it sets redundant or unimportant predictors to exactly 0, thereby simplifying interpretation.

In particular, we use a prior on the coefficients θ_c of the form

$$\log p(\theta_c) = \log \frac{\sqrt{\gamma}}{2} - \sqrt{\gamma} \sum_{i=1}^{N_{\theta_c}} |\theta_{c,i}|, \quad (31)$$

with a hyper-parameter γ which can be set by either applying some cross-validation scheme or more efficiently by minimization of Stein’s unbiased risk estimate [20, 27]. A straightforward implementation of this prior is described in [7].

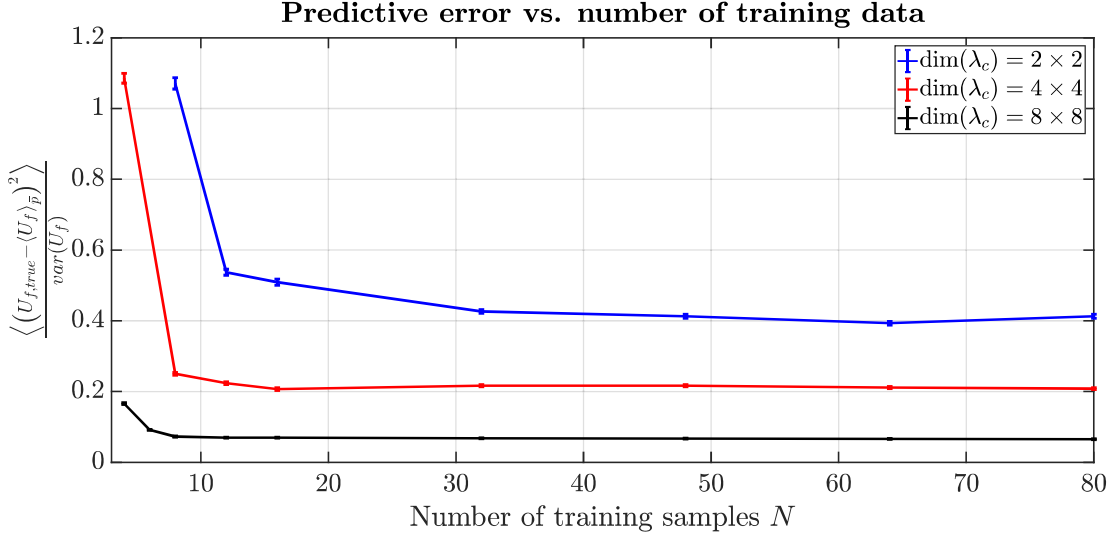


Figure 4: Relative squared prediction error $\frac{\langle (\mathbf{U}_{f,\text{true}} - \langle \mathbf{U}_f \rangle_{\bar{p}})^2 \rangle}{\text{var}(\mathbf{U}_f)}$ versus number training data samples N and different CG mesh sizes $N_{\text{el},c} = \dim(\lambda_c)$. We set $\phi_{\text{hi}} = 0.2$, $l = 0.0781$ and $c = 10$.

5 RESULTS

5.1 Required training data

We use a total number of 306 feature functions χ_j and a Laplacian prior (31) with a hyperparameter γ we find by cross-validation. We set the length scale parameter to $l = 0.0781$ and the expected volume fraction of the high conducting phase to $\phi_{\text{hi}} = 0.2$. For the contrast, we take $c = \frac{\lambda_{\text{hi}}}{\lambda_{\text{lo}}} = 10$, where we set $\lambda_{\text{lo}} = 1$. The full- and reduced-order models are both computed on regular quadrilateral finite element meshes of size 256×256 for the FG and 2×2 , 4×4 and 8×8 for the CG.

Our goal is to measure the predictive capabilities of the described model. To that end, we compute the mean squared distance

$$d^2 = \frac{1}{N_{\text{test}} n_f} \sum_{j=1}^{n_f} \sum_{i=1}^{N_{\text{test}}} \left(U_{f,\text{true},j}^{(i)} - \langle U_{f,j}^{(i)} \rangle_{\bar{p}(\mathbf{U}_f^{(i)} | \lambda_f^{(i)}, \theta)} \right)^2 \quad (32)$$

of the predictive mean $\langle \mathbf{U}_f^{(i)} \rangle_{\bar{p}(\mathbf{U}_f^{(i)} | \lambda_f^{(i)}, \theta)}$ to the true FG output $U_f^{(i)}$ on a test set of $N_{\text{test}} = 256$ samples. Predictions are carried out by drawing 10,000 samples of λ_c from the learned distribution $p_c(\log \lambda_c | \lambda_f, \theta_c^*, \sigma) = \mathcal{N}(\log(\lambda_c) | \Phi(\lambda_f) \theta_c^*, \text{diag}(\sigma^2))$, solving the coarse model $\mathbf{U}_c^{(i)} = \mathbf{U}_c(\lambda_c^{(i)})$ and drawing a sample from $p_{cf}(U_f^{(i)} | \mathbf{U}_c^{(i)}, \theta_{cf}^*) = \mathcal{N}(U_f^{(i)} | \mathbf{W} \mathbf{U}_c^{(i)}, \mathbf{S}^*)$ for every test data point. Monte Carlo noise is small enough to be neglected.

As a reference value for the computed error d^2 , we compute the mean variance of the FG output

$$\text{var}(U_f) = \frac{1}{n_f} \sum_{i=1}^{n_f} (\langle U_{f,i}^2 \rangle - \langle U_{f,i} \rangle^2), \quad (33)$$

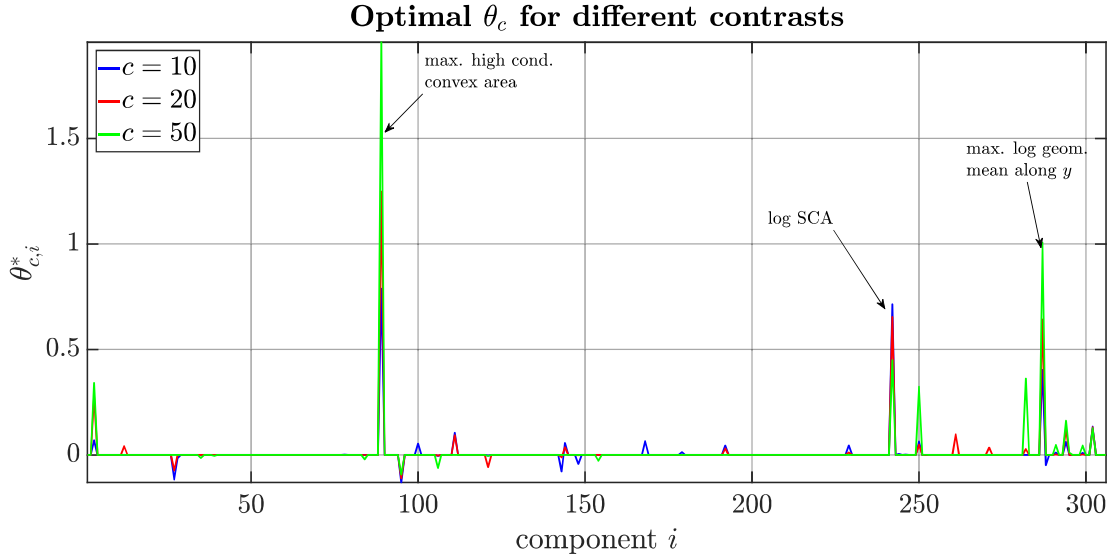


Figure 5: Components of the optimal θ_c^* for different values of contrast $c = \{10, 20, 50\}$. Most feature functions are deactivated by the sparsity prior. We observe that with increasing contrast, the importance of the log SCA diminishes at the expense of more geometric features.

where expectation values $\langle \cdot \rangle$ are estimated on a set of 1024 FG samples such that errors due to Monte Carlo can be neglected. Figure 4 shows the relative squared prediction error $\frac{d^2}{\text{var}(U_f)}$ in dependence of the number of training data samples for different coarse mesh sizes $\dim(\lambda_c) = 2 \times 2, 4 \times 4$ and 8×8 . We observe that the predictive error converges to a finite value already for relatively few training data. This is due to the inevitable information loss during the coarsening process $\lambda_f \rightarrow \lambda_c$ as well as finite element discretization errors. In accordance with that, we see that the error drops with the dimension of the coarse mesh, $\dim(\lambda_c)$.

5.2 Activated microstructural feature functions

For the same volume fraction and microstructural length scale parameter as above ($\phi_{\text{hi}} = 0.2, l = 0.0781$), a varying contrast of $c = \{10, 20, 50\}$, a coarse model dimension of $\dim(\lambda_c) = 4 \times 4$ and a training set of $N = 128$ samples, we find the optimal θ_c^* shown in Figure 5. Due to the application of a sparsity enforcing prior as described in section 4.3, we observe that most components of θ_c^* are exactly 0. Comparability between different feature functions can be ensured by standardization or normalization of feature function outputs on the training data. For all three contrast values, we see that the three most important feature functions are given by the maximum convex area of blobs of conductivity λ_{hi} , the maximum log of the geometric mean of conductivities along a straight line from boundary to boundary in y -direction and the log of the self-consistent effective medium approximation as described above. The maximum convex area feature returns the largest convex area of all blobs found within a coarse element. convex area. It is an interesting to note that although, in the set of 306 feature functions χ_j , the (max/min/mean/variance of) the blob area (for both high and low conducting phases) are also included, it is only the convex area which is activated.

In Figure 5, it is observed that with increasing contrast, the coefficient $\theta_{c,j}^*$ belonging to the log self-consistent approximation is decreasing in contrast to increasing values of $\theta_{c,j}^*$'s cor-

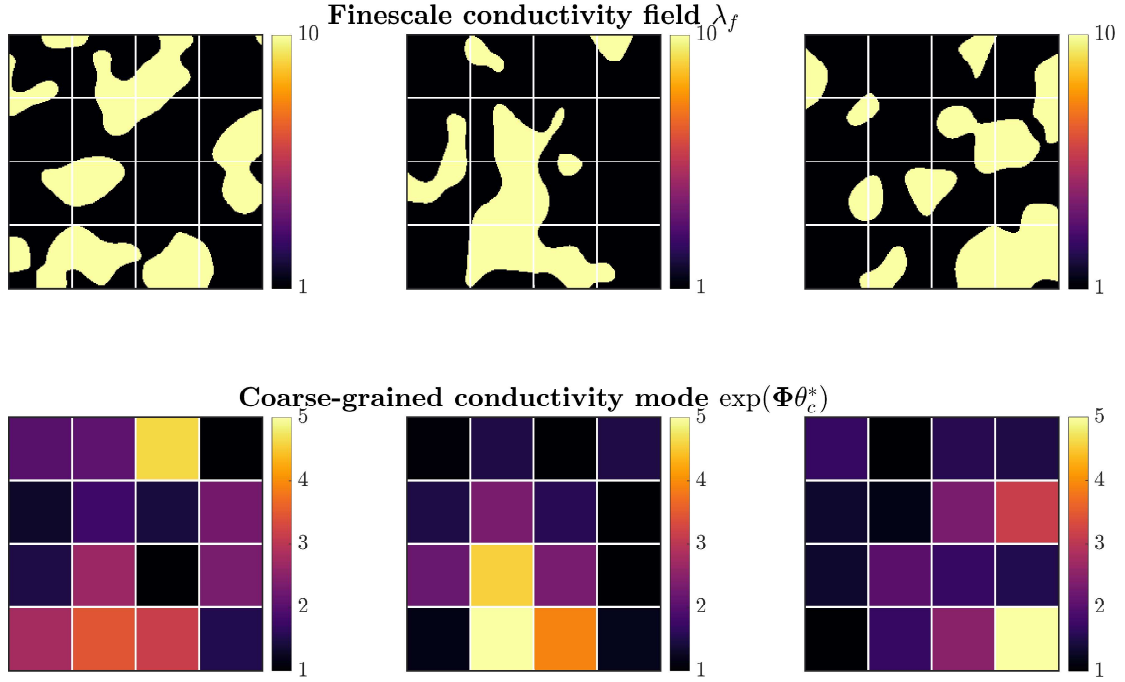


Figure 6: Three test samples and the corresponding mode $\exp(\Phi(\lambda_f)\theta_c^*)$ of p_c .

responding the max. high conducting convex area and the max. log geometric mean along y -direction. We believe that this is because, the higher the contrast c is, the more the exact geometry and connectedness of the microstructure plays a role for predicting the effective properties. As the SCA only considers the volume fraction and the conductivities of both phases, it disregards such information.

5.3 Learned effective property λ_c

Figure 6 shows three test samples of $\phi_{\text{hi}} = 0.2$, $l = 0.0781$ and $c = 10$ (top row) along with the mode $\exp(\Phi(\lambda_f)\theta_c^*)$ of the learned distribution of the effective conductivity $p_c(\lambda_c|\lambda_f, \theta_c^*)$, which is connected to the mean by $\langle \lambda_c \rangle_{p_c} = \exp(\Phi(\lambda_f)\theta_c^* + \frac{1}{2}\sigma^{*2})$. We emphasize again that the latent variable λ_c is not a lower-dimensional compression of λ_f with the objective of most accurately reconstructing λ_f , but of providing good predictions of $U_f(\lambda_f)$. Even though Figure 6 gives the impression of a simple local averaging relation between λ_f and λ_c , this is not always be the case. In particular, $p_c(\lambda_c|\lambda_f, \theta_c^*)$ was found to have non-vanishing probability mass for $\lambda_{c,i} < \lambda_{\text{lo}}$ or $\lambda_{c,i} > \lambda_{\text{hi}}$, especially for more general models where the coarse-to-fine mapping (22) is not fixed to be the shape function interpolation \mathbf{W} .

5.4 Predictive uncertainty




Figure 7: Variance parameters σ^{*2} and s^* of p_c and p_{cf} , respectively.

The predictive uncertainty is composed of the uncertainty in having an accurate encoding of λ_f in λ_c which is described by σ_k^2 in p_c (21), as well as the uncertainty in the reconstruction process from U_c to U_f , which is given by the diagonal covariance $S = \text{diag}(s)$ in p_{cf} (22). Both are shown after training ($\phi_2 = 0.2, l = 0.0781, c = 10$) in Figure 7. For σ^{*2} , we observe that values in the corner elements always converge to very tight values, whereas some non-corner elements can converge to comparably large values. The exact location of these elements is data dependent. The coarse-to-fine reconstruction variances s^* is depicted on the right column of Figure 7. As expected, we see that the estimated coarse-to-fine reconstruction error is largest in the center of coarse elements i.e. at large distances from the coarse model finite element nodes.

5.5 Predictions

As in section 5.1, for a model with $N_{\text{train}} = 32$ and $\dim(\boldsymbol{\lambda}_c) = 8 \times 8$ and random microstructures with parameter values $\phi_{\text{hi}} = 0.2$, $l = 0.0781$, $c = 10$, we consider predictions by sampling from $\bar{p}(U_f | \boldsymbol{\lambda}_f, \boldsymbol{\theta}_{cf}, \boldsymbol{\theta}_c)$ using 10,000 samples. The predictive histogram for the temperature $U_{f,lr}$ in the lower right corner of the domain can be seen in Figure 8. Figure 9 shows a surface plot of the true solution (colored), the predictive mean (blue) $\pm \sigma$ (gray). As can be seen, the true solution U_f is nicely included in \bar{p} everywhere.

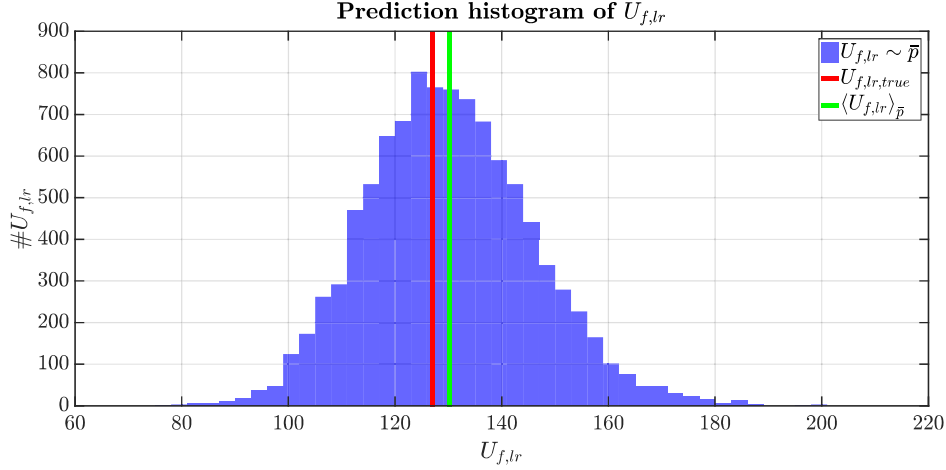


Figure 8: Predictive histogram (samples from $\bar{p}(U_{f,lr} | \boldsymbol{\lambda}_f, \boldsymbol{\theta}^*)$) for the temperature $U_{f,lr}$ of the lower right corner of the domain. The true solution $U_{f,lr,true}$ is nicely captured by the distribution \bar{p} .




Figure 9: Predictions over the whole domain on four different test samples. The true solution (colored) lies in between $\pm \sigma$ (grey). The predictive mean (blue) is very close to the true solution.

6 CONCLUSION

In this paper, we described a generative Bayesian model which is capable of giving probabilistic predictions to an expensive fine-grained model (FG) based only on a small finite number of training data and multiple solutions of a fast, but less accurate reduced order model (CG).

In particular, we consider the discretized solution of stochastic PDEs with random coefficients where the FG corresponds to a fine-scale discretization. Naturally, this comes along with a very high-dimensional vector of input uncertainties. The proposed model is capable of extracting the most relevant features of those input uncertainties and gives a mapping to a much lower dimensional space of effective properties (encoding). These lower dimensional effective properties serve as the input to the CG, which solves the PDE on a much coarser scale. The last step consists of a probabilistic reconstruction mapping from the coarse- to the fine-scale solution (decoding).

We demonstrated features and capabilities of the model proposed for a $2D$ steady-state heat problem, where the fine scale of the conductivity implies (upon discretization) a random input vector of dimension 256×256 and the solution of a discretized system of equations of comparable size. In combination with a sparsity-enforcing prior, the proposed model identified the most salient features of the fine-scale conductivity field and allowed accurate predictions of the FG response using a CGs of size only 2×2 , 4×4 and 8×8 . The predictive distribution always included the true FG solution as well as provided uncertainty bounds arising from the information loss taking place during the coarse-graining process.

REFERENCES

- [1] J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith, and M. West. Bayesian factor regression models in the large p , small n paradigm. *Bayesian statistics*, 7:733–742, 2003.
- [2] C. M. Bishop. *Neural networks for pattern recognition*. Oxford university press, 1995.
- [3] C. M. Bishop. *Latent Variable Models*, pages 371–403. Springer Netherlands, Dordrecht, 1998.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning*, volume 4. 2006.
- [5] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [6] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- [7] M. A. T. Figueiredo. Adaptive sparseness for supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1150–1159, 2003.
- [8] R. G. Gahem, P. D. Spanos, R. G. Ghanem, and P. D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. 2003.
- [9] M. D. Hoffman, D. M. Blei, C. Wang, and J. W. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- [10] M. C. Kennedy and A. O’Hagan. Predicting the output from a complex computer code when fast approximations are available, 2000.
- [11] P. Koutsourelakis. Probabilistic characterization and simulation of multi-phase random media. *Probabilistic Engineering Mechanics*, 21(3):227–234, 2006.
- [12] P.-S. Koutsourelakis. Accurate Uncertainty Quantification Using Inaccurate Computational Models. *SIAM Journal on Scientific Computing*, 31(5):3274–3300, 2009.
- [13] P.-S. Koutsourelakis and G. Deodatis. Simulation of binary random fields with applications to two-phase random media. *Journal of Engineering Mechanics*, 131(4):397–412, 2005.
- [14] X. Ma and N. Zabaras. Kernel principal component analysis for stochastic input model generation. *Journal of Computational Physics*, 230(19):7311–7331, Aug. 2011.
- [15] MATLAB. *version 9.1.0.441655 (R2016b)*. The MathWorks Inc., Natick, Massachusetts, 2016.
- [16] P. Perdikaris, D. Venturi, J. O. Royset, and G. E. Karniadakis. Multi-fidelity modelling via recursive co-kriging and Gaussian Markov random fields Subject Areas :. 2015.
- [17] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.

- [18] A. Roberts and M. Teubner. Transport properties of heterogeneous materials derived from gaussian random fields: bounds and simulation. *Physical Review E*, 51(5):4141, 1995.
- [19] M. Schöberl, N. Zabaras, and P.-S. Koutsourelakis. Predictive coarse-graining. *J. Comput. Physics*, 333:49–77, 2017.
- [20] C. Stein. Stein-1981.pdf. *The Annals of statistics*, 9(6):1135–1151, 1981.
- [21] R. Tibshirani. Royal Statistical Society. *Journal of the Royal Statistical Society B*, 58(1):267–288, 1996.
- [22] N. Tishby, F. C. N. Pereira, and W. Bialek. The information bottleneck method. *CoRR*, physics/0004057, 2000.
- [23] S. Torquato. *Random Heterogeneous Materials - Microstructure and Macroscopic Properties*. 2001.
- [24] M. J. Wainwright, M. I. Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- [25] E. Weinan. *Principles of multiscale modeling*. Cambridge University Press, 2011.
- [26] C. E. R. Williams and C. K. I. *Gaussian Processes for Machine Learning*. 2005.
- [27] H. Zou, T. Hastie, and R. Tibshirani. On the "degrees of freedom" of the lasso. *Annals of Statistics*, 35(5):2173–2192, 2007.