

Variational Bayesian strategies for high-dimensional, stochastic design problems

P.S. Koutsourelakis^{a,*}

^a*Professur für Kontinuumsmechanik, Technische Universität München, Boltzmannstrasse 15, 85747 Garching (b. München), Germany*

Abstract

This paper is concerned with a lesser-studied problem in the context of model-based, uncertainty quantification (UQ), that of optimization/design/control under uncertainty. The solution of such problems is hindered not only by the usual difficulties encountered in UQ tasks (e.g. the high computational cost of each forward simulation, the large number of random variables) but also by the need to solve a nonlinear optimization problem involving large numbers of design variables and potentially constraints. We propose a framework that is suitable for a large class of such problems and is based on the idea of recasting them as probabilistic inference tasks. To that end, we propose a Variational Bayesian (VB) formulation and an iterative VB-Expectation-Maximization scheme that is also capable of identifying a low-dimensional set of directions in the design space, along which, the objective exhibits the largest sensitivity. We demonstrate the validity of the proposed approach in the context of two numerical examples involving $\mathcal{O}(10^3)$ random and design variables. In all cases considered the cost

*Corresponding Author. Tel: +49-89-289-16690

Email address: p.s.koutsourelakis@tum.de (P.S. Koutsourelakis)

URL: <http://www.contmech.mw.tum.de> (P.S. Koutsourelakis)

of the computations in terms of calls to the forward model was of the order $\mathcal{O}(10^2)$. The accuracy of the approximations provided is assessed by appropriate information-theoretic metrics.¹

Keywords: Uncertainty Quantification, Variational Bayes, Optimization, Dimensionality reduction, Dictionary Learning

1. Introduction-Motivation

With the increased computational capabilities afforded by the utilization of peta- and exa-scale computing resources throughout engineering and the physical sciences, the issue of confidence in simulation results has come at the center of current research. The objective of obtaining a nominal computational representation of a physical process is being replaced by the new paradigm of *predictive simulations* where the analysis delivers a quantification of uncertainty due to randomness in parameters, data or models. Decisions that are based on high-fidelity computational simulations due to their potential economic or societal impact cannot be accepted without quantitative information on the confidence in the computed result.

The field of model-based, uncertainty quantification has seen marked advances in recent years. Naturally, the majority of the efforts have been directed towards *forward* uncertainty propagation i.e. the computation of output statistics given input uncertainties. While several important challenges still remain unanswered, the ultimate objective of the analysis of physical processes and engineering systems is to enable their control and optimization with respect to

¹ This paper is based on the homonymous talk given during the international symposium on "Big Data and Predictive Computational Modeling" that took place in 18-21 May 2015 at TUM-IAS, Munich Germany.

design objectives. Problems of optimization in the presence of uncertainty have attracted much less attention. On one hand, this is because they encompass all the difficulties encountered in uncertainty propagation. First and foremost the complexity of the forward problem and the increased computational expense associated with each call to the forward solver. It is generally the number of such forward solves that determines the overall computational cost. Secondly, the high-dimensionality of the vector of random variables. Especially in cases where spatiotemporal discretizations of random processes and fields are necessary, one must frequently deal with thousands of random variables. Furthermore, in stochastic optimization problems, there is the additional need to solve a demanding, nonlinear optimization problem which might itself involve thousands of design variables as well as equality/inequality constraints,

Significant advances have been achieved in deterministic optimization and control of complex systems particularly with the development of adjoint-based techniques [1, 2, 3] as well as by making use of reduced-order modeling techniques [4, 5]. Nevertheless their direct application in the stochastic counterparts of these problems would be infeasible or impractical as the integration with respect to uncertainties poses an insurmountable task.

While decision-making under uncertainty was pioneered in the 1950s [6], applications to large-scale physical models are scarce due to the inherent computational difficulties. Advances in stochastic/robust control and optimization [7, 8, 9] or reliability-based design optimization [10, 11] are generally applicable to small systems or rely on specific system structure. Techniques using surrogate models and response surfaces [12] or generalized Polynomial Chaos expansions [13] might fail to provide good approximations if the number of uncertainties

is large, irreducible or non-Gaussian. Furthermore, there is a difficulty in quantifying the error introduced due to the discrepancy between the surrogate and reference model. A critical problem in that respect is the ability to deal with noisy evaluations of the objective functions, its gradient and higher-order derivatives.

The stochastic optimization framework advocated in the present paper is motivated by the following desiderata:

- The ability to seamlessly utilize deterministic (legacy) simulators and deterministic optimization components such as a first and second order parametric derivatives of model outputs.
- The ability to deal with high-dimensional vectors of random and design variables.
- Least possible number of forward solutions
- The ability to quantify the robustness of the identified optimum and provide information on the design features that exhibit the largest sensitivity.
- The ability to utilize even highly-approximate, reduced-order models or surrogates in order to expedite the solution process.

The objective functions considered in this paper can be written in a general form as:

$$V(\mathbf{z}) = \int U(\boldsymbol{\theta}, \mathbf{z}) p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (1)$$

where $\boldsymbol{\theta} \in \mathbb{R}^{d_{\boldsymbol{\theta}}}$ denotes the vector of random variables with a probability density function $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ and $\mathbf{z} \in \mathbb{R}^{d_z}$ denotes the vector of design variables. The function $U(\boldsymbol{\theta}, \mathbf{z})$ depends on the *output* of the mathematical model and in turn, implicitly

depends on random and design variables. Each evaluation of $U(\boldsymbol{\theta}, \mathbf{z})$ implies a forward model solution which is assumed *expensive as in most challenging applications*. Naturally the optimization problem can be augmented with constraints with regards to the design variables as it will be demonstrated in the stochastic topology optimization problem that will be considered in the last section. We adopt the term *utility function* (opposite of a loss function) for $U(\boldsymbol{\theta}, \mathbf{z})$ and expected utility for $V(\mathbf{z})$ and, without loss of generality, pose the corresponding problem as one of *maximization*.

The formulation above is quite general and can be readily adapted to cases of practical interest. For example if $U(\boldsymbol{\theta}, \mathbf{z}) = \mathbf{1}_{\mathcal{A}}(\boldsymbol{\theta}, \mathbf{z})$ is the indicator function of an event \mathcal{A} of interest (e.g. failure, or exceedance of a response threshold) then maximizing $V(\mathbf{z})$ in Equation (1) is equivalent to the maximization of the probability associated with the event \mathcal{A} (similarly one can minimize the probability of event \mathcal{A} by employing the indicator function of the complementary event \mathcal{A}^c in place of U in Equation (1)). The case that would be of principal concern in this paper involves utility functions of the following form ²:

$$U(\boldsymbol{\theta}, \mathbf{z}) = \exp\left\{-\frac{1}{2} \parallel \mathbf{Q}^{1/2}(\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\theta}, \mathbf{z})) \parallel^2\right\} \quad (2)$$

where $\mathbf{u}(\boldsymbol{\theta}, \mathbf{z}) \in \mathbb{R}^n$ denotes an output vector of interest (i.e. displacements, velocities, temperature etc), $\mathbf{u}_{target} \in \mathbb{R}^n$ a target/desired response and \mathbf{Q} a positive definite matrix of choice (in the current examples $\mathbf{Q} = \tau_Q \mathbf{I}_n$). Maximizing the corresponding expected utility implies finding \mathbf{z} for which the response

²As it will become apparent in the subsequent derivations, the exponent in Equation (2) is used in order to simplify the presentation and several other options to the same effect are possible.

quantities of interest are, *on average*, as close (in the norm defined by \mathbf{Q}) to the target values \mathbf{u}_{target} . Similar objective functions have been employed by [14] to identify random composites with target effective/homogenized properties and in [15] in the context of computational mechanics. In addition, related stochastic design/control objectives have been proposed in [16] and [17].

The obvious strategy for maximizing the expected utility in Equation (1) is stochastic approximations such as noisy gradient ascent which, in its simplest form, iterates as follows:

$$\mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \eta_t \hat{\mathbf{g}}_t \quad (3)$$

where $\hat{\mathbf{g}}_t$ is a noisy (unbiased) estimator of the gradient:

$$\nabla_{\mathbf{z}} V(\mathbf{z}^{(t)}) = \int \frac{\partial U(\boldsymbol{\theta}, \mathbf{z}^{(t)})}{\partial \mathbf{z}} p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (4)$$

and η_t a sequence of learning rates that satisfy $\sum_{t=0}^{\infty} \eta_t = +\infty$, $\sum_{t=0}^{\infty} \eta_t^2 < +\infty$ [18, 19]. While convergence to a (local) maximum is assured under fairly weak conditions [20, 21] even when a single sample of $\boldsymbol{\theta}$ from $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is used in the context of a basic Monte Carlo estimate of $\hat{\mathbf{g}}_t$, the convergence rate can be slow requiring an exuberant number of forward calls to evaluate U and/or $\frac{\partial U}{\partial \mathbf{z}}$.

An alternative perspective to the problem was proposed in [22] where it was recast as a *probabilistic inference* task. In particular one defines an auxiliary probability density $p_{aux}(\boldsymbol{\theta}, \mathbf{z})$, jointly on random and design variables, as follows:

$$p_{aux}(\boldsymbol{\theta}, \mathbf{z}) \propto U(\boldsymbol{\theta}, \mathbf{z})p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \quad (5)$$

The marginal $p_{aux}(\mathbf{z}) \propto \int p_{aux}(\boldsymbol{\theta}, \mathbf{z}) d\boldsymbol{\theta}$ is clearly proportional to $V(\mathbf{z})$. If for example one could sample from the joint density $p_{aux}(\boldsymbol{\theta}, \mathbf{z})$, the \mathbf{z} -coordinates will be marginally distributed according to $V(\mathbf{z})$ and populate regions where this attains its maximum value(s).

The proposed reformulation allows for a uniform treatment of random $\boldsymbol{\theta}$ and design variables \mathbf{z} . More importantly, being able to infer $p_{aux}(\mathbf{z})$ (or a good approximation thereof) will not only lead to point estimates for the maxima of the expected utility $V(\mathbf{z})$ (which coincide with the maxima of $p_{aux}(\mathbf{z})$) but also provide valuable information about the sensitivity of the latter with respect to \mathbf{z} and therefore the *robustness* of the selected optimal design [23]. Sequential Monte Carlo strategies have been previously employed [24, 25, 23] with significant success in identifying multiple local maxima as well as utilizing approximate, surrogate models to expedite the inference task. Nevertheless the computational cost can still be significant as they potentially require a few thousand forward calls.

In this work we advocate an alternative probabilistic inference framework, namely Variational Bayes (VB) [26, 27]. Such methods have risen into prominence for probabilistic inference tasks in the machine learning community [28, 29, 30]. They provide *approximate* inference results by solving an optimization problem over a family of appropriately selected probability densities with the ob-

³For the definition of p_{aux} to be valid, it suffices that U is non-negative. The formulation can also account for U that take negative values as long as it is bounded from below i.e. $U(\boldsymbol{\theta}, \mathbf{z}) \geq U_0 > -\infty$ ($U_0 < 0$), in which case one can use $U(\boldsymbol{\theta}, \mathbf{z}) - U_0$ in place of $U(\boldsymbol{\theta}, \mathbf{z})$

jective of minimizing the Kullback-Leibler divergence [31] with the target density (in our case p_{aux}). The success of such an approach hinges upon the selection of appropriate densities that have the capacity of providing good approximations while enabling efficient (and preferably) closed-form optimization with regards to their parameters.

A pivotal role in Variational Bayesian (VB) strategies or any other inference method, is dimensionality reduction i.e. the identification of lower-dimensional features that provide the strongest signature to the random variables and associated distributions. Discovering a sparse set of features has attracted great interest in many applications as in the representation of natural images [32] and a host of algorithms have been developed not only for finding such representations but also an appropriate dictionary for achieving this goal [33]. While all these tools are pertinent to the present problem they differ in a fundamental way. They are based on several data/observations/instantiations of the vector that we seek to represent. In our problem however we do not have such direct observations i.e. the data available pertains to the output of a model which is nonlinearly and implicitly dependent on the vector of latent variables. Furthermore we are primarily interested in approximating the distribution associated with this vector rather than the dimensionality reduction itself. More importantly, only dimensionality reductions that are informative about the optimization objectives should be sought.

A premise validated in a series of papers on the so-called “sloppy” models [34] is that in several cases there exists a limited number of parameter combinations to which the outputs are sensitive. The overwhelming majority of directions are *sloppy* i.e. they embody parameter correlations that have minor influence in the

response and correspond to removable degrees of freedom. In the context of inverse problems it was found [35, 36, 37] that such features of the parameters can be associated with the eigenvectors of an appropriate Hessian or Fisher Information matrix corresponding to small eigenvalues. Along these lines and by using a fully probabilistic argumentation we develop a *reciprocal probabilistic PCA* ⁴ scheme where eigenvectors of *smallest* variance are iteratively computed and are employed not only for solving the probabilistic inference problem but for identifying the most sensitive design parameter combinations for the stochastic optimization objective.

The rest of the paper is organized as follows: The next section (Section 2) presents the essential ingredients of the VB framework advocated, the dimensionality reduction scheme proposed and an iterative, coordinate-ascent algorithm that enables the identification of all the unknowns. Section 3 demonstrates the performance and features of the proposed methodology in two problems from heat conduction and solid mechanics involving $\mathcal{O}(10^3)$ random and design variables.

2. Methodology

As discussed in the introduction we formulate the optimization-under-uncertainty problem as one of probabilistic inference. To that end our goal is two-fold. Firstly, to compute efficiently an accurate approximation of the marginal density on the design variables \mathbf{z} which provides a representation of the expected utility $V(\mathbf{z})$.

⁴ We use the term *reciprocal* to distinguish from probabilistic PCA schemes [38] where one is interested in identifying the directions with the largest variance. In contrast, in the current setting as it will be explained later on, we are interested in the directions with lowest variance

Secondly, to identify a lower-dimensional subspace with regards to the design variables \mathbf{z} that provides an assessment of the solution's robustness by discovering the most sensitive directions i.e. the directions along which, variations in \mathbf{z} will cause the largest decrease in the expected utility $V(\mathbf{z})$. Such directions have been proven useful in *deterministic* design tasks [39]. Apart from their obvious utility, they can also facilitate the inference task discussed previously. More importantly perhaps we propose a unified framework where the identification of the aforementioned lower-dimensional subspace is performed *simultaneously* with the inference of the associated densities under the same Variational Bayesian objective. This yields not only a highly efficient algorithm (in terms of the number of forward solves) but also a highly extendable framework as discussed in the conclusions.

We discuss first the parametrization advocated, identify latent variables and model parameters (Section 2.1) and subsequently demonstrate how the associated inference and learning tasks can be simultaneously addressed in the VB framework (Section 2.5). We finally present validation metrics that quantitatively assess the quality of the approximations derived (Section 2.6).

2.1. Parametrization - Dimensionality Reduction

Consider the auxiliary density $p_{aux}(\boldsymbol{\theta}, \mathbf{z})$ defined in Equation (5). This can be further extended by the introduction of an additional density $p_z(\mathbf{z})$ as follows:

$$p_{aux}(\boldsymbol{\theta}, \mathbf{z}) = \frac{U(\boldsymbol{\theta}, \mathbf{z})p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) p_z(\mathbf{z})}{Z}, \quad Z = \int U(\boldsymbol{\theta}, \mathbf{z})p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) p_z(\mathbf{z}) d\boldsymbol{\theta} d\mathbf{z} \quad (6)$$

where $p_z(\mathbf{z})$ is the analog of the regularization term in a deterministic optimization problem. In many ways Equation (6) is a restatement of Bayes' rule with

respect to θ, z :

$$p(\theta, z|data) = \frac{p(data|\theta, z) p(\theta, z)}{p(data)} \quad (7)$$

where $p_\theta(\theta)$ $p_z(z)$ play the role of the *prior*, $U(\theta, z)$ is the *likelihood* and $p_{aux}(\theta, z)$ is the *posterior*. The connection is more apparent when one considers the utility function of interest in this work (Equation (2)) in which case:

$$p_{aux}(\theta, z|u_{target}) = \frac{e^{-\frac{1}{2}\|Q^{1/2}(u_{target}-u(\theta, z))\|^2} p_\theta(\theta) p_z(z)}{Z} \quad (8)$$

Clearly the target response u_{target} is the direct analog of the “data” in Equation (7) and the role of marginal likelihood or model evidence term $p(data)$ is played by the normalization constant Z [16]. We make use of this connection frequently to motivate the modeling choices made, particularly with regards to the regularizations or priors which are terms that we use interchangeably.

The inference task in such a case would be formidable given the high dimensionality of θ, z and the cost associated with each evaluation of U as previously discussed. To address this, we propose the following decomposition of the design variables $z \in \mathbb{R}^{d_z}$:

$$\underbrace{z}_{d_z \times 1} = \underbrace{\mu_z}_{d_z \times 1} + \underbrace{W}_{d_z \times d_y} \underbrace{y}_{d_y \times 1} + \underbrace{\eta_z}_{d_z \times 1} \quad (9)$$

The motivation behind such a decomposition is quite intuitive as it resembles a Principal Component Analysis (PCA) model [38]. The vector μ_z captures the central/mean value of z , y are the reduced (and latent) coordinates of z along the linear subspace spanned by the d_y columns of the matrix W and η_z the residual “noise”. As in PCA, the premise is that a few y i.e. $d_y \ll d_z$ suffice to capture the density of z . In contrast though with PCA where the reduced

coordinates are associated with the principal directions of largest variance, the \mathbf{y} employed here should do the exact opposite i.e. identify directions with smallest variance that imply largest sensitivity. We explain this in more detail in the next Section.

The linear decomposition of a high-dimensional vector such as \mathbf{z} has received a lot of attention in several different fields. Most commonly \mathbf{z} represents a high-dimensional signal (e.g. an image, an audio/video recording) and \mathbf{W} consists of an over- or under-complete basis set [32, 40] which attempts to encode the signal as *sparingly* as possible. Significant advances in Compressed Sensing [41] or Sparse Bayesian Learning [42] have been achieved in recent years along these lines. A host of deterministic [43] or probabilistic [44] algorithms have been developed for identifying the reduced-coordinates \mathbf{y} as well as techniques for learning the most appropriate set of basis \mathbf{W} (dictionary learning) i.e. the one that can lead to the sparsest possible representation.

We adopt a simpler representation for the input random variables $\boldsymbol{\theta} \in \mathbb{R}^{d_\theta}$:

$$\underbrace{\boldsymbol{\theta}}_{d_\theta \times 1} = \underbrace{\boldsymbol{\mu}_\theta}_{d_\theta \times 1} + \underbrace{\boldsymbol{\eta}_\theta}_{d_\theta \times 1} \quad (10)$$

the usefulness of which will become apparent in the sequel. In a fully probabilistic setting all the aforementioned parameters ($\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z, \mathbf{W}, \boldsymbol{\mu}_\theta$) and the corresponding densities arising from Equation (6) would be sought. Such an inference problem would in general be formidable particularly with regards to $\boldsymbol{\mu}_z, \mathbf{W}$, whose dimension is dominated by $d_z \gg 1$. To address this difficulty we propose computing point estimates for $\boldsymbol{\mu}_z, \mathbf{W}, \boldsymbol{\mu}_\theta$ while quantifying the appropriate densities for $\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta$. We distinguish therefore between:

- the *latent variables* $\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta$.
- and *model parameters* $\mathbf{R} = \{\boldsymbol{\mu}_z, \mathbf{W}, \boldsymbol{\mu}_\theta\}$.

The computation of appropriate distributions for the latent variables and point estimates for \mathbf{R} will be addressed simultaneously under the VB framework discussed in the sequel.

2.2. Variational Bayesian approximation

Given the re-parametrization of the primal variables $\boldsymbol{\theta}, \mathbf{z}$ in Equations (9), (10), one can write the target auxiliary density as:

$$p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R}) = \frac{U(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta)p_y(\mathbf{y})p_{\eta_z}(\boldsymbol{\eta}_z)p_{\mu_z}(\boldsymbol{\mu}_z)p_W(\mathbf{W})}{Z} \quad (11)$$

where in place of the regularization $p_z(\mathbf{z})$ on \mathbf{z} we employ regularizations (priors) on the corresponding parameters $\mathbf{y}, \boldsymbol{\eta}_z$ and $\boldsymbol{\mu}_z, \mathbf{W}$ (Equation (9)). As discussed earlier rather than approximating the whole p_{aux} which would pose significant difficulties, we seek point estimates for \mathbf{R} by maximizing the (marginal) density $p_{aux}(\mathbf{R})$:

$$p_{aux}(\mathbf{R}) = \int p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R}) d\mathbf{y} d\boldsymbol{\eta}_z d\boldsymbol{\eta}_\theta \quad (12)$$

Such a maximization would amount to an analog of Maximum-A-Posteriori (MAP) estimates in a Bayesian setting.

To that end, for any density $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ on the latent variables and by em-

playing Jensen's inequality we obtain that [30]:

$$\begin{aligned}
\log p_{aux}(\mathbf{R}) &= \log \int p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R}) d\mathbf{y} d\boldsymbol{\eta}_z d\boldsymbol{\eta}_\theta \\
&= \log \int q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} d\mathbf{y} d\boldsymbol{\eta}_z d\boldsymbol{\eta}_\theta \\
&\geq \int q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) \log \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} d\mathbf{y} d\boldsymbol{\eta}_z d\boldsymbol{\eta}_\theta \\
&= \mathcal{F}(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta), \mathbf{R})
\end{aligned} \tag{13}$$

The variational lower bound \mathcal{F} given above has an intimate connection with the KL-divergence between $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ and the *conditional* density $p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R}) = \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{p_{aux}(\mathbf{R})}$ which can be expressed as:

$$\begin{aligned}
0 \leq KL(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) || p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})) &= -E_q \left[\log \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \\
&= -E_q \left[\frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{p_{aux}(\mathbf{R}) q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \\
&= \log p_{aux}(\mathbf{R}) - \mathcal{F}(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta), \mathbf{R})
\end{aligned} \tag{14}$$

We note that when $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) \equiv p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})$ the KL-divergence attains its minimum value 0, while \mathcal{F} attains its maximum value with respect to q (given $\mathbf{R} = (\boldsymbol{\mu}_z, \mathbf{W}, \boldsymbol{\mu}_\theta)$) and becomes equal to $\log p_{aux}(\mathbf{R})$. On the other hand the poorer the approximation that $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ provides to $p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})$, the larger the KL-divergence and the smaller \mathcal{F} (as a function of q) becomes.

The aforementioned discussion suggests an iterative optimization scheme that resembles the Variational Bayes - Expectation-Maximization (VB-EM) methods that have appeared in Machine Learning literature [26]. At each iteration t , one alternates between (Figure 1):

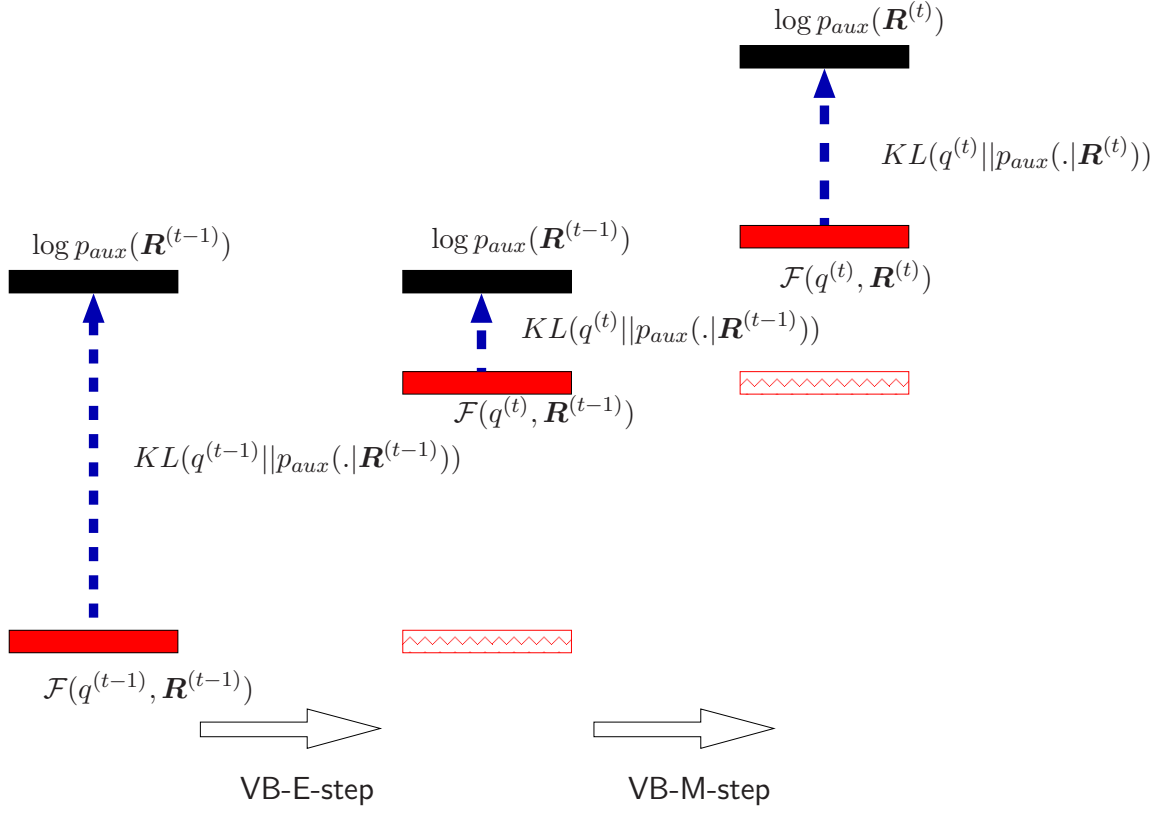


Figure 1: During the VB-E step, optimization with respect to the approximating distribution q takes place, whereas during the VB-M step, \mathcal{F} is optimized with respect to the model parameters \mathbf{R} (adapted from [45])

- **VB-Expectation:** Given $\mathbf{R}^{(t-1)}$, find:

$$q^{(t)}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) = \arg \max_q \mathcal{F}(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta), \mathbf{R}^{(t-1)}) \quad (15)$$

- **VB-Maximization:** Given $q^{(t)}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$, find:

$$\mathbf{R}^{(t)} = \arg \max_{\mathbf{R}} \mathcal{F}(q^{(t)}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta), \mathbf{R}) \quad (16)$$




Figure 2: Variational Bayesian Expectation-Maximization (VB-EM, [26])

In plain terms, the strategy advocated in order to carry out the inference task explained can be described as a generalized coordinate ascent with regards to \mathcal{F} (Figure 2).

2.3. Approximations

The variational lower bound \mathcal{F} (Equation (13)) is the objective function in the proposed scheme. In this Section we discuss its form for the utility function of interest (Equation (2)) and an isotropic $\mathbf{Q} = \tau_Q \mathbf{I}_n$. Furthermore we discuss necessary approximations that enable the VB-EM steps. We defer discussions on the validity and quantitative assessment of these approximations for Section 2.6.

In particular, from Equations (11) and (13), we have:

$$\begin{aligned}
\mathcal{F}(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta), \mathbf{R}) &= \int q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) \log \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} d\mathbf{y} d\boldsymbol{\eta}_z d\boldsymbol{\eta}_\theta \\
&= E_q \left[\frac{U(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta) p_y(\mathbf{y}) p_{\eta_z}(\boldsymbol{\eta}_z) p_{\mu_z}(\boldsymbol{\mu}_z) p_W(\mathbf{W})}{Z q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \\
&= E_q [U(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)] + E_q \left[\frac{p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta) p_y(\mathbf{y}) p_{\eta_z}(\boldsymbol{\eta}_z)}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \\
&\quad + \log p_{\mu_z}(\boldsymbol{\mu}_z) + \log p_W(\mathbf{W}) \\
&= \mathcal{F}_U + \mathcal{F}_{reg} + \log p_{\mu_z}(\boldsymbol{\mu}_z) + \log p_W(\mathbf{W})
\end{aligned} \tag{17}$$

where we distinguish the individual terms:

$$\mathcal{F}_U = E_q [U(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)] \tag{18}$$

$$\mathcal{F}_{reg} = E_q \left[\frac{p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta) p_y(\mathbf{y}) p_{\eta_z}(\boldsymbol{\eta}_z)}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \tag{19}$$

We note that in the aforementioned expressions we omit $\log Z$ as this does not depend on q nor \mathbf{R} and therefore does not affect any of the VB-EM results.

Furthermore, we note that:

$$\begin{aligned}
\mathcal{F}_U &= E_q [U(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)] \\
&= -\frac{\tau_Q}{2} E_q [| \mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) |^2]
\end{aligned} \tag{20}$$

is not only analytically intractable but also poses significant difficulties due to the computational expense associated with each forward call for the evaluation of $\mathbf{u}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)$. To alleviate these issues we propose a *linearization*

of the output vector \mathbf{u} around for $\boldsymbol{\theta} = \boldsymbol{\mu}_\theta$ and $\mathbf{z} = \boldsymbol{\mu}_z$. In particular:

$$\begin{aligned} \mathbf{u}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) &\approx \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z) + \mathbf{G}_\theta \boldsymbol{\eta}_\theta + \mathbf{G}_z (\mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) \\ &= \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z) + \mathbf{G}_\theta \boldsymbol{\eta}_\theta + \mathbf{G}_z \mathbf{W}\mathbf{y} + \mathbf{G}_z \boldsymbol{\eta}_z \end{aligned} \quad (21)$$

where $\mathbf{G}_\theta = \frac{\partial \mathbf{u}}{\partial \boldsymbol{\theta}}$, $\mathbf{G}_z = \frac{\partial \mathbf{u}}{\partial \mathbf{z}}$ evaluated at $(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)$. These derivatives can be computed using adjoint formulations when the forward model is a system of PDEs as in the examples considered in Section 3. Such a linearization will lead to a quadratic, Gauss-Newton-type, expression upon substitution in the log-utility function:

$$\begin{aligned} |\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)|^2 &\approx |\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)|^2 + \boldsymbol{\eta}_\theta^T \mathbf{G}_\theta^T \mathbf{G}_\theta \boldsymbol{\eta}_\theta \\ &\quad + \mathbf{y}^T \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z^T \mathbf{G}_z^T \mathbf{G}_z \boldsymbol{\eta}_z \\ &\quad - 2(\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z))^T (\mathbf{G}_\theta \boldsymbol{\eta}_\theta + \mathbf{G}_z \mathbf{W}\mathbf{y} + \mathbf{G}_z \boldsymbol{\eta}_z) \\ &\quad + 2\boldsymbol{\eta}_\theta^T \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z^T \mathbf{G}_z^T (\mathbf{G}_\theta \boldsymbol{\eta}_\theta + \mathbf{G}_z \mathbf{W}\mathbf{y}) \end{aligned} \quad (22)$$

We note here that a quadratic approximation could also be obtained using a 2^{nd} order Taylor series expansion of $|\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta, \boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)|^2$ directly. This would require the computation of the Hessian matrix which is also possible using adjoint formulations albeit at a significant additional cost [1]. Furthermore, for very large $d_\theta, d_z \gg 1$ the storage of the Hessian might be impractical. The reason for the quadratic approximation advocated is that it leads to closed-form expressions for the density q in the VB-Expectation step (Equation (15)) as it will become apparent in Section 2.5. Higher-order approximations would also be suitable as long as the latter requirement is satisfied.

A quadratic expression can be obtained by a 2^{nd} -order Taylor series expansion

of $\log p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta)$ around $\boldsymbol{\mu}_\theta$ without significant cost (most often than not, analytically) i.e.:

$$\log p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta) \approx \log p_\theta(\boldsymbol{\mu}_\theta) + \boldsymbol{\eta}_\theta^T \frac{\partial \log p_\theta}{\partial \boldsymbol{\theta}}|_{\boldsymbol{\mu}_\theta} + \frac{1}{2} \boldsymbol{\eta}_\theta^T \frac{\partial \log p_\theta}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T}|_{\boldsymbol{\mu}_\theta} \boldsymbol{\eta}_\theta \quad (23)$$

In the case that $p_\theta(\boldsymbol{\theta})$ is a multivariate Gaussian as in the examples considered i.e. $\mathcal{N}(\boldsymbol{\mu}_{\theta 0}, \mathbf{C}_{\theta 0})$, then the quadratic expression is exact and attains the form:

$$\log p_\theta(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta) = -\frac{1}{2}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta - \boldsymbol{\mu}_{\theta 0})^T \mathbf{C}_{\theta 0}^{-1}(\boldsymbol{\mu}_\theta + \boldsymbol{\eta}_\theta - \boldsymbol{\mu}_{\theta 0}) \quad (24)$$

2.4. Prior specification for latent variables and model parameters

The latent, reduced coordinates $\mathbf{y} \in \mathbb{R}^{d_y}$ capture the variation of \mathbf{z} around its mean $\boldsymbol{\mu}_z$ along the directions of \mathbf{W} as implied by Equation (9). It is therefore reasonable to assume that, a priori, these should have zero mean and should be uncorrelated [38]. For that purpose we adopt a multivariate Gaussian prior (denoted by $p_y(\mathbf{y})$ in the Equations of the previous section) with a diagonal covariance denoted by $\mathbf{C}_{y0} = \text{diag}(\tau_{0,i}^{-1}), i = 1, \dots, d_y$. In the examples presented in Section 3, $\tau_{0,i}$ are set to the same value τ_{y0} i.e.:

$$\mathbf{C}_{y0} = \tau_{y0}^{-1} \mathbf{I}_{d_y} \quad (25)$$

Alternatively, one can select $\tau_{0,1}^{-1} < \tau_{0,2}^{-1} < \dots \tau_{0,d_y}^{-1}$ which induces a stochastic ordering of the reduced coordinates \mathbf{y} since \mathbf{z} is invariant to permutations of the entries of the \mathbf{y} and the columns of \mathbf{W} (Equation (9)).

The remaining latent variables $\boldsymbol{\eta}_z$ account for the part of \mathbf{z} that is not captured by $\boldsymbol{\mu}_z + \mathbf{W}\mathbf{y}$ (Equation (9)) and should therefore account for the variance

in the subspace orthogonal to \mathbf{W} . The premise in the formulation advocated (Figure 3) is that \mathbf{y} capture the most sensitive directions (locally) around $\boldsymbol{\mu}_z$ which are much smaller in number than the dimensionality of \mathbf{z} i.e. $d_y \ll d_z$. The variance of \mathbf{y} should therefore be **the smallest** amongst all possible directions. The remaining directions where the variance is much larger should be captured by $\boldsymbol{\eta}_z$. We use therefore an isotropic Gaussian as a “prior” for $\boldsymbol{\eta}_z$ i.e. $p_{\boldsymbol{\eta}_z}(\boldsymbol{\eta}_z)$ is $\mathcal{N}(\mathbf{0}, \tau_{z0}^{-1}(\mathbf{I} - \mathbf{W}\mathbf{W}^T))$ ⁵ where the prior variance is set much larger than the prior variances of \mathbf{y} e.g.:

$$\tau_{z0}^{-1} = \frac{\tau_{y0}^{-1}}{\epsilon^2} \quad (26)$$

where $\epsilon^2 \ll 1$ (Section 3). We point out that this is the premise invoked also in the context of Sloppy Models [35, 36, 37], whose behavior depends only on a few stiff combinations of parameters (accounted here by \mathbf{W} and \mathbf{y}), with many sloppy parameter directions largely unimportant for model predictions (accounted here by $\boldsymbol{\eta}_z$). We also note here the fundamental difference with PCA decompositions which attain the same form as Equation(9). In PCA, \mathbf{W} and the latent variables \mathbf{y} capture the directions of **largest** variance and $\boldsymbol{\eta}_z$ account for the remaining variance which is isotropic, smaller, and superimposed on the directions \mathbf{W} .

With regards to the regularization (prior) specification $p_{\mathbf{W}}(\mathbf{W})$ on \mathbf{W} we note that its d_y columns \mathbf{w}_i , $i = 1, \dots, d_y$ span the subspace over which an approximation of \mathbf{z} is sought. We note that \mathbf{z} depends on the product $\mathbf{W}\mathbf{y}$ which would remain invariant by appropriate rescaling of each pair of $\mathbf{w}'_i = \alpha_i \mathbf{w}_i$

⁵The covariance $(\mathbf{I} - \mathbf{W}\mathbf{W}^T)$ is obviously improper as it has d_y zero eigenvalues to reflect the fact that $\boldsymbol{\eta}_z$ is inherently $(d_z - d_y)$ -dimensional

and $y'_i = \frac{1}{\alpha_i} y_i$ for any α_i . Hence, to resolve identifiability issues we require that \mathbf{W} is *orthogonal* i.e. $\mathbf{W}^T \mathbf{W} = \mathbf{I}_{d_y}$ where \mathbf{I}_{d_y} is the d_y -dimensional identity matrix. This is equivalent to employing a uniform prior on \mathbf{W} on the Stiefel manifold $V_{d_y}(\mathbb{R}^{d_z})$ [46].

The final aspect of the prior model pertains to $p_{\mu_z}(\boldsymbol{\mu}_z)$. As this is closely related to the physical meaning of the design variables \mathbf{z} we make this specific for each of the examples considered in Section 3.

2.5. VB- Expectation-Maximization: Update equations for $q(\boldsymbol{\eta}_\theta, \mathbf{y}, \boldsymbol{\eta}_z)$ and \mathbf{R}

We consider Gaussian families of approximating distributions $q(\boldsymbol{\eta}_\theta, \mathbf{y}, \boldsymbol{\eta}_z)$ with the following mean and covariance characteristics:

$$\begin{bmatrix} E_q[\boldsymbol{\eta}_\theta] = \mathbf{0} \\ E_q[\mathbf{y}] = \mathbf{0} \\ E_q[\boldsymbol{\eta}_z] = \mathbf{0} \end{bmatrix}, \quad \begin{bmatrix} E_q[\boldsymbol{\eta}_\theta \boldsymbol{\eta}_\theta^T] = \mathbf{C}_{\theta\theta} & E_q[\boldsymbol{\eta}_\theta \mathbf{y}^T] = \mathbf{C}_{\theta y} & E_q[\boldsymbol{\eta}_\theta \boldsymbol{\eta}_z^T] = \mathbf{0} \\ E_q[\mathbf{y} \boldsymbol{\eta}_\theta^T] = \mathbf{C}_{\theta y}^T & E_q[\mathbf{y} \mathbf{y}^T] = \mathbf{C}_{yy} & E_q[\mathbf{y} \boldsymbol{\eta}_z^T] = \mathbf{0} \\ E_q[\boldsymbol{\eta}_z \boldsymbol{\eta}_\theta^T] = \mathbf{0} & E_q[\boldsymbol{\eta}_z \mathbf{y}^T] = \mathbf{0} & E_q[\boldsymbol{\eta}_z \boldsymbol{\eta}_z^T] = \tau_z^{-1}(\mathbf{I} - \mathbf{W} \mathbf{W}^T) \end{bmatrix} \quad (27)$$

This form is postulated for the following reasons:

- $\boldsymbol{\eta}_\theta$ expresses variations of $\boldsymbol{\theta}$ from its mean $\boldsymbol{\mu}_\theta$ (Equation (10)) and should therefore have a mean zero.
- \mathbf{y} and $\boldsymbol{\eta}_z$ express variations of \mathbf{z} from its mean $\boldsymbol{\mu}_z$ (Equation (9)) and should also have a mean zero.
- $\boldsymbol{\eta}_z$ expresses residual variation (noise) of \mathbf{z} from $\boldsymbol{\mu}_z + \mathbf{W} \mathbf{y}$. Apart from having a zero mean, it is assumed to be uncorrelated with $\boldsymbol{\eta}_\theta$ as well as \mathbf{y} .
- $\boldsymbol{\eta}_z$ accounts for variance in the subspace orthogonal to \mathbf{W} . Along this it is assumed that the variance is isotropic and equal to τ_z^{-1} (to be determined)

Before embarking in the presentation of the expressions for the aforementioned parameters, we note that q provides an approximation to p_{aux} and therefore its marginal with respect to $\mathbf{y}, \boldsymbol{\eta}_z$ can be used to approximate the marginal on \mathbf{z} i.e. $p_{aux}(\mathbf{z})$ which is proportional to the expected utility $V(\mathbf{z})$. We note that based on Equation (27), the marginal $q(\mathbf{y}, \boldsymbol{\eta}_z)$ will also be a Gaussian and there the approximate $p_{aux}(\mathbf{z})$ will be a Gaussian with the following mean and covariance:

$$E[\mathbf{z}] = \boldsymbol{\mu}_z, \quad E[\mathbf{z}\mathbf{z}^T] = \mathbf{C}_{zz} = \mathbf{W}\mathbf{C}_{yy}\mathbf{W}^T + \tau_z^{-1}(\mathbf{I} - \mathbf{W}\mathbf{W}^T) \quad (28)$$

We can therefore approximate (up to a multiplicative constant) the expected utility $V(\mathbf{z})$ as:

$$V(\mathbf{z}) \approx V(\boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) \propto e^{-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu}_z)^T \mathbf{C}_{zz}^{-1}(\mathbf{z}-\boldsymbol{\mu}_z)} \quad (29)$$

Let $\hat{\mathbf{w}}_j$, $j = 1, \dots, d_z$ denote the eigenvectors of \mathbf{C}_{zz} and $\sigma_1 < \sigma_2^2 < \dots < \sigma_{d_z}^2$ the corresponding eigenvalues in ascending order (Figure 3). Consider (local) variations $\Delta \mathbf{z}_j$ of \mathbf{z} from $\boldsymbol{\mu}_z$ along the distinct directions $\hat{\mathbf{w}}_j$ i.e.:

$$\Delta \mathbf{z}_j = \mathbf{z} - \boldsymbol{\mu}_z = \alpha \hat{\mathbf{w}}_j \quad (30)$$

Then:

$$V(\Delta \mathbf{z}_1) \propto e^{-\frac{1}{2} \frac{\alpha^2}{\sigma_1^2}} < V(\Delta \mathbf{z}_2) \propto e^{-\frac{1}{2} \frac{\alpha^2}{\sigma_2^2}} < \dots < V(\Delta \mathbf{z}_{d_z}) \propto e^{-\frac{1}{2} \frac{\alpha^2}{\sigma_{d_z}^2}} \quad (31)$$

Hence the expected utility of competing designs \mathbf{z}_j will decrease faster for variations along directions with the *smaller* variances/eigenvalues which represent the

directions of higher sensitivity.

The methodology developed is based on the postulate that the number of *sensitive directions* is small compared to d_z and can be captured by the $d_y \ll d_z$ latent variables \mathbf{y} and the vectors in \mathbf{W} . Most of the remaining directions are assumed to be sloppy i.e. have a much higher variance and can be represented by $\boldsymbol{\eta}_z$. The d_y sensitive directions can be found by diagonalizing $\mathbf{C}_{yy} = \mathbf{U} \text{diag}(\sigma_{1 \div d_y}^2) \mathbf{U}^T$ and as a result:

$$\hat{\mathbf{W}} = \begin{bmatrix} \hat{\mathbf{w}}_1 & \hat{\mathbf{w}}_2 & \dots & \hat{\mathbf{w}}_{d_y} \end{bmatrix} = \mathbf{W} \mathbf{U} \quad (32)$$

From the expressions of $\mathcal{F}_U, \mathcal{F}_{reg}$ in Equations (17) (18), (19) and the approximations in Equations (22) and (24), we obtain (up to additive constants):

$$\begin{aligned} \mathcal{F}(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta), \mathbf{R}) = & -\frac{\tau_Q}{2} (|\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)|^2 & (\text{from } E_q[U]) \\ & + \mathbf{G}_\theta^T \mathbf{G}_\theta : \mathbf{C}_{\theta\theta} + \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{yy} + \tau_z^{-1} \mathbf{G}_z^T \mathbf{G}_z : (\mathbf{I} - \mathbf{W} \mathbf{W}^T) \\ & + 2 \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{\theta y}) \\ & - \frac{1}{2} (\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0})^T \mathbf{C}_{\theta 0}^{-1} (\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0}) & (\text{from } E_q[p_\theta]) \\ & - \frac{1}{2} \mathbf{C}_{\theta 0}^{-1} : \mathbf{C}_{\theta\theta} \\ & - \frac{\tau_{y0}}{2} \mathbf{I} : \mathbf{C}_{yy} & (\text{from } E_q[p(\mathbf{y})]) \\ & - \frac{d_z - d_y}{2} \frac{\tau_{z0}}{\tau_z} & (\text{from } E_q[p(\boldsymbol{\eta}_z)]) \\ & + \frac{1}{2} \log \begin{vmatrix} \mathbf{C}_{\theta\theta} & \mathbf{C}_{\theta y} \\ \mathbf{C}_{\theta y}^T & \mathbf{C}_{yy} \end{vmatrix} & (\text{from } E_q[q(\boldsymbol{\eta}_\theta, \mathbf{y})]) \\ & + \frac{d_z - d_y}{2} \log \tau_z & (\text{from } E_q[q(\boldsymbol{\eta}_z)]) \\ & + \log p_{\mu_z}(\boldsymbol{\mu}_z) + \log p_W(\mathbf{W}) \end{aligned} \quad (33)$$




Figure 3: Illustration in two dimensions ($d_z = 2$) of $p_{aux}(z)$ (which approximates expected utility $V(z)$) in the vicinity of (local) maximum μ_z . The most sensitive/stiff directions are captured by \mathbf{W} and their variance σ_1^2 is accounted by \mathbf{y} . The most insensitive/sloppy directions are orthogonal to \mathbf{W} (along \mathbf{W}^\perp) and their variance $\sigma_2^2 \gg \sigma_1^2$ is accounted by η_z .

The iterative VB-EM optimization with regards to q and \mathbf{R} (Figure 2) proceeds then as follows:

- **VB-Expectation:** Given the current \mathbf{R} find the optimal q (i.e. the optimal $C_{\theta\theta}, C_{\theta y}, C_{yy}, \tau_z$):

$$\begin{bmatrix} C_{\theta\theta}^{opt} & C_{\theta y}^{opt} \\ sym. & C_{yy}^{opt} \end{bmatrix}^{-1} = \begin{bmatrix} \tau_Q \mathbf{G}_\theta^T \mathbf{G}_\theta + C_{\theta\theta}^{-1} & \tau_Q \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} \\ sym. & \tau_Q \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} + \tau_{y0} \mathbf{I} \end{bmatrix} \quad (34)$$

and:

$$\tau_z^{opt} = \tau_{z0} + \frac{1}{d_z - d_y} \tau_Q \mathbf{G}_z^T \mathbf{G}_z : (\mathbf{I} - \mathbf{W} \mathbf{W}^T) \quad (35)$$

- **VB-Maximization:** Given the current q (i.e. $\mathbf{C}_{\theta\theta}, \mathbf{C}_{\theta y}, \mathbf{C}_{yy}, \tau_z$), find the optimal $\mathbf{R} = \{\boldsymbol{\mu}_z, \mathbf{W}, \boldsymbol{\mu}_\theta\}$. To carry out this task, it suffices to consider only the terms of \mathcal{F} in Equation (33), that depend on the parameters of interest i.e.:

$$(\boldsymbol{\mu}_z^{opt}, \boldsymbol{\mu}_\theta^{opt}) = \arg \max_{\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta} \hat{\mathcal{F}}_\mu(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta), \quad \mathbf{W}^{opt} = \arg \max_{\mathbf{W}} \hat{\mathcal{F}}_W(\mathbf{W}) \quad (36)$$

where:

$$\begin{aligned} \mathcal{F}_\mu(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta) = & -\frac{\tau_Q}{2} (|\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)|^2) \\ & -\frac{1}{2}(\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0})^T \mathbf{C}_{\theta 0}^{-1}(\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0}) + \log p_{\mu_z}(\boldsymbol{\mu}_z) \end{aligned} \quad (37)$$

and:

$$\begin{aligned} \mathcal{F}_W(\mathbf{W}) = & -\frac{\tau_Q}{2} \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{yy} - \frac{\tau_Q}{2} \tau_z^{-1} \mathbf{G}_z^T \mathbf{G}_z : (\mathbf{I} - \mathbf{W} \mathbf{W}^T) \\ & -\frac{\tau_Q}{2} 2 \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{\theta y} + \log p_W(\mathbf{W}) \\ = & -\frac{\tau_Q}{2} \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} : (\mathbf{C}_{yy} - \tau_z^{-1} \mathbf{I}) \\ & -\frac{\tau_Q}{2} 2 \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{\theta y} + \log p_W(\mathbf{W}) \end{aligned} \quad (38)$$

Some remarks are warranted at this stage:

- The maximization of \mathcal{F}_μ with respect to $(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)$ can be carried out using any nonlinear optimization scheme. We discuss in Appendix B a Gauss-Newton-type scheme which requires only first-order derivatives of \mathbf{u} . We note that this is the only part of the VB-EM scheme proposed that requires

calls to the forward solver for the computation of \mathbf{u} and its derivatives. Hence in the context of large-scale, complex models this step controls, to a large extent, the overall cost of the proposed algorithm.

- The updates of $(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)$ in Equation (37) are *decoupled* from the rest i.e. q and \mathbf{W} . This is a direct consequence of the assumption on q that $E_q[\mathbf{y}] = E_q[\boldsymbol{\eta}_z] = E_q[\boldsymbol{\eta}_\theta] = \mathbf{0}$ (Equation (27)) which was described earlier. As a result, the optimal $(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)$ can be computed *beforehand* and the rest of the VB-EM steps would involve only iterative updates of q (i.e. $\mathbf{C}_{\theta\theta}, \mathbf{C}_{\theta y}, \mathbf{C}_{yy}, \tau_z$) and \mathbf{W} .
- The optimization of \mathcal{F}_W with regards to the orthogonal matrix \mathbf{W} requires appropriate nonlinear constrained optimization tools. A highly efficient such tool is discussed in Appendix A. We reiterate that this step does not require any further calls to the forward solver.
- We also point out that the proposed algorithm inherits all the favorable traits of Expectation-Maximization algorithms as discussed in [47]. As explained therein it suffices that the updates for q (VB-E-step) or \mathbf{R} (VB-M-step) lead to an improvement of the variational bound \mathcal{F} rather than being the locally optimally values. This would for example allow for only partial updates of q (e.g. updating $\mathbf{C}_{\theta\theta}$ only every few iterations) or incremental improvements of \mathbf{W} that simply lead to an increase in \mathcal{F}_W without finding the local maximum. Such strategies could expedite significantly the computations involved.
- Finally we note that implicit to the aforementioned derivations is the assumption of a unimodal density on the latent variables and as a result

a unique, global maximum for z (Equation (28)). This assumption can be relaxed by employing a mixture of Gaussians (e.g. [48]) that will enable the approximation of highly non-Gaussian and potentially multi-modal p_{aux} which in turn can reveal multiple *local maxima* of the expected utility $V(z)$. Such approximations could also be combined with the employment of different basis sets \mathbf{W} for each of the mixture component i.e. different sensitive/sloppy directions for each local optimum. We defer further discussions along these lines to future work.

2.6. Validation - Assessing the accuracy of approximations

Thus far we have employed the variational lower bound in order to identify the optimal dimensionality reduction and to infer the latent variables that approximate $p_{aux}(z)$ (Equation (28)) and the expected utility $V(z)$ (Equation (29)). The goal in this section is to propose quantitative indicators that assess the accuracy of the VB approximation. To that end we consider the Kullback-Leibler divergence $KL(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) || p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R}))$ (Equation (14)) that motivated the VB-EM scheme discussed. In particular:

$$\begin{aligned}
KL(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) || p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})) &= -E_q \left[\log \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \\
&= -E_q \left[\frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})}{p_{aux}(\mathbf{R}) q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \\
&= \log p_{aux}(\mathbf{R}) - E_q \left[\frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})}{p_{aux}(\mathbf{R}) q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right]
\end{aligned} \tag{39}$$

where $p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})$ is given in Equation (11) and $\log p_{aux}(\mathbf{R})$ in Equation (12). We propose estimating both terms in Equation (39) using Importance Sampling with $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ as the Importance Sampling density [26]. If we denote

by:

$$w(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) = \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)}. \quad (40)$$

the (un-normalized) importance weights, then by drawing samples $\{\mathbf{y}^{(m)}, \boldsymbol{\eta}_z^{(m)}, \boldsymbol{\eta}_\theta^{(m)}\}_{m=1}^M$ from $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ we obtain that:

$$\log \langle w \rangle = \log \left(\frac{1}{M} \sum_{m=1}^M w(\mathbf{y}^{(m)}, \boldsymbol{\eta}_z^{(m)}, \boldsymbol{\eta}_\theta^{(m)}) \right) \longrightarrow \log p_{aux}(\mathbf{R}) \quad (41)$$

and:

$$\langle \log w \rangle = \frac{1}{M} \sum_{m=1}^M \log w(\mathbf{y}^{(m)}, \boldsymbol{\eta}_z^{(m)}, \boldsymbol{\eta}_\theta^{(m)}) \longrightarrow E_q \left[\log \frac{p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta, \mathbf{R})}{q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)} \right] \quad (42)$$

In summary, by employing Importance Sampling we can estimate:

$$KL(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) || p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R})) \approx \log \langle w \rangle - \langle \log w \rangle \quad (43)$$

We note that sampling from $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ is straightforward due to its Gaussian form but the evaluation of the weights require the computation of the actual utility i.e. running the exact forward model. We point out however that this is done *solely* for the purposes of validation. Given that the KL -divergence is not bounded from above and in order to compare it when considering various values of d_y i.e. the dimension of the reduced coordinates y , we propose normalizing it with the entropy $H(q)$ of the multivariate Gaussian $q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta)$ which can be

exactly computed as:

$$H(q) = -\frac{d_\theta + d_y}{2} \log 2\pi - \frac{1}{2} \log \left| \begin{array}{cc} \mathbf{C}_{\theta\theta}^{opt} & \mathbf{C}_{\theta y}^{opt} \\ sym. & \mathbf{C}_{yy}^{opt} \end{array} \right| - \frac{d_z - d_y}{2} \log \frac{2\pi}{\tau_z^{opt}} \quad (44)$$

In the examples that follow we report therefore the following *normalized* KL-divergence:

$$nKL = \frac{KL(q(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta) || p_{aux}(\mathbf{y}, \boldsymbol{\eta}_z, \boldsymbol{\eta}_\theta | \mathbf{R}))}{H(q)} \quad (45)$$

In all the expressions above we use \mathbf{R}^{opt} as found at the last iteration of the VB-EM scheme.

3. Numerical Illustrations

In this section we discuss the numerical results obtained in the analysis of two examples. In both cases the forward model consisted of an elliptic PDE. The discretization of the forward problem for the computation of outputs \mathbf{u} as well as of the adjoint problem for the computation of the derivatives $\mathbf{G}_\theta, \mathbf{G}_z$ was performed using standard finite element tools. Both problems involved a very high number of random variables d_θ arising from the discretization of a random field with small correlation length (in relation to the problem domain). Especially the second example involved a very large number of design variables d_z . A summary of the basic dimensions/quantities is contained in Table 1.

With regards to the regularization (prior) terms, in both problems we employed $\tau_{y0}^{-1} = 10^4$ (Equation (25)) and $\epsilon^2 = 10^{-10}$ (Equation (26)). Details about the $p_{\mu_z}(\boldsymbol{\mu}_z)$ are given for each example separately. With regards to the VB-EM scheme employed we note that at each iteration, 100 \mathbf{W} -updates were performed according to the equations detailed in Appendix A.

	Random Variables θ	Design Variables z	τ_Q^{-1}	$n = \dim(\mathbf{u}_{target})$
Num. Illustration 1	$d_\theta = 1600$	$d_z = 21$	0.01	11
Num. Illustration 2	$d_\theta = 3536$	$d_z = 3536$	5×10^{-6}	8

Table 1: Basic quantities/dimensions

3.1. Numerical Illustration 1

The goal of this problem is to optimally select the input to a random, heterogeneous medium so as to maximize an expected utility related to the response. In particular we consider the rectangular domain $\Omega = [-1, 1] \times [0, 1]$ of Figure 4 and the steady-state heat diffusion with a governing PDE:

$$\nabla \cdot (-\lambda(\mathbf{x}) \nabla u(\mathbf{x})) = 0, \quad \mathbf{x} \in \text{int}(\Omega) \quad (46)$$

The boundary conditions are $u = 0$ on Γ_D , $-\lambda(\mathbf{x}) \frac{\partial T(\mathbf{x})}{\partial n} = 0$ on Γ_N . The design variables z parametrize the flux on the left hand boundary.

Figure 4: Problem Configuration for Numerical Illustration 1

The uncertainties θ parametrize the conductivity field $\lambda(\mathbf{x})$. In particular we consider a statistically-homogeneous, log-normally-distributed random field with mean 1 and coefficient of variation 0.50. This is defined through a transformation of a statistically-homogeneous Gaussian field $\lambda_g(\mathbf{x})$ as:

$$\lambda(\mathbf{x}) = e^{\lambda_g(\mathbf{x})} \quad (47)$$

The following autocovariance $C_g(\Delta x_1, \Delta x_2)$ for $\lambda_g(\mathbf{x})$ is employed:

$$C_g(\Delta x_1, \Delta x_2) = \sigma_g^2 \exp\left\{-\frac{\sqrt{\Delta x_1^2 + \Delta x_2^2}}{x_0}\right\}, \quad \sigma_g^2 = 0.223 \quad (48)$$

- | | |
|----------------|---------------|
| (i) sample 1 | (ii) sample 2 |
| (iii) sample 3 | (iv) sample 4 |

Figure 5: Sample realizations of the conductivity (example 1) - Young's modulus (example 2) field $\lambda(x)$ as prescribed in Equation (47)

where a correlation length of $x_0 = 0.1$ is used. We note that the correlation length is small in relation to the dimensions of the problem domain and as a result a large number of random variables θ are required. In particular we discretize the problem domain into 1600 triangular, finite elements ⁶ and model with θ the value of $\lambda_g(x)$ at the centroid of each element. This gives rise to $d_\theta = 1600$ and a p_θ (Equation (24)) with mean $\mu_{\theta 0} = -0.112$, variance $\sigma_g^2 = 0.223$ and covariance matrix $C_{\theta 0}$ obtained from Equation (48). Sample realizations of the conductivity field $\lambda(x)$ are depicted in Figure 5 for illustrative purposes.

We employ a design variable z for each node along the left-hand boundary of the problem domain (Figure 4) resulting in $d_z = 21$ design variables. Finally, with regards to the utility function U , we use temperatures along $x_1 = 0, x_2 \in [-.25, 0.75]$ (red line in Figure 4) and in particular at 11 equidistant points with $x_{2,k} = 0.25 + 0.05(k - 1), k = 1 \div 11$. The target temperature vector u_{target} is set to:

$$u_{target,k} = 20 - 40|x_{2,k} - 0.5| \quad (49)$$

and $\tau_Q^{-1} = 0.01$ (Equation (8)). We finally note that a vague Gaussian regularization/prior was employed for μ_z such that $p_{\mu_z}(\mu_z) \equiv \mathcal{N}(\mathbf{0}, C_{z0} = 10^{10} \mathbf{I})$.

Figure 6 depicts the computed μ_z as a function of the number of iterations

⁶we consider a 40×20 regular grid and each rectangle is divided along its diagonal into two triangles

Figure 6: Computed μ_z (see Appendix B) as a function of the iteration number. In example 1 this expresses the flux on the left boundary $x_1 = 0$, $x_2 \in [0, 1]$. Each iteration involves a forward call for the computation of the output \mathbf{u} and its derivatives.

(Appendix B). As it can be seen, convergence is attained with as few as 20 forward calls. We re-emphasize that these are the only forward solutions required for the computation of the outputs and their derivatives. We note that while the linearization in Equation (21) with respect to \mathbf{z} is exact, this is not the case with regards to the random variables $\boldsymbol{\theta}$. This is due primarily to the nonlinear dependence of the response on the conductivity field $\lambda(\mathbf{x})$.

The evolution of the the variational lower-bound \mathcal{F} (Equation (33)) with regards to the iterations alternating between q and W updates is shown in Figure 7i. We note that these iterations *do not entail any additional forward calls*. Figure 7ii depicts the evolution of the identified σ_j^2 per VB-EM iteration where as it is clearly seen, there exist 3 “stiff” generalized eigenvectors with small values for the corresponding generalized eigenvalues. One also notes that the variances top-off at the prior value $\tau_{y0}^{-1} = 10^4$. These 3 most sensitive generalized eigenvectors $\hat{\mathbf{W}}$ (Equation (32)) and the associated variances are shown in Figure 8. The numbers in parentheses were the computed variances when the calculation was repeated for exactly the same problem but by assuming a coefficient of variation of $0.71 = \sqrt{0.5}$ (instead of 0.50) for the conductivity field $\lambda(\mathbf{x})$. The most sensitive eigenvectors were identical (and therefore not plotted) but, as expected, their sensitivity is reduced or equivalently the corresponding variances were larger. Figure 9 compares the μ_z computed for these two cases where one notes that while the shape is the same the amplitude/range is different.

(i) Evolution of \mathcal{F} (Equation (33)).

(ii) Evolution of σ_j^2

Figure 7: VB-EM Each iteration corresponds to one q (Equations (34), (35)) and one \mathbf{W} (Equation (38)) update

(i) $\sigma_1^2 = 4.0 \times 10^{-2}$
(5.9×10^{-2})

(ii) $\sigma_2^2 = 1.5 \times 10^3$
(1.6×10^3)

(iii) $\sigma_3^2 = 6.6 \times 10^3$
(6.8×10^3)

Figure 8: First three most sensitive eigenvectors $\{\hat{\mathbf{w}}_j\}_{j=1}^3$ (Equation (32)) and associated variances σ_j^2 . We note that $\sigma_3^2/\sigma_1^2 = \mathcal{O}(10^5)$

Figure 10 depicts sample designs drawn from $q(\boldsymbol{\mu}_z + \mathbf{W}y)$ (which approximates the expected utility $V(\boldsymbol{\mu}_z + \mathbf{W}y)$) corresponding to different (relative) levels of the the expected utility. While in the approximation advocated $\boldsymbol{\mu}_z$ represents the optimal design for which $V(z)$ attains its (locally) maximum value, by considering expected utility values $V(z)$ less than the optimal we can identify an infinity of alternative designs but also assess the sensitivity of the solution.

Finally in Table 2 we record the normalized KL-divergence as discussed in Section 2.6 and note that this decays for increasing d_y to relatively small values indicating a good quality in the approximation found.

3.2. Numerical Illustration 2: Stochastic Topology Optimization

The vast majority of studies in the context of stochastic topology optimization consider uncertainties in the loads (i.e. input) of linear systems [49]. This allows one to find closed-form expressions for the random response and perform the integrations needed much more easily. Recently notable efforts have been

Figure 9: Comparison of $\boldsymbol{\mu}_z$ computed when the conductivity field $\lambda(\mathbf{x})$ has a coefficient of variation (cov) of 0.50 and 0.71.

$$(i) \frac{V(z)}{V(\mu_z)} = 0.95$$

$$(ii) \frac{V(z)}{V(\mu_z)} = 0.75$$

$$(iii) \frac{V(z)}{V(\mu_z)} = 0.50$$

$$(iv) \frac{V(z)}{V(\mu_z)} = 0.25$$

Figure 10: Alternative designs z at various levels of expected utility $\frac{V(z)}{V(\mu_z)}$ as compared to the optimal μ_z

d_y	nKL (Equation (45))
1	1.5×10^{-1}
2	1.2×10^{-1}
5	4.7×10^{-2}
10	2.5×10^{-2}
20	9.8×10^{-3}

Table 2: Normalized KL-divergence from Equation (45) for example 1

made towards addressing the significantly more complicated problem involving geometric and/or material uncertainties [50]. Some of the proposed solution strategies employed perturbations techniques [51, 52] whose performance decays as the random variability around the mean and/or the number of random variables increases. Other attempts have made use of intrusive [53] and non-intrusive [54, 55] versions of (generalized) Polynomial Chaos (gPC) in order to address the stochastic components.

We consider the two-dimensional domain $\Omega = [0, 1.6] \times [0, 1]$ in Figure 11i. The goal is to identify where the material of interest should be placed in order to achieve the objectives (subject to appropriate constraints) to be discussed. We can therefore partition Ω into Ω_1 which contains all points where material is placed and $\Omega_0 = \Omega \setminus \Omega_1$ which corresponds to the points without any material

(void). The governing differential is that of elastostatics:

$$\begin{aligned} \nabla \cdot (\mathbf{D}(\mathbf{x})\boldsymbol{\epsilon}(\mathbf{u}(\mathbf{x}))) &= \mathbf{0}, \quad \mathbf{x} \in \text{int}(\Omega) \\ \boldsymbol{\epsilon}(\mathbf{u}(\mathbf{x})) &= \begin{bmatrix} \frac{\partial u_1}{\partial x_1} \\ \frac{\partial u_2}{\partial x_2} \\ \frac{\partial u_1}{\partial x_2} + \frac{\partial u_2}{\partial x_1} \end{bmatrix} \end{aligned} \quad (50)$$

where $\mathbf{u}(\mathbf{x}) = \begin{bmatrix} u_1(x_1, x_2) \\ u_2(x_1, x_2) \end{bmatrix}$ is the displacement field, \mathbf{D} is the (plane-stress)

elasticity matrix⁷ i.e. $\mathbf{D}(\mathbf{x}) = \frac{E(\mathbf{x})}{1-\nu^2} \begin{bmatrix} 1 & \nu & 0 \\ \nu & 1 & 0 \\ 0 & 0 & 1-\nu \end{bmatrix}$ and $E(\mathbf{x})$ is the Young's

modulus. Its spatial variation can be modeled as:

$$E(\mathbf{x}) = E_{min} + 1_{\Omega_1}(\mathbf{x})(\lambda(\mathbf{x}) - E_{min}), \quad 1_{\Omega_1}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in \Omega_0 \\ 1 & \text{if } \mathbf{x} \in \Omega_1 \end{cases} \quad (51)$$

The value of $E_{min} = 10^{-10}$ (instead of 0) is used to avoid numerical issues in the solution of the governing equations. With regards to boundary conditions it is assumed that $\mathbf{u} = \mathbf{0}$ along Γ_D and traction-free along Γ_N (Figure 11i) with the exception of a point force $P = 10^{-3}$ at $(x_1 = 1.6, x_2 = 0)$.

In deterministic formulations, $\lambda(\mathbf{x})$ is assumed constant. In the context of the analysis pursued in this study we are interested in exploring the case where $\lambda(\mathbf{x})$ not only varies spatially but also exhibits *stochastic variability* i.e. $\lambda(\mathbf{x})$ is a random field. The model adopted for $\lambda(\mathbf{x})$ is identical to that in Example 1

⁷ $\nu = 0.3$ (constant) in this study

which we repeat here for completeness. In particular we define $\lambda(\mathbf{x})$ through a transformation of a statistically-homogeneous, Gaussian random field $\lambda_g(\mathbf{x})$ as in Equation (47). The latter has a mean (constant) $\mu_g = -0.112$ and autocovariance $C_g(\Delta x_1, \Delta x_2)$ as prescribed in Equation (48) with a correlation length $x_0 = 0.1$ and a variance $\sigma_g^2 = 0.223$. This gives rise to a log-normally distributed $\lambda(\mathbf{x})$ with mean 1 and coefficient of variation 0.50.

The problem domain Ω is discretized using a regular mesh of 3536 triangular elements ⁸. The vector of random variables $\boldsymbol{\theta}$ represents the values of $\lambda_g(\mathbf{x})$ at the centroid of each element. This gives rise to $d_\theta = 3536$ and a p_θ (Equation (24)) with mean $\boldsymbol{\mu}_{\theta 0} = -0.112$, variance $\sigma_g^2 = 0.223$ and covariance matrix $\mathbf{C}_{\theta 0}$ obtained from Equation (48). We note that, as in Example 1, a small correlation length is selected giving rise to a large number of random variables $\boldsymbol{\theta}$.

Normally the design variables \mathbf{z} should be binary and discretize the indicator function $1_{\Omega_1}(\mathbf{x})$ in Equation (51) ⁹. As in deterministic topology optimization schemes [58] and in order to be able to compute meaningful derivatives with respect to the design variables we adopt a relaxation of the problem. In order to represent the variations of the elastic modulus $E(\mathbf{x})$ (Equation (51)), we employ the sigmoid function to transform a real-valued field $z(\mathbf{x})$ as follows:

$$E(\mathbf{x}) = E_{min} + \frac{1}{1 + e^{-z(\mathbf{x})}}(\lambda(\mathbf{x}) - E_{min}) \quad (52)$$

While the sigmoid function ensures that $E(\mathbf{x}) \in [E_{min}, \lambda(\mathbf{x})]$ as in Equation

⁸we consider a 52×34 regular grid and each rectangle is divided along its diagonal into two triangles

⁹We note that in deterministic formulations level-set-based representation have also been adopted e.g. [56, 57]

(51) it does not necessarily yield a *hard* partitioning ($0 - 1$) of Ω as required in such problems. To achieve this i.e. to promote solutions where $z(\mathbf{x}) \rightarrow -\infty$ (i.e. $E(\mathbf{x}) \rightarrow E_{min}$) or $z(\mathbf{x}) \rightarrow +\infty$ (i.e. $E(\mathbf{x}) \rightarrow \lambda(\mathbf{x})$) we adopt an appropriate *hierarchical prior/regularization* $p_z(\mathbf{z})$ that is discussed in detail in Appendix C. Naturally the vector of design variables \mathbf{z} represents the values of the field $z(\mathbf{x})$ at the centroid of each finite element (as we did for the random variables $\boldsymbol{\theta}$) resulting in $d_z = 3536$ *design variables* (Table 1).

More importantly though the problem formulation is only meaningful with the introduction of a constraint on the volume of material that should be used i.e. the volume fraction $VF = \frac{area(\Omega_1)}{area(\Omega)}$. This in turn implies an *equality constraint* for the design variables \mathbf{z} which can be written as:

$$c(\mathbf{z}) = \frac{1}{d_z} \sum_{j=1}^{d_z} \frac{1}{1 + e^{-z_j}} - VF = 0, \quad (53)$$

where VF is the targeted volume fraction¹⁰. In order to account for this nonlinear constraint in the proposed framework where the design variables \mathbf{z} are treated as random variables, we propose expanding the target, auxiliary $p_{aux}(\boldsymbol{\theta}, \mathbf{z})$ (Equation (6)) as follows:

$$p_{aux}(\boldsymbol{\theta}, \mathbf{z}) \propto e^{-\frac{c^2(\mathbf{z})}{2\epsilon_c^2}} U(\boldsymbol{\theta}, \mathbf{z}) p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) p_z(\mathbf{z}) \quad (54)$$

Clearly this represents a *soft, probabilistic* enforcement of the aforementioned constraint where for small ϵ_c^2 , the target density p_{aux} , and therefore the associated

¹⁰In the example considered, the area of each finite element is the same. If this does not hold, the constraint has to be adjusted appropriately without loss of generality

- | | |
|---------------------------|---|
| (i) Problem configuration | (ii) Deterministic solution for $VF = 0.4$ obtained by setting $Var[\boldsymbol{\theta}] = 0$ |
|---------------------------|---|

Figure 11: Problem domain, boundary conditions and deterministic solution

z , are contained in the vicinity of the manifold implied by Equation (53). The additional term in p_{aux} in Equation (54) partially alters the associated update equations of the VB-EM scheme previously presented. We discuss these in detail in Appendix C as well. In the examples presented the value $\epsilon_c^2 = 10^{-10}$ was used.

For the complete definition of the problem, we note that the target response vector \mathbf{u}_{target} consisted of the *vertical displacement* u_2 at 8 points along the bottom boundary i.e. with $x_2 = 0$ and $x_1 = 0.2 k$, $k = 1 \div 8$ such that:

$$u_{target,k} = 6.25 \times 10^{-3} k \quad (55)$$

and $\tau_Q^{-1} = 5 \times 10^{-6}$ (Equation (8)). For comparison purposes, the deterministic problem was solved for $VF = 0.4$. To that end, the exact same algorithmic scheme for finding $\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z$ was employed (Appendix C) by assuming that the variance of the random variables $\boldsymbol{\theta}$ was zero and their mean exactly the same as detailed above. The resulting $\boldsymbol{\mu}_z$ which is shown in Figure 11ii was obtained after (approximately) 50 iterations and exhibits two diagonal ribs that are obviously critical in stiffening the system. As it is easily understood, the objective function is not (in general) concave and multiple local maxima could exist.

Figure 12 depicts the estimated $\boldsymbol{\mu}_z$ for the *stochastic problem* and for two volume fractions considered i.e. $VF = 0.4$ and $VF = 0.2$. The first was obtained with 35 forward calls whereas the second with 54. As compared to

(i) $VF = 0.4$

(ii) $VF = 0.2$

Figure 12: Computed μ_z (see Appendix B and Appendix C) as a function of the iteration number. Each iteration involves a forward call for the computation of the output \mathbf{u} and its derivatives. For $VF = 0.4$ and $VF = 0.2$ the computation required 35 and 54 such calls respectively.

(i) $VF = 0.4$

(ii) $VF = 0.2$

Figure 13: Evolution of \mathcal{F} (Equation (33)). Each iteration corresponds to one q (Equations (34), (35)) and one \mathbf{W} (Equation (38)) update

the deterministic solution in Figure 11ii with the two diagonal stiffening ribs, one notes that in Figure 12i only one is present. This could be attributed to a different local maximum or it could be the result of the random variability in the properties of the material.

More importantly the algorithm proposed can identify the most sensitive directions around the local maximum. These are obtained through successive iterations between q (Equations (34), (35)) and \mathbf{W} updates (Equation (38)). The evolution of the the variational lower-bound \mathcal{F} (Equation (33)) with regards to these iterations is depicted in Figure 13. We note that these iterations *do not entail any additional forward calls*. Some of the generalized eigenvectors identified $\hat{\mathbf{W}}$ (Equation (32)) and the associated variances are shown in Figure 15. Due to the presence of the constraint, the first (most sensitive) such eigenvector is determined by the gradient of the constraint at μ_z and the associated variance σ_1^2 (in parentheses, Figure 15) by the user-specified parameters ϵ_c (Equation (54)). Figure 14 depicts the evolution of the identified σ_j per VB-EM iteration where as it is clearly seen the first, most sensitive generalized eigenvectors are identified in the first few iterations. One also notes that the variances top-off at the prior

(i) $VF = 0.4$

(ii) $VF = 0.2$

Figure 14: Evolution of σ_j^2

(i) $(\sigma_1^2 = 7.31 \times 10^{-1})$

(v) $(\sigma_1^2 = 3.24 \times 10^0)$

(ii) $\sigma_2^2 = 1.25 \times 10^2$

(vi) $\sigma_2^2 = 1.93 \times 10^2$

(iii) $\sigma_5^2 = 2.78 \times 10^3$

(vii) $\sigma_5^2 = 1.91 \times 10^3$

(iv) $\sigma_7^2 = 1.36 \times 10^4$

(viii) $\sigma_9^2 = 1.99 \times 10^4$

Figure 15: Generalized eigenvectors $\hat{\mathbf{w}}_j$ for $VF = 0.4$ (left column) and $VF = 0.2$ (right column)

value $\tau_{y0}^{-1} = 10^4$.

Figure 16 depicts the squared values $(\hat{\mathbf{w}}_j)^2$ (shown in Figure 15) in a log-scale. This allows one to see how the sensitivity associated with each generalized eigenvector is spatially distributed. Finally Figure 18 depicts the outlines of sample designs drawn from $q(\boldsymbol{\mu}_z + \mathbf{W}y)$ (which approximates the expected utility $V(\boldsymbol{\mu}_z + \mathbf{W}y)$) corresponding to different (relative) levels of the the expected utility. In the approximation advocated, $\boldsymbol{\mu}_z$ represents the optimal design for which $V(\mathbf{z})$ attains its (locally) maximum value. By considering $V(\mathbf{z})$ less than the optimal, we can identify an infinity of alternative designs but also assess the sensitivity of the solution.

Finally in Table 3 we record the normalized KL-divergence as discussed in Section 2.6 and note that this decays for increasing d_y to relatively small values indicating a good quality in the approximation found, particularly for $VF = 0.4$.

$$\begin{array}{ll}
\text{(i)} (\hat{\mathbf{w}}_1)^2 & \text{(v)} (\hat{\mathbf{w}}_1)^2 \\
\text{(ii)} (\hat{\mathbf{w}}_2)^2 & \text{(vi)} (\hat{\mathbf{w}}_2)^2 \\
\text{(iii)} (\hat{\mathbf{w}}_5)^2 & \text{(vii)} (\hat{\mathbf{w}}_5)^2 \\
\text{(iv)} (\hat{\mathbf{w}}_7)^2 & \text{(viii)} (\hat{\mathbf{w}}_9)^2
\end{array}$$

Figure 16: The squares of the entries of each of the generalized eigenvectors $\hat{\mathbf{w}}_j$ (log scale) for $VF = 0.4$ (left column) and $VF = 0.2$ (right column)

$$\begin{array}{lll}
\text{(i)} \frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)} = 0.75 & \text{(ii)} \frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)} = 0.50 & \text{(iii)} \frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)} = 0.25
\end{array}$$

Figure 17: Outline of alternative designs \mathbf{z} at various levels of expected utility $\frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)}$ as compared to the optimal $\boldsymbol{\mu}_z$ ($VF = 0.4$)

d_y	nKL (Equation (45))	
	$VF = 0.4$	$VF = 0.2$
5	1.5×10^{-2}	3.4×10^{-1}
10	8.7×10^{-3}	1.9×10^{-1}
15	3.9×10^{-3}	1.3×10^{-1}
20	6.0×10^{-4}	6.8×10^{-2}

Table 3: Normalized KL-divergence from Equation (45) for example 2

4. Conclusions

We present a framework for solving a large class of model-based, optimization-under-uncertainty problems. The overarching idea is that of recasting the problem as one of probabilistic inference. This enables the uniform treatment of both random and design variables and is capable of furnishing not only a (local)

$$\begin{array}{lll}
\text{(i)} \frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)} = 0.75 & \text{(ii)} \frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)} = 0.50 & \text{(iii)} \frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)} = 0.25
\end{array}$$

Figure 18: Outline of alternative designs \mathbf{z} at various levels of expected utility $\frac{V(\mathbf{z})}{V(\boldsymbol{\mu}_z)}$ as compared to the optimal $\boldsymbol{\mu}_z$ ($VF = 0.2$)

maximum (i.e. a point estimate) but also the sensitivity of the objective to the design variables. To achieve this objective, we propose a Variational Bayesian framework that operates on two fronts. Firstly, it attempts to compute efficiently an accurate approximation of the joint density of interest. Secondly, it seeks a lower-dimensional subspace with regards to the design variables \mathbf{z} that provides an assessment of the solutions robustness by discovering the most sensitive directions i.e. the directions along which, variations in \mathbf{z} will cause the largest decrease in the expected utility. This is based on the same premise as the so-called Sloppy Models whose behavior depends only on a few stiff combinations of parameters, with many sloppy parameter directions largely unimportant for model behavior. The identification of this lower-dimensional subspace, enables the analyst to compute, apart from the optimal design, an infinity of alternative designs which achieve a lower value of the expected utility. Interestingly enough, addressing the probabilistic inference task under the Variational Bayesian perspective involves the solution of an optimization problem. To that end we propose an iterative VB-Expectation-Maximization scheme.

The aforementioned claims have been validated in the context of two numerical examples involving $\mathcal{O}(10^3)$ random and design variables. In all cases considered the cost of the computations in terms of calls to the forward model was of the order $\mathcal{O}(10 \div 10^2)$. The accuracy of the approximations provided is assessed by appropriate information-theoretic metrics.

The framework proposed cannot currently account for the possibility of multiple local maxima, as the approximation constructed is based on unimodal Gaussian densities. Nevertheless, the formulation can be readily extended by employing *mixture of Gaussians* that will enable not only approximations for multi-modal

cases but also produce better results for unimodal, but highly non-Gaussian densities. We note finally the possibility of using approximate, surrogate or reduced-order models in order to expedite computations. All the algorithmic steps discussed can be readily performed by using these less-expensive forward solvers. As long as these convey some information about the expensive, reference forward model, then they can provide a good starting point for further computations that would require fewer expensive calls to converge.

Appendix A. Maximization of \mathcal{F}_W

As discussed earlier, in order to update \mathbf{W} it suffices to consider only $\mathcal{F}_W(\mathbf{W})$ (Equation (38)):

$$\begin{aligned} \mathcal{F}_W(\mathbf{W}) &= -\frac{\tau_Q}{2} \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} : (\mathbf{C}_{yy} - \tau_z^{-1} \mathbf{I}) \\ &\quad - \frac{\tau_Q}{2} 2 \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{\theta y} + \log p_W(\mathbf{W}) \end{aligned} \quad (\text{A.1})$$

While the first part is quadratic with respect to \mathbf{W} the difficulty arises from the orthogonality constraint $\mathbf{W}^T \mathbf{W} = \mathbf{I}$ which can be enforced directly or through the regularization term $p_W(\mathbf{W})$ as previously discussed. To address this constrained optimization problem, we employ the iterative algorithm proposed in [59] which is highly efficient not only in terms of the number of iterations needed but also in terms of the the cost per iteration. It is based on the constraint-preserving Cayley transform according to which the current \mathbf{W} is updated to \mathbf{W}' as follows:

$$\mathbf{W}' = (\mathbf{I} + \frac{a}{2} \mathbf{A})^{-1} (\mathbf{I} - \frac{a}{2} \mathbf{A}) \mathbf{W} \quad (\text{A.2})$$

where:

$$\mathbf{A} = \mathbf{J} \mathbf{W}^T - \mathbf{W} \mathbf{J}^T \quad (\text{A.3})$$

and $\mathbf{J} = \frac{\partial \mathcal{F}_W}{\partial \mathbf{W}}$. The latter can be readily obtained from Equation (A.1):

$$\mathbf{J} = -\tau_Q \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} (\mathbf{C}_{yy} - \tau_z^{-1} \mathbf{I}) - \tau_Q \mathbf{G}_z^T \mathbf{G}_\theta \mathbf{C}_{\theta y} \quad (\text{A.4})$$

It can be shown that \mathbf{W}' satisfies automatically the orthogonality constraint and that and for $a = 0$, \mathbf{W}' is an ascent direction of \mathcal{F}_W . Several options exist for selecting the step size a . In the numerical illustrations we made use of the

Barzilai-Borwein scheme detailed in [60] which results in a non-monotone line search algorithm. We note that the inversion of the $d_z \times d_z$ matrix $(\mathbf{I} + \frac{a}{2}\mathbf{A})$ can be efficiently performed by inverting a matrix of dimension $2d_y \times 2d_y$ which is much smaller than d_z [59]. We finally re-emphasize that the updates of \mathbf{W} require no forward calls. The updates/iterations are terminated when no further improvement to the objective \mathcal{F}_W is possible.

Appendix B. Maximization of \mathcal{F}_μ

As it was previously discussed, in order to update $\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z$ it suffices to consider only $\mathcal{F}_\mu(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta)$ (Equation (37)):

$$\begin{aligned} \mathcal{F}_\mu(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta) = & -\frac{\tau_Q}{2} (|\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)|^2) \\ & -\frac{1}{2}(\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0})^T \mathbf{C}_{\theta 0}^{-1}(\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0}) + \log p_{\mu_z}(\boldsymbol{\mu}_z) \end{aligned} \quad (\text{B.1})$$

This represent a nonlinear, unconstrained optimization problem that can be solved with any of the well-known algorithms [61, 62]. We present here a Gauss-Newton type algorithm that we employed and produced the results discussed in Section 3. For clarity of the presentation we consider first the case in the first numerical illustration where the regularization/prior $p_{\mu_z}(\boldsymbol{\mu}_z)$ was a Gaussian $\mathcal{N}(\mathbf{0}, \mathbf{C}_{z0})$ in which case:

$$\begin{aligned} \mathcal{F}_\mu(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta) = & -\frac{\tau_Q}{2} (|\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_z)|^2) \\ & -\frac{1}{2}(\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0})^T \mathbf{C}_{\theta 0}^{-1}(\boldsymbol{\mu}_\theta - \boldsymbol{\mu}_{\theta 0}) - \frac{1}{2}\boldsymbol{\mu}_z^T \mathbf{C}_{z0}^{-1}\boldsymbol{\mu}_z \end{aligned} \quad (\text{B.2})$$

If $(\boldsymbol{\mu}_z^{(t)}, \boldsymbol{\mu}_\theta^{(t)})$ denote the values at iteration t and $(\boldsymbol{\mu}_z^{(t+1)} = \boldsymbol{\mu}_z^{(t)} + \Delta\boldsymbol{\mu}_z^{(t)}, \boldsymbol{\mu}_\theta^{(t+1)} = \boldsymbol{\mu}_\theta^{(t)} + \Delta\boldsymbol{\mu}_\theta^{(t)})$, then a first-order Taylor series yields the following approximation:

$$\begin{aligned} \mathcal{F}_\mu(\Delta\boldsymbol{\mu}_z^{(t)}, \Delta\boldsymbol{\mu}_\theta^{(t)}) \approx & -\frac{\tau_Q}{2} \left(|\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_z^{(t)}, \boldsymbol{\mu}_\theta^{(t)}) - \mathbf{G}_{\theta,t}\Delta\boldsymbol{\mu}_\theta^{(t)} - \mathbf{G}_{z,t}\Delta\boldsymbol{\mu}_z^{(t)}|^2 \right) \\ & -\frac{1}{2}(\boldsymbol{\mu}_\theta^{(t)} + \Delta\boldsymbol{\mu}_\theta^{(t)} - \boldsymbol{\mu}_{\theta 0})^T \mathbf{C}_{\theta 0}^{-1}(\boldsymbol{\mu}_\theta^{(t)} + \Delta\boldsymbol{\mu}_\theta^{(t)} - \boldsymbol{\mu}_{\theta 0}) \\ & -\frac{1}{2}(\boldsymbol{\mu}_z^{(t)} + \Delta\boldsymbol{\mu}_z^{(t)})^T \mathbf{C}_{z 0}^{-1}(\boldsymbol{\mu}_z^{(t)} + \Delta\boldsymbol{\mu}_z^{(t)}) \end{aligned} \quad (\text{B.3})$$

where $\mathbf{G}_{\theta,t} = \frac{\partial \mathbf{u}}{\partial \boldsymbol{\theta}}|_{\boldsymbol{\theta}=\boldsymbol{\mu}_\theta^{(t)}}$ and $\mathbf{G}_{z,t} = \frac{\partial \mathbf{u}}{\partial \mathbf{z}}|_{\mathbf{z}=\boldsymbol{\mu}_z^{(t)}}$. Differentiating with respect to $(\Delta\boldsymbol{\mu}_z^{(t)}, \Delta\boldsymbol{\mu}_\theta^{(t)})$ leads to the following system of coupled linear equations:

$$\begin{bmatrix} \frac{\partial \mathcal{F}_\mu}{\partial \Delta\boldsymbol{\mu}_\theta^{(t)}} \\ \frac{\partial \mathcal{F}_\mu}{\partial \Delta\boldsymbol{\mu}_z^{(t)}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix} \rightarrow \mathbf{H}_t \begin{bmatrix} \Delta\boldsymbol{\mu}_\theta^{(t)} \\ \Delta\boldsymbol{\mu}_z^{(t)} \end{bmatrix} = \mathbf{h}_t \quad (\text{B.4})$$

where:

$$\mathbf{H}_t = \begin{bmatrix} \tau_Q \mathbf{G}_{\theta,t}^T \mathbf{G}_{\theta,t} + \mathbf{C}_{\theta 0}^{-1} & \tau_Q \mathbf{G}_{\theta,t}^T \mathbf{G}_{z,t} \\ \tau_Q \mathbf{G}_{z,t}^T \mathbf{G}_{\theta,t} & \tau_Q \mathbf{G}_{z,t}^T \mathbf{G}_{z,t} + \mathbf{C}_{z 0}^{-1} \end{bmatrix} \quad (\text{B.5})$$

and:

$$\mathbf{h}_t = \begin{bmatrix} \tau_Q \mathbf{G}_{\theta,t}^T (\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_z^{(t)}, \boldsymbol{\mu}_\theta^{(t)})) - \mathbf{C}_{\theta 0}^{-1}(\boldsymbol{\mu}_\theta^{(t)} - \boldsymbol{\mu}_{\theta 0}) \\ \tau_Q \mathbf{G}_{z,t}^T (\mathbf{u}_{target} - \mathbf{u}(\boldsymbol{\mu}_z^{(t)}, \boldsymbol{\mu}_\theta^{(t)})) - \mathbf{C}_{z 0}^{-1}\boldsymbol{\mu}_z^{(t)} \end{bmatrix} \quad (\text{B.6})$$

We note that at each iteration the forward solver needs to be called for the computation of the output vector \mathbf{u} and its derivatives $\mathbf{G}_{\theta,z}$. Iterations are terminated when no further improvement is possible i.e. $\frac{|\Delta\boldsymbol{\mu}_\theta^{(t)}|}{|\boldsymbol{\mu}_\theta^{(t)}|}, \frac{|\Delta\boldsymbol{\mu}_z^{(t)}|}{|\boldsymbol{\mu}_z^{(t)}|} < (tolerance) = 10^{-5}$.

Appendix C. Regularization of μ_z and update equation for Num. Illustration 2

We discuss in this section the definition of the regularization/prior $p_{\mu_z}(\mu_z)$ for numerical illustration 2 (Section 3.2) and the resulting changes in the optimization scheme for μ_z in Appendix B. Given the physical interpretation of the design variables μ_z as binary variables which for each pixel indicate the presence or not of material, we adopt a regularization for μ_z that promotes the discovery of such solutions but also exhibits the requisite spatial correlation. To that end we propose a hierarchical prior where in addition to $\mu_z = \{\mu_{z,j}\}_{j=1}^{3536}$ we introduce the binary hyperparameters $\phi = \{\phi_j = \pm 1\}_{j=1}^{3536}$ such that:

$$p_{\mu_z}(\mu_z|\phi) = \prod_{j=1}^{3536} p(\mu_{z,j}|\phi_j) \quad (\text{C.1})$$

where $p(\mu_{z,j}|\phi_j = -1) = \mathcal{N}(-m, s^2)$ and $p(\mu_{z,j}|\phi_j = +1) = \mathcal{N}(m, s^2)$. The value of m was selected so that in combination with the sigmoid function (Equation (52)) produces solutions close to the binary images we would like to achieve:

$$\frac{1}{1 + e^m} = 10^{-3} \approx 0, \quad \frac{1}{1 + e^{-m}} = 1 - 10^{-3} \approx 1 \quad (\text{C.2})$$

This yields $m = -6.9$ and the resulting, bimodal, hierarchical prior is depicted in Figure C.19i. In order to account for the spatial dependence of neighboring $\mu_{z,j}$ we employ an auto-logistic hyperprior on ϕ of the following form [63, 64]:

$$p(\phi|\beta) \propto e^{-\frac{\beta}{2} \sum_j \sum_{k \sim j} \phi_j \phi_k} \quad (\text{C.3})$$

The second sum in the expression above is over all indices k which correspond to sites neighboring to j (neighborhood relation denoted by \sim). Given the triangular mesh used, we consider 3 neighbors for each site j as shown in Figure C.19ii. The hyperparameter β controls the strength of spatial correlation. At one extreme, if $\beta \rightarrow +\infty$, neighboring ϕ_j prefer to have different values (i.e. $-1/+1$ or $+1/-1$) as this yields a higher hyperprior value. At the other extreme, if $\beta \rightarrow -\infty$, neighboring ϕ_j prefer to have the same values (i.e. $-1/-1$ or $+1/+1$). For $\beta = 0$ no correlation is present. We note that the aforementioned prior in ϕ imbues indirectly spatial correlation in $\mu_{z,j}$.

In summary, the prior $p_{\mu_z}(\mu_z)$ can be found by integrating out the hyperparameters ϕ and β as:

$$p_{\mu_z}(\mu_z) = \int p_{\mu_z}(\mu_z|\phi)p(\phi|\beta) d\phi d\beta \quad (\text{C.4})$$

The integration above cannot be performed analytically and for that reason we employed an Expectation-Maximization scheme [65, 27] whereby at the Expectation step a Metropolized-Gibbs scheme is used to sample the hyperparameters ϕ and β from their conditional posterior (given the current value of μ_z). This does not require any forward calls and can be very efficiently performed. The samples generated can be used to estimate $\log p_{\mu_z}(\mu_z)$ and its derivatives as needed for the update equations in Appendix B. If we denote with $\langle \rangle$ expectations with regards to the posterior samples of ϕ described above and by keeping only terms that depend on μ_z we obtain that:

$$\begin{aligned} \log p_{\mu_z}(\mu_z) &= -\frac{1}{2s^2} \sum_j \langle (\mu_{z,j} - m\phi_j)^2 \rangle \\ &= -\frac{1}{2s^2} (\mu_z^T \mu_z - 2m\mu_z^T \langle \phi \rangle + \langle \phi^T \phi \rangle) \end{aligned} \quad (\text{C.5})$$

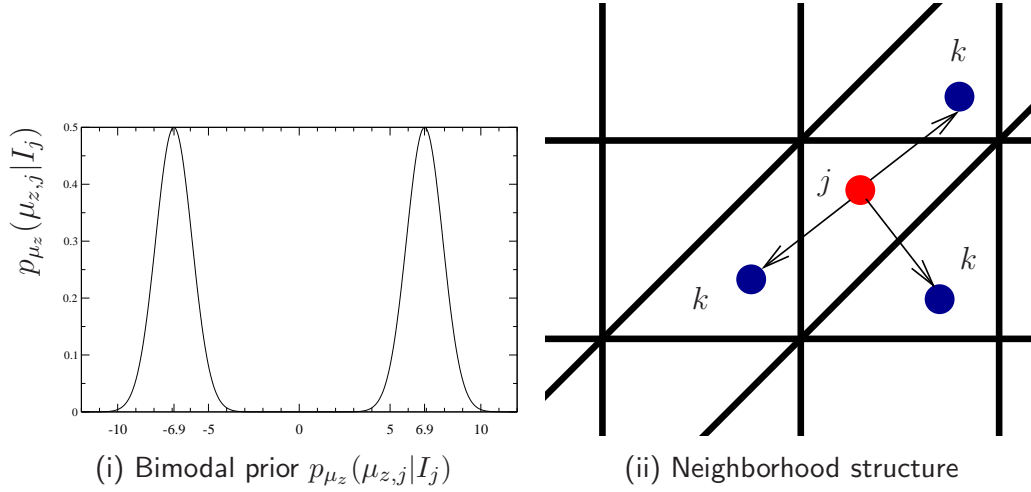


Figure C.19: Definition of $p_{\mu_z}(\mu_z)$

The quadratic form of this expression implies that the only changes in the update Equation (B.4) in Appendix B will be in:

$$\mathbf{H}_t = \begin{bmatrix} \tau_Q \mathbf{G}_{\theta,t}^T \mathbf{G}_{\theta,t} + \mathbf{C}_{\theta 0}^{-1} & \tau_Q \mathbf{G}_{\theta,t}^T \mathbf{G}_{z,t} \\ \tau_Q \mathbf{G}_{z,t}^T \mathbf{G}_{\theta,t} & \tau_Q \mathbf{G}_{z,t}^T \mathbf{G}_{z,t} + \frac{1}{s^2} \mathbf{I} \end{bmatrix} \quad (\text{C.6})$$

and:

$$\mathbf{h}_t = \begin{bmatrix} \tau_Q \mathbf{G}_{\theta,t}^T (\mathbf{u}_{target} - \mathbf{u}(\mu_z^{(t)}, \mu_z^{(t)})) - \mathbf{C}_{\theta 0}^{-1} (\mu_{\theta}^{(t)} - \mu_{\theta 0}) \\ \tau_Q \mathbf{G}_{z,t}^T (\mathbf{u}_{target} - \mathbf{u}(\mu_z^{(t)}, \mu_z^{(t)})) - \frac{1}{s^2} (\mu_z^{(t)} - m < \phi >) \end{bmatrix} \quad (\text{C.7})$$

The aforementioned equations should be augmented by the equality constraint in Equation (53). We enforce this constraint directly on μ_z so as the optimal design (μ_z) satisfies it. From an algorithmic point, the process adopted is similar to Sequential Quadratic Programming (SQP, [61]) where the quadratized objective (Equation (B.3)) at each iteration t is augmented by the linearized

constraint:

$$0 = c(\boldsymbol{\mu}_z^{(t)} + \Delta \boldsymbol{\mu}_z^{(t)}) \approx c(\boldsymbol{\mu}_z^{(t)}) + \mathbf{f}_t^T \Delta \boldsymbol{\mu}_z^{(t)} \quad (\text{C.8})$$

where $\mathbf{f}_t = \frac{\partial c}{\partial \mathbf{z}}|_{\mathbf{z}=\boldsymbol{\mu}_z^{(t)}}$.

In order to account for the constraint in the rest of the auxiliary density p_{aux} , the scheme described in Equation (54) is adopted which induces a soft/probabilistic enforcement. The term $-\frac{c^2(\mathbf{z})}{2\epsilon_c^2}$ will therefore yield an additional contribution in the variational lower-bound \mathcal{F} detailed in Equation (17). If we denote by \mathcal{F}_c this additional term, then:

$$\begin{aligned} \mathcal{F}_c &= -\frac{1}{2\epsilon_c^2} E_q[c^2(\mathbf{z})] \\ &= -\frac{1}{2\epsilon_c^2} E_q[c^2(\boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z)] \end{aligned} \quad (\text{C.9})$$

Given the nonlinear form of $c(\mathbf{z})$, we employ another linearization around $\boldsymbol{\mu}_z$:

$$\begin{aligned} c(\boldsymbol{\mu}_z + \mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) &\approx c(\boldsymbol{\mu}_z) + \mathbf{f}^T(\mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) \\ &= \mathbf{f}^T(\mathbf{W}\mathbf{y} + \boldsymbol{\eta}_z) \quad (\text{since } c(\boldsymbol{\mu}_z) = 0) \end{aligned} \quad (\text{C.10})$$

where $\mathbf{f} = \frac{\partial c}{\partial \mathbf{z}}|_{\mathbf{z}=\boldsymbol{\mu}_z}$. As a result of this and the form of q (Equation (27)), Equation (C.9) becomes:

$$\mathcal{F}_c = -\frac{1}{2\epsilon_c^2} (\mathbf{W}^T \mathbf{f} \mathbf{f}^T \mathbf{W} : \mathbf{C}_{yy} + \tau_z^{-1} \mathbf{f} \mathbf{f}^T : (\mathbf{I} - \mathbf{W}\mathbf{W}^T)) \quad (\text{C.11})$$

which when combined with the rest of the terms in \mathcal{F} in Equation (33), leads to the following changes in the update equations in the VB-EM scheme:

- **VB-Expectation:**

$$\begin{bmatrix} \mathbf{C}_{\theta\theta}^{opt} & \mathbf{C}_{\theta y}^{opt} \\ sym. & \mathbf{C}_{yy}^{opt} \end{bmatrix}^{-1} = \begin{bmatrix} \tau_Q \mathbf{G}_\theta^T \mathbf{G}_\theta + \mathbf{C}_{\theta 0}^{-1} & \tau_Q \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} \\ sym. & \tau_Q \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} + \tau_{y0} \mathbf{I} + \frac{1}{\epsilon_c^2} \mathbf{W}^T \mathbf{f} \mathbf{f}^T \mathbf{W} \end{bmatrix} \quad (\text{C.12})$$

and:

$$\tau_z^{opt} = \tau_{z0} + \frac{1}{d_z - d_y} (\tau_Q \mathbf{G}_z^T \mathbf{G}_z + \frac{1}{\epsilon_c^2} \mathbf{f} \mathbf{f}^T) : (\mathbf{I} - \mathbf{W} \mathbf{W}^T) \quad (\text{C.13})$$

- **VB-Maximization:**

$$\mathbf{W}^{opt} = \arg \max_{\mathbf{W}} \hat{\mathcal{F}}_W(\mathbf{W}) \quad (\text{C.14})$$

where:

$$\begin{aligned} \mathcal{F}_W(\mathbf{W}) = & -(\frac{\tau_Q}{2} \mathbf{W}^T \mathbf{G}_z^T \mathbf{G}_z \mathbf{W} + \frac{1}{2\epsilon_c^2} \mathbf{W}^T \mathbf{f} \mathbf{f}^T \mathbf{W}) : (\mathbf{C}_{yy} - \tau_z^{-1} \mathbf{I}) \\ & - \frac{\tau_Q}{2} 2 \mathbf{G}_\theta^T \mathbf{G}_z \mathbf{W} : \mathbf{C}_{\theta y} + \log p_W(\mathbf{W}) \end{aligned} \quad (\text{C.15})$$

References

- [1] L. T. Biegler, O. Ghattas, M. Heinkenschloss, B. v. B. Waanders, Large-Scale PDE-Constrained Optimization: An Introduction, in: L. T. Biegler, M. Heinkenschloss, O. Ghattas, B. v. B. Waanders (Eds.), Large-Scale PDE-Constrained Optimization, no. 30 in Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 2003, pp. 3–13.
- [2] V. Akcelik, G. Biros, O. Ghattas, D. Keyes, K. Ko, L.-Q. Lee, E. G. Ng, Adjoint methods for electromagnetic shape optimization of the low-loss cavity for the International Journal of Physics: Conference Series 16 (1) (2005) 435.
doi:10.1088/1742-6596/16/1/059.
URL <http://iopscience.iop.org/1742-6596/16/1/059>
- [3] M. Hinze, R. Pinnau, M. Ulbrich, S. Ulbrich, Optimization with PDE constraints, Springer, New York, 2009.
- [4] M. A. Dihlmann, B. Haasdonk, Certified PDE-constrained parameter optimization using reduced basis Computational Optimization and Applications 60 (3) (2014) 753–787.
doi:10.1007/s10589-014-9697-1.
URL <http://link.springer.com/article/10.1007/s10589-014-9697-1>
- [5] L. F. Yongjin Zhang, Accelerating PDE constrained optimization by the reduced basis method: application to batch chromatography, International Journal for Numerical Methods in Engineeringdoi:10.1002/nme.4950.
- [6] G. B. Dantzig, Linear Programming under Uncertainty, Management Science 1 (3-4) (1955) 197–206. doi:10.1287/mnsc.1.3-4.197.
URL <http://mansci.journal.informs.org/content/1/3-4/197>

- [7] Y. Tsompanakis, N. D. Lagaros, P. , Structural design optimization considering uncertainties, Taylor & Francis, London; New York, 2008.
- [8] N. D. Lagaros, V. Plevris, M. Papadrakakis, Multi-objective design optimization using cascade evolutionary computations, *Computer Methods in Applied Mechanics and Engineering* (194) (2005) 3496–3515.
- [9] H.-G. Beyer, B. Sendhoff, Robust optimization - A comprehensive survey, *Computer Methods in Applied Mechanics and Engineering* 196 (33-34) (2007) 3190–3218, wOS:000248605400009. doi:10.1016/j.cma.2007.03.003.
- [10] M. Papadrakakis, N. D. Lagaros, V. Plevris, Design optimization of steel structures considering uncertainties, *Engineering Structures* (27) (2005) 1408–1418.
- [11] M. Papadrakakis, N. D. Lagaros, Reliability-based structural optimization using neural networks and Monte Carlo simulation, *Computer Methods in Applied Mechanics and Engineering* (191) (2002) 3491–3507.
- [12] S. Sankaran, C. Audet, A. L. Marsden, A method for stochastic constrained optimization using derivative-free surrogate pattern search and *Journal of Computational Physics* 229 (12) (2010) 4664–4682. doi:10.1016/j.jcp.2010.03.005.
URL <http://www.sciencedirect.com/science/article/pii/S0021999110001154>
- [13] S. Sankaran, Stochastic optimization using a sparse grid collocation scheme, *Probabilistic Engineering Mechanics* (24) (2009) 382–396.

- [14] S. Hyun, S. Torquato, Designing composite microstructures with targeted properties, *Journal of Materials Research* 16 (01) (2001) 280–285.
doi:10.1557/JMR.2001.0042.
URL <http://adsabs.harvard.edu/abs/2001JMatR..16..280H>
- [15] E. Capiez-Lernout, C. Soize, Robust Design Optimization in Computational Mechanics, *Journal of Applied Mechanics* 75 (2) (2008) 021001–021001.
doi:10.1115/1.2775493.
URL <http://dx.doi.org/10.1115/1.2775493>
- [16] N. Zabaras, B. Ganapathysubramanian, A scalable framework for the solution of stochastic inverse problems using a sparse grid collocation approach, *Journal of Computational Physics* (227) (2008) 4697–4735.
- [17] P. Seshadri, P. Constantine, G. Iaccarino, Aggressive Design Under Uncertainty, in: 16th AIAA Non-Deterministic Approaches Conference, American Institute of Aeronautics and Astronautics.
URL <http://arc.aiaa.org/doi/abs/10.2514/6.2014-1007>
- [18] H. Robbins, S. Monro, A stochastic approximation method, *The Annals of Mathematical Statistics* 22 (3) (1951) 400–407.
URL <http://www.jstor.org/stable/2236626>
- [19] J. Kiefer, J. Wolfowitz, Stochastic Estimation of the Maximum of a Regression Function, *The Annals of Mathematical Statistics* 23 (3) (1952) 462–466.
doi:10.1214/aoms/1177729392.
URL <http://projecteuclid.org/euclid.aoms/1177729392>

- [20] J. C. Spall, Introduction to Stochastic Search and Optimization, Estimation, Simulation, and Control. Wiley.
- [21] H. Kushner, G. Yin, Stochastic approximation and recursive algorithms and applications, Springer, 2003.
- [22] P. Muller, Simulation based optimal design, in: Bayesian Statistics 6, 1998.
- [23] R. Sternfels, P.-S. Koutsourelakis, Stochastic Design and Control in Random Heterogeneous Materials, International Journal for Multiscale Computational Engineering 9 (4) (2011) 425–443, wOS:000296218100006.
- [24] B. Amzal, F. Y. Bois, E. Parent, C. P. Robert, Bayesian-Optimal Design via Interacting Particle Systems, Journal of the American Statistical Association 101 (474) (2006) 773–785. doi:10.1198/016214505000001159.
URL <http://www.tandfonline.com/doi/abs/10.1198/016214505000001159>
- [25] H. Kuck, N. de Freitas, A. Doucet, SMC Samplers for Bayesian Optimal Nonlinear Design, in: 2006 IEEE Nonlinear Statistical Signal Processing Workshop, 2006, pp. 99–102. doi:10.1109/NSSPW.2006.4378829.
- [26] M. J. Beal, Z. Ghahramani, The Variational Bayesian EM Algorithm for Incomplete Data: with Application to Scoring Graphical Model Structures, Bayesian Statistics (7).
- [27] C. Bishop, Pattern Recognition and Machine Learning, 1st Edition, Springer, New York, 2007.

- [28] M. Jordan, Z. Ghahramani, T. Jaakola, L. Saul, An Introduction to Variational Methods for Graphical Models, *Machine Learning* 37 (1999) 183–233.
- [29] H. Attias, A Variational Bayesian Framework for Graphical Models, in: *In Advances in Neural Information Processing Systems 12*, MIT Press, 2000, pp. 209–215.
- [30] M. Wainwright, M. Jordan, Graphical models, exponential families, and variational inference, in: *Foundations and Trends in Machine Learning*, Vol. 1 of 1-305, 2008, pp. 1–305.
- [31] T. Cover, J. Thomas, *Elements of Information Theory*, John Wiley & Sons, 1991.
- [32] B. A. Olshausen, D. J. Field, Sparse coding with an overcomplete basis set: A strategy employed by Vision Research 37 (23) (1997) 3311–3325.
doi:10.1016/S0042-6989(97)00169-7.
URL <http://www.sciencedirect.com/science/article/pii/S0042698997001697>
- [33] M. Lewicki, T. J. Sejnowski, Learning overcomplete representations, *Neural Computation* 12 (2) (2000) 337–365.
- [34] K. S. Brown, J. P. Sethna, Statistical mechanical approaches to models with many poorly known parameters, *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics* 68 (2 Pt 1) (2003) 021904.
- [35] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, J. P. Sethna, Universally Sloppy Parameter Sensitivities in Systems Biology Models,

- PLoS Comput Biol 3 (10) (2007) e189.
doi:10.1371/journal.pcbi.0030189.
URL <http://dx.plos.org/10.1371/journal.pcbi.0030189>
- [36] J. F. Apgar, D. K. Witmer, F. M. White, B. Tidor, Sloppy models, parameter uncertainty, and the role of experimental design, *Molecular BioSystems* 6 (10) (2010) 1890. doi:10.1039/b918098b.
URL <http://xlink.rsc.org/?DOI=b918098b>
- [37] B. B. Machta, R. Chachra, M. K. Transtrum, J. P. Sethna, Parameter Space Compression Underlies Emergent Theories and Predictive Models, *Science* 342 (6158) (2013) 604–607. doi:10.1126/science.1238723.
URL <http://www.sciencemag.org/content/342/6158/604>
- [38] M. Tipping, C. M. Bishop, Probabilistic Principal Component Analysis, *Journal of the Royal Statistical Society B* 61 (1999) 611–622.
- [39] T. W. Lukaczyk, P. Constantine, F. Palacios, J. J. Alonso, Active Subspaces for Shape Optimization, in: 10th AIAA Multidisciplinary Design Optimization Conference, American Institute of Aeronautics and Astronautics.
URL <http://arc.aiaa.org/doi/abs/10.2514/6.2014-1171>
- [40] N. Dobigeon, J.-Y. Tournet, Bayesian Orthogonal Component Analysis for Sparse Representation, *IEEE Transactions on Signal Processing* 58 (5) (2010) 2675–2685. doi:10.1109/TSP.2010.2041594.
- [41] E. Candes, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal

- reconstruction from highly incomplete frequency information, *IEEE Trans. Information Theory* 52 (2006) 489–509.
- [42] D. P. Wipf, B. D. Rao, Sparse Bayesian learning for basis selection, *IEEE Transactions on Signal Processing* 52 (8) (2004) 2153–2164.
- [43] H. Lee, A. Battle, R. Raina, A. Y. Ng, Efficient sparse coding algorithms, in: B. Schoelkopf, J. C. Platt, T. Hoffman (Eds.), *Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems*, Vancouver, British Columbia, Canada, December 4-7, 2006, MIT Press, 2006, pp. 801–808.
- [44] M. Seeger, D. Wipf, Variational Bayesian Inference Techniques, *IEEE Signal Processing Magazine* 27 (6) (2010) 81–91. doi:10.1109/MSP.2010.938082.
- [45] M. J. Beal, Variational Algorithms for Approximate Bayesian Inference, Ph.D. thesis, Gatsby Computational Neuroscience Unit, University College London (2003).
URL <http://www.cse.buffalo.edu/faculty/mbeal/thesis/index.html>
- [46] R. J. Muirhead, *Aspects of Multivariate Statistical Theory*, Wiley, 1982.
- [47] R. Neal, G. Hinton, A view of the EM algorithm that justifies incremental, sparse, and other variants, in: M. Jordan (Ed.), *Learning in Graphical Models*, Dordrecht: Kluwer Academic Publishers, 1998, pp. 355–368.
- [48] R. A. Choudrey, S. J. Roberts, Variational mixture of Bayesian independent component analyzers, *Neural Computation* 15 (1) (2003) 213–252. doi:10.1162/089976603321043766.

- [49] J. K. Guest, T. Igusa, Structural optimization under uncertain loads and nodal locations, *Computer Methods in Applied Mechanics and Engineering* (198) (2008) 116–124.
- [50] M. Schevenels, B. S. Lazarov, O. Sigmund, Robust topology optimization accounting for spatially varying manufacturing errors, *Computer Methods in Applied Mechanics and Engineering* 200 (49-52) (2011) 3613–3627, wOS:000297494700009. doi:10.1016/j.cma.2011.08.006.
- [51] M. Jalalpour, T. Igusa, J. K. Guest, Optimal design of trusses with geometric imperfections: Accounting for global instability, *International Journal of Solids and Structures* 48 (21) (2011) 3011–3019, wOS:000295422900002. doi:10.1016/j.ijsolstr.2011.06.020.
- [52] B. S. Lazarov, M. Schevenels, O. Sigmund, Topology optimization with geometric uncertainties by perturbation techniques, *International Journal for Numerical Methods in Engineering* 90 (11) (2012) 1321–1336, wOS:000304446400001. doi:10.1002/nme.3361.
- [53] M. Tootkaboni, A. Asadpoure, J. K. Guest, Topology optimization of continuum structures under uncertainty A Polynomial Chaos approach, *Computer Methods in Applied Mechanics and Engineering* 201204 (2012) 263–275. doi:10.1016/j.cma.2011.09.009.
URL <http://www.sciencedirect.com/science/article/pii/S0045782511002982>
- [54] N. H. Kim, H. Wang, N. V. Queipo, Efficient Shape Optimization Under Uncertainty Using Polynomial Chaos Expansions and Local Sensitivities, *AIAA Journal* (44) (2006) 1112–1115.

- [55] B. S. Lazarov, M. Schevenels, O. Sigmund, Topology optimization considering material and geometric uncertainties using stochastic collocation, *Structural and Multidisciplinary Optimization* 46 (4) (2012) 597–612.
doi:10.1007/s00158-012-0791-7.
URL <http://link.springer.com/10.1007/s00158-012-0791-7>
- [56] S. Conti, H. Held, M. Pach, M. Rumpf, R. Schultz, Shape Optimization Under Uncertainty-a Stochastic Programming Perspective, *Siam Journal on Optimization* 19 (4) (2009) 1610–1632, wOS:000264353500005.
doi:10.1137/070702059.
- [57] S. Chen, W. Chen, A new level-set based approach to shape and topology optimization under geometric uncertainty, *Structural and Multidisciplinary Optimization* 44 (1) (2011) 1–18, wOS:000291254600001.
doi:10.1007/s00158-011-0660-9.
- [58] M. Bendsoe, O. Sigmund, *Topology Optimization - Theory, Methods, and Applications*, Springer, 2003.
- [59] Z. Wen, W. Yin, A feasible method for optimization with orthogonality constraints, *Mathematical Programming* 142 (1-2) (2012) 397–434.
doi:10.1007/s10107-012-0584-1.
URL <http://link.springer.com/article/10.1007/s10107-012-0584-1>
- [60] J. Barzilai, J. M. Borwein, Two-Point Step Size Gradient Methods, *IMA Journal of Numerical Analysis* 8 (1) (1988) 141–148.
doi:10.1093/imanum/8.1.141.
URL <http://imajna.oxfordjournals.org/content/8/1/141>

- [61] J. Nocedal, S. J. Wright, Numerical Optimization, Springer Science & Business Media, 1999.
- [62] S. P. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004.
- [63] J. Besag, Nearest-Neighbour Systems and Auto-Logistic Model for Binary Data, Journal of the Royal Statistical Society Series B-Statistical Methodology 34 (1) (1972) 75–84, wOS:A1972M842900004.
- [64] J. Besag, Spatial Interaction and Statistical-Analysis of Lattice Systems, Journal of the Royal Statistical Society Series B-Methodological 36 (2) (1974) 192–236, wOS:A1974U703600004.
- [65] A. Dempster, N. Laird, D. Rubin, Maximum likelihood from incomplete data via the EM algorithm (with discussion), J. Roy. Statist. Soc. Ser. B 39 (1) (1977) 1–38.