MIT Open Access Articles

An approximate empirical Bayesian method for large-scale linear-Gaussian inverse problems

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Zhou, Qingping, et al. "An approximate empirical Bayesian method for large-scale linear-Gaussian inverse problems." *Inverse Problems* 34, 9 (June 2018) © 2018 IOP Publishing Ltd.

As Published: <http://dx.doi.org/10.1103/10.1088/1361-6420/aac287>

Publisher: IOP Publishing

Persistent URL: <https://hdl.handle.net/1721.1/122927>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

An approximate empirical Bayesian method for large-scale linear-Gaussian inverse problems

Qingping Zhou

Department of Mathematics and Institute of Natural Sciences, Shanghai Jiao Tong University, Shanghai 200240, China

E-mail: zhou2015@sjtu.edu.cn

Wenqing Liu

Department of Mathematics and Zhiyuan College, Shanghai Jiao Tong University, Shanghai 200240, China

E-mail: wenqingliu@sjtu.edu.cn

Jinglai Li

Institute of Natural Sciences, Department of Mathematics and MOE Key Laboratory of Scientific and Engineering Computing, Shanghai Jiao Tong University, Shanghai 200240, China

E-mail: jinglaili@sjtu.edu.cn

Youssef M. Marzouk

Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

E-mail: ymarz@mit.edu

March 2018

Abstract. We study Bayesian inference methods for solving linear inverse problems, focusing on hierarchical formulations where the prior or the likelihood function depend on unspecified hyperparameters. In practice, these hyperparameters are often determined via an empirical Bayesian method that maximizes the marginal likelihood function, i.e., the probability density of the data conditional on the hyperparameters. Evaluating the marginal likelihood, however, is computationally challenging for large-scale problems. In this work, we present a method to approximately evaluate marginal likelihood functions, based on a low-rank approximation of the update from the prior

covariance to the posterior covariance. We show that this approximation is optimal in a minimax sense. Moreover, we provide an efficient algorithm to implement the proposed method, based on a combination of the randomized SVD and a spectral approximation method to compute square roots of the prior covariance matrix. Several numerical examples demonstrate good performance of the proposed method.

1. Introduction

Bayesian inference approaches have become increasingly popular as a tool to solve inverse problems [33, 17, 32]. In this setting, the parameters of interest are treated as random variables, endowed with a prior distribution that encodes information available before the data are observed. Observations are modeled by their joint probability distribution conditioned on the parameters of interest, which defines the likelihood function; this distribution incorporates the forward model and a stochastic description of measurement or model errors. The posterior distribution of the unknown parameters is then obtained via Bayes’ formula, which combines the information provided by the prior and the data. Here we focus on a particular class of Bayesian inverse problems, where the forward model is linear and both the prior and the error distributions are Gaussian. Such *linear-Gaussian* inverse problems arise in many real-world applications [33, 17].

In practice, a difficulty in implementing Bayesian inference for such problems is that the prior and/or the likelihood function may contain unspecified hyperparameters. For instance, the prior distribution usually encodes some kind of correlation among the inversion parameters, but the correlation length is often unknown. Alternatively, the variance and/or the mean of the observational errors may not be available, thus yielding unspecified hyperparameters in the likelihood function. A natural idea to address such problems is to construct a *hierarchical* Bayesian model for both the inversion parameters and the hyperparameters. Then both sets of parameters can be inferred from data through the characterization of their joint posterior distribution. This inferential approach is often called “fully Bayesian.” A fully Bayesian method can be prohibitively expensive for large-scale inverse problems, however, as it requires sampling from a posterior distribution that is higher dimensional and no longer Gaussian. An alternative strategy is the empirical Bayes (EB) approach [7, 6], which first estimates the hyperparameters by maximizing their marginal likelihood function, i.e., the probability density of the data conditioned only on the hyperparameters, and then plugs in the estimated values to compute the posterior of the inversion parameters.

The use of the EB method has been theoretically justified [23, 28]: roughly speaking, these theoretical results show that, under certain conditions and for a sufficiently large sample size, the EB method leads to similar inferential conclusions as fully Bayesian inference. At the same time, the EB method is more computationally efficient than the fully Bayesian approach, as it does not require characterizing the joint posterior distribution.

Nonetheless, the numerical implementation of EB methods remains computationally taxing when dimensionality of the inversion parameters is high, as it involves maximizing the marginal likelihood function, which in turn requires “integrating out” the inversion parameters for each evaluation. (Details about the computational cost of evaluating the marginal likelihood function are given in Section 2.) The goal of this work is to present an *approximate EB method* that can significantly reduce the computational cost of solving linear inverse problems with hyperparameters. Specifically, our method evaluates the marginal likelihood function by using a *low-rank update* approximation of the posterior covariance matrix, introduced in previous work for the non-hierarchical setting, e.g., [11, 3, 21, 1, 30, 29]. The intuition behind the low-rank update approximation is that the data may be informative, relative to the prior, only on a *low-dimensional subspace* of the entire parameter space. As a result, one can consider approximations of the posterior covariance matrix in the form of low-rank negative semidefinite updates of the prior covariance matrix, where the update is obtained by solving a generalized eigenproblem involving the log-likelihood Hessian and the prior precision matrix. The optimality properties of this method are studied in [30], which shows that this low-rank approximation is optimal with respect to a specific class of loss functions.

Using the approximation of the posterior covariance developed in [11, 3, 30], we introduce a new method to efficiently compute the marginal likelihood function. We prove that the proposed method yields an *optimal* approximation of the marginal likelihood function in a minimax sense. Our last contribution lies in the numerical implementation. Unlike the inverse problems considered in the work mentioned above, where the prior is fixed, our problems require repeatedly computing the square roots of different prior covariance matrices—which can be prohibitively expensive when the dimensionality of the problem is high. To address the issue, we use the spectral approximation developed in [15] to compute the square root of the prior covariance matrix, resulting in a very efficient algorithm for evaluating (and maximizing) the marginal likelihood function.

The rest of this paper is organized as follows. In Section 2 we introduce the basic

setup of the empirical Bayesian method for solving inverse problems. The low-rank update approximation for evaluating the log marginal likelihood function is developed and analyzed in Section 3, and the detailed numerical implementation of the proposed method is discussed in Section 4. Section 5 provides two numerical examples to demonstrate the performance of the proposed method, and Section 6 offers concluding remarks.

2. Problem setup

Consider a linear inverse problem in a standard setting:

$$\mathbf{y} = G\mathbf{x} + \boldsymbol{\eta}$$

where $\mathbf{y} \in \mathbb{R}^m$ is the data, $\mathbf{x} \in \mathbb{R}^n$ is the unknown, G is a $n \times m$ matrix, often called the forward operator, and $\boldsymbol{\eta}$ is the observation noise. Such an inverse problem can be solved with a Bayesian inference method: namely we assume that the prior probability density of $\mathbf{x}$ is $\pi(\mathbf{x})$ and the likelihood function is $\pi(\mathbf{y}|\mathbf{x})$, and thus the posterior is given by Bayes' theorem:

$$\pi(\mathbf{x}|\mathbf{y}) = \frac{\pi(\mathbf{y}|\mathbf{x})\pi(\mathbf{x})}{\pi(\mathbf{y})}.$$

Throughout this paper, we use the notation $\pi(\cdot)$ as a generic representation of probability density; the specific density is made clear by its arguments. We further assume a Gaussian likelihood and a Gaussian prior with a non-singular covariance matrix $\Gamma_{\text{pr}} \succ 0$ and, without loss of generality, zero mean:

$$\mathbf{y} \mid \mathbf{x} \sim \mathcal{N}(G\mathbf{x}, \Gamma_{\text{obs}}), \quad \mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Gamma_{\text{pr}}). \quad (1)$$

In this setting, the resulting posterior distribution is also Gaussian

$$\mathbf{x} \mid \mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_{\text{pos}}, \Gamma_{\text{pos}}),$$

with mean and covariance matrix given by

$$\boldsymbol{\mu}_{\text{pos}} = \Gamma_{\text{pos}} G^\top \Gamma_{\text{obs}}^{-1} \mathbf{y} \quad \text{and} \quad \Gamma_{\text{pos}} = (H + \Gamma_{\text{pr}}^{-1})^{-1}, \quad (2)$$

where

$$H = G^\top \Gamma_{\text{obs}}^{-1} G \quad (3)$$

is the Hessian of the log-likelihood.

Now we consider a more complex situation, where the matrices G , Γ_{pr} and Γ_{obs} (or some of them) depend on a vector of unspecified hyperparameters $\boldsymbol{\theta} \in \mathbb{R}^p$. As

mentioned in the introduction, a popular method to deal with such problems is the empirical Bayesian (EB) approach, which determines $\boldsymbol{\theta}$ by maximizing the marginal likelihood function:

$$\max_{\boldsymbol{\theta}} \pi(\mathbf{y}|\boldsymbol{\theta}) = \int \pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x} \quad (4)$$

or the marginal posterior density:

$$\max_{\boldsymbol{\theta}} \pi(\boldsymbol{\theta}|\mathbf{y}) = p(\boldsymbol{\theta}) \int \pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x}$$

where $p(\boldsymbol{\theta})$ is the prior density of $\boldsymbol{\theta}$. Note that the computations of these two objectives are very similar and so we shall only discuss (4). It is easy to see that the optimization problem in (4) is equivalent to minimizing the negative log marginal likelihood:

$$\begin{aligned} \min_{\boldsymbol{\theta}} -\log \pi(\mathbf{y}|\boldsymbol{\theta}) \\ = \frac{1}{2} \mathbf{y}^T \Gamma_{\text{obs}}^{-1} \mathbf{y} + \frac{1}{2} \log |\Gamma_{\text{obs}}| - \frac{1}{2} \mathbf{z}^T \Gamma_{\text{pos}} \mathbf{z} + \frac{1}{2} \log \frac{|\Gamma_{\text{pr}}|}{|\Gamma_{\text{pos}}|}, \end{aligned} \quad (5)$$

where G , Γ_{pr} and Γ_{obs} depend on $\boldsymbol{\theta}$ and $\mathbf{z} = G^T \Gamma_{\text{obs}}^{-1} \mathbf{y}$. Note that an important special case arises when only the prior depends on the hyperparameter $\boldsymbol{\theta}$ and the likelihood function is fixed; in this case, we can simply minimize

$$L(\boldsymbol{\theta}, \mathbf{z}) := -\frac{1}{2} \mathbf{z}^T \Gamma_{\text{pos}} \mathbf{z} + \frac{1}{2} \log \frac{|\Gamma_{\text{pr}}|}{|\Gamma_{\text{pos}}|}. \quad (6)$$

For conciseness, below we will present our method for (6), while noting that all the results can be trivially extended to (5).

Direct evaluation of (6) is not desirable for large scale problems, as it requires several operations with $O(n^3)$ complexity. In what follows, we present an accurate and efficient—with $O(n^2 r)$ complexity for some $r \ll n$ —method to approximate $L(\boldsymbol{\theta})$, based on a rank- r update approximation of Γ_{pos} .

3. The optimal approximation method

3.1. Optimal low-rank update approximation of the posterior covariance

The proposed method begins with the optimal low-rank update approximation of posterior covariance developed in [3, 30], which is briefly reviewed here. Note that Γ_{pos} can be written as a non-positive update of Γ_{pr} :

$$\Gamma_{\text{pos}} = \Gamma_{\text{pr}} - K K^T,$$

where

$$KK^\top = \Gamma_{\text{pr}} G^\top \Gamma_y^{-1} G \Gamma_{\text{pr}}$$

and $\Gamma_y = \Gamma_{\text{obs}} + G \Gamma_{\text{pr}} G^\top$ is the covariance of the marginal distribution of $\mathbf{y}$. This update of Γ_{pr} is negative semidefinite because the data add information; they cannot increase the prior variance in any direction. As has been discussed in previous work [30], in many practical problems, the low-rank structure often lies in the update of Γ_{pr} that yields Γ_{pos} , rather than in Γ_{pos} itself. Hence, [30] proposes to use the following class of positive definite matrices to approximate Γ_{pos} :

$$\mathcal{M}_r = \{ \Gamma_{\text{pr}} - BB^\top \succ 0 : \text{rank}(B) \leq r \}. \quad (7)$$

To establish optimality statements regarding the approximation of a covariance matrix, [30] adopts the following distance between symmetric positive definite (SPD) matrices of size n :

$$d_{\mathcal{F}}^2(A, B) = \text{tr} [\ln^2(A^{-1/2} B A^{-1/2})] = \sum_{i=1}^n \ln^2(\sigma_i),$$

where $(\sigma_i)_{i=1}^n$ are the generalized eigenvalues of the pencil (A, B) . This metric was first introduced by Rao in [26], and can measure the difference between two Gaussian distributions with the same mean. We direct interested readers to [29] for a detailed discussion and other applications of this metric. It is important to note that this distance is generalized in [30] to be

$$\mathcal{D}(A, B) = \sum_{i=1}^n f(\sigma_i),$$

where f is a function in $C^1(\mathbb{R}^+)$ that satisfies $f'(x)(1-x) < 0$ for any $x \neq 1$, and $\lim_{x \rightarrow \infty} f(x) = \infty$. This generalization will be used for the proof of our optimality statement in next section. Thus we seek a low-rank approximation $\hat{\Gamma}_{\text{pos}}$ of the covariance matrix Γ_{pos} , such that

$$\hat{\Gamma}_{\text{pos}} = \arg \min_{\Gamma \in \mathcal{M}} \mathcal{D}(\Gamma_{\text{pos}}, \Gamma). \quad (8)$$

The solution of (8) can be derived analytically regardless of the specific choice of f . Specifically, let S_{pr} be any square root of the prior covariance matrix such that $\Gamma_{\text{pr}} = S_{\text{pr}} S_{\text{pr}}^\top$. We define the prior-preconditioned Hessian as

$$\hat{H} = S_{\text{pr}}^\top H S_{\text{pr}}, \quad (9)$$

which plays an essential role in our numerical method. Now let $(\delta_i^2, \mathbf{w}_i)$ be the eigenvalue-eigenvector pairs of $\hat{H}$ with the ordering $\delta_i^2 \geq \delta_{i+1}^2$; then a solution of (8) is given by:

$$\hat{\Gamma}_{\text{pos}} = \Gamma_{\text{pr}} - BB^\top, \quad BB^\top = \sum_{i=1}^r \delta_i^2 (1 + \delta_i^2)^{-1} \hat{\mathbf{w}}_i \hat{\mathbf{w}}_i^\top, \quad \hat{\mathbf{w}}_i = S_{\text{pr}} \mathbf{w}_i. \quad (10)$$

The corresponding minimum distance is

$$\mathcal{D}(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}}) = f(1)r + f \sum_{i=r+1}^n f(1/(1 + \delta_i^2)).$$

The minimizer (10) is unique if the first r eigenvalues of $\hat{H}$ are distinct.

3.2. Approximating the log-likelihood function

Now we apply the optimal low rank approximation to our problem. The idea is rather straightforward: we approximate L in (6) with

$$\hat{L}(\boldsymbol{\theta}, \mathbf{z}) := -\frac{1}{2} \mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z} + \frac{1}{2} \log \frac{|\Gamma_{\text{pr}}|}{|\hat{\Gamma}_{\text{pos}}|}, \quad (11)$$

for some approximate posterior covariance matrix $\hat{\Gamma}_{\text{pos}}$. Next we shall derive the appropriate choice of $\hat{\Gamma}_{\text{pos}}$. To do this, we need to impose an additional assumption on the approximate posterior covariance matrix: $\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}} \succeq 0$, which means that the approximation itself should not create any new information. As a result, the class of matrices for approximating Γ_{pos} becomes

$$\mathcal{M}'_r = \left\{ \hat{\Gamma}_{\text{pos}} = (\Gamma_{\text{pr}} - BB^\top) : \hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}} \succeq 0, \text{rank}(B) \leq r \right\} \quad (12)$$

for some maximum rank r . Next we shall consider the approximations of the two terms $\log(|\Gamma_{\text{pr}}|/|\Gamma_{\text{pos}}|)$ and $\mathbf{z}^\top \Gamma_{\text{pos}} \mathbf{z}$ in (6) separately.

3.2.1. Approximating the log-determinant term. First, we consider finding a matrix $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$ to approximate $\log(|\Gamma_{\text{pr}}|/|\Gamma_{\text{pos}}|)$ with $\log(|\Gamma_{\text{pr}}|/|\hat{\Gamma}_{\text{pos}}|)$. In this setting, it is easy to see that the approximation error in the log-marginal likelihood (6) is $|\log(|\Gamma_{\text{pos}}|/|\hat{\Gamma}_{\text{pos}}|)|$, and a natural way to determine $\hat{\Gamma}_{\text{pos}}$ is to seek a $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$ that minimizes this approximation error. To this end, we have the following theorem:

Theorem 3.1 Suppose that we approximate $\log(|\Gamma_{\text{pr}}|/|\Gamma_{\text{pos}}|)$ with $\log(|\Gamma_{\text{pr}}|/|\hat{\Gamma}_{\text{pos}}|)$ for some $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$. The matrix $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$ that minimizes the resulting approximation error, i.e., the solution of

$$\min_{\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r} \left| \log \frac{|\Gamma_{\text{pos}}|}{|\hat{\Gamma}_{\text{pos}}|} \right|, \quad (13)$$

is given by (10). Moreover, the optimal approximation and the associated error are, respectively,

$$\log \frac{|\Gamma_{\text{pr}}|}{|\hat{\Gamma}_{\text{pos}}|} = \sum_{i=1}^r \log(1 + \delta_i^2) \text{ and } \log \frac{|\hat{\Gamma}_{\text{pos}}|}{|\Gamma_{\text{pos}}|} = \sum_{i=r+1}^n \log(1 + \delta_i^2). \quad (14)$$

Proof We prove this theorem using the optimality results in [30]. To start, we choose a particular distance metric by letting

$$f(x) = |\log x|.$$

We denote the resulting distance metric as $\mathcal{D}_1(A, B)$ to indicate that such a metric is actually the 1-norm of $(\ln(\sigma_1), \dots, \ln(\sigma_n))$ while $d_{\mathcal{F}}$ is the 2-norm of the same sequence. It can be verified that

$$\log \frac{|\Gamma_{\text{pr}}|}{|\Gamma_{\text{pos}}|} = \sum_{i=1}^n \log(1 + \delta_i^2) = \mathcal{D}_1(\Gamma_{\text{pr}}, \Gamma_{\text{pos}}),$$

as $\log(1 + \delta_i^2) \geq 0$ for all $i \in \mathbb{N}$. Since the approximate posterior covariance $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$, we can show

$$\mathcal{D}_1(\Gamma_{\text{pr}}, \Gamma_{\text{pos}}) = \mathcal{D}_1(\Gamma_{\text{pr}}, \hat{\Gamma}_{\text{pos}}) + \mathcal{D}_1(\hat{\Gamma}_{\text{pos}}, \Gamma_{\text{pos}}), \quad (15)$$

where $\mathcal{D}_1(\Gamma_{\text{pr}}, \hat{\Gamma}_{\text{pos}}) = \log(|\Gamma_{\text{pr}}|/|\hat{\Gamma}_{\text{pos}}|)$ is the approximation and $\mathcal{D}_1(\hat{\Gamma}_{\text{pos}}, \Gamma_{\text{pos}}) = \log(|\hat{\Gamma}_{\text{pos}}|/|\Gamma_{\text{pos}}|)$ is the error associated with it. Note that (15) does not hold without the assumption $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$. Thus (13) can be rewritten as,

$$\min_{\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r} \mathcal{D}_1(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}}). \quad (16)$$

Recall that the solution of $\min_{\hat{\Gamma}_{\text{pos}} \in \mathcal{M}_r} \mathcal{D}_1(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}})$ is given by (10), and it is easy to verify that the matrix $\hat{\Gamma}_{\text{pos}}$ given by (10) is in $\mathcal{M}'_r$, which implies that (10) also provides the solution of (16). As a result, the optimal approximation and the associated approximation error are given by (14), which completes the proof.

3.2.2. Approximating the quadratic term. Similarly, we can also find an approximate posterior covariance $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$ and approximate $\mathbf{z}^\top \Gamma_{\text{pos}} \mathbf{z}$ with $\mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z}$. This problem is a bit more complicated: since $\mathbf{z}$ (which is a linear transformation of the data $\mathbf{y}$) is random, we cannot determine the matrix $\hat{\Gamma}_{\text{pos}}$ by directly minimizing the approximation error. In this case, a useful treatment is to apply the minimax criterion, i.e., to seek a matrix $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$ that minimizes the maximum approximation error with respect to $\mathbf{z}$. In particular, for the maximum error to exist, we shall require $\mathbf{z}$ to be bounded: $\mathbf{z} \in Z_c = \{\mathbf{z} : \|\mathbf{z}\|_2 \leq c\}$ for a constant $c > 0$. (See Remark 3.1 for a discussion of this boundedness assumption.) The following theorem provides the optimal solution to this problem.

Theorem 3.2 *Suppose that we approximate $\mathbf{z}^\top \Gamma_{\text{pos}} \mathbf{z}$ with $\mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z}$, for some $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$. The matrix $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$ that achieves the minimax approximation error, i.e., the solution of*

$$\min_{\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r} \max_{\mathbf{z} \in Z_c} |\mathbf{z}^\top \Gamma_{\text{pos}} \mathbf{z} - \mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z}| \quad (17)$$

is given by (10). Moreover, the resulting approximation is

$$\mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z} = \mathbf{z}^\top \Gamma_{\text{pr}} \mathbf{z} - \mathbf{z}^\top \mathbf{b}, \quad \mathbf{b} = \sum_{i=1}^r \frac{\delta_i^2}{1 + \delta_i^2} (\hat{\mathbf{w}}_i^\top \mathbf{z}) \hat{\mathbf{w}}_i, \quad (18)$$

and the associated approximation error is

$$|\mathbf{z}^\top (\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}}) \mathbf{z}| = \mathbf{z}^\top \left(\sum_{i=r+1}^n \frac{\delta_i^2}{1 + \delta_i^2} (\hat{\mathbf{w}}_i^\top \mathbf{z}) \hat{\mathbf{w}}_i \right). \quad (19)$$

Proof For any given $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$, it is easy to see that the solution of

$$\max_{\mathbf{z} \in Z_c} |\mathbf{z}^\top \Gamma_{\text{pos}} \mathbf{z} - \mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z}| = \mathbf{z}^\top (\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}}) \mathbf{z}$$

is $\mathbf{z} = c \mathbf{v}_{\text{max}}$ where $\mathbf{v}_{\text{max}}$ is the eigenvector of the largest eigenvalue of $\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}}$, and the maximum error is $c^2 \rho(\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}})$ where $\rho(\cdot)$ is the spectral radius. Thus (17) becomes $\min_{\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r} \rho(\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}})$, or equivalently,

$$\min_{\text{rank}(B) \leq r} \rho(KK^\top - BB^\top),$$

and it follows immediately that the optimal $\hat{\Gamma}_{\text{pos}}$ is given by (10). Substituting (10) into $\mathbf{z}^\top \hat{\Gamma}_{\text{pos}} \mathbf{z}$ and $\mathbf{z}^\top (\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}}) \mathbf{z}$ yields (18) and (19), respectively.

In principle, of course, $\|\mathbf{z}\|_2$ is not bounded from above, since $\mathbf{z}$ is normally distributed. However, imposing the boundedness assumption, which considerably simplifies the theoretical analysis, does not limit the applicability of the method or affect the optimal solution. Specifically, we have the following remarks:

Remark 3.1 *First, since $\mathbf{z}$ follows a Gaussian distribution, one can always choose a constant c such that the inequality $\|\mathbf{z}\|_2 \leq c$ holds with probability arbitrarily close to one. Second, the minimax solution $\hat{\Gamma}_{\text{pos}}$ is independent of the value of c .*

Now we combine the two approximate treatments, which is essentially to approximate (6) by (11) with $\hat{\Gamma}_{\text{pos}}$ given by (10). It is easy to see that the approximation error is

$$\Delta L(\boldsymbol{\theta}, \mathbf{z}) = \frac{1}{2} \mathbf{z}^\top (\hat{\Gamma}_{\text{pos}} - \Gamma_{\text{pos}}) \mathbf{z} + \frac{1}{2} \log(|\hat{\Gamma}_{\text{pos}}|/|\Gamma_{\text{pos}}|), \quad (20)$$

and we have the following result regarding its optimality:

Corollary 3.3 *Suppose that we approximate (6) with (11) for some matrix $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r$. The matrix $\hat{\Gamma}_{\text{pos}}$ given by (10) achieves the minimax approximation error, i.e., it solves*

$$\min_{\hat{\Gamma}_{\text{pos}} \in \mathcal{M}'_r} \max_{\mathbf{z} \in Z_c} |\Delta L(\boldsymbol{\theta}, \mathbf{z})|.$$

As both terms on the right hand side of (20) are nonnegative, the corollary is a direct consequence of Theorems 3.1 and 3.2.

4. Numerical implementation

Here we discuss the numerical implementation of our approximate method to evaluate (6). In principle, this involves two computationally intensive components—both requiring $O(n^3)$ computations under standard numerical treatments. The first task is to compute the eigenvalues of $\hat{H}$. Recall that our method only requires the r leading eigenvalues and associated eigenvectors of $\hat{H}$, which can be computed efficiently with a randomized algorithm for the singular value decomposition (SVD) [20, 13]. The second task is to compute S_{pr} , the square root of Γ_{pr} . As will become clear later, we do not necessarily need S_{pr} ; rather, our algorithm only requires the ability to evaluate $S_{\text{pr}}\Omega$ for a given matrix Ω . To this end, we resort to the approximation method proposed in [15]. We provide a brief description of the two adopted methods here, tailored according to our specific purposes.

4.1. Randomized SVD

The main idea behind the randomized SVD approach is to identify a subspace that captures most of the action of a matrix, using random sampling [20, 13, 10]. The original matrix is then projected onto this subspace, yielding a smaller matrix, and a standard SVD is then applied to the smaller matrix to obtain the leading eigenvalues of the original matrix. A randomized algorithm for constructing the rank- r' SVD of an $n \times n$ matrix $\hat{H}$, given in [13], proceeds as follows:

- (i) Draw an $n \times r'$ Gaussian random matrix Ω ;
- (ii) Form the $n \times r'$ sample matrix $Y = \hat{H}\Omega$;
- (iii) Form an $n \times r'$ matrix Q with orthonormal columns, such that $Y = QR$;
- (iv) Form the $r' \times n$ matrix $B = Q^T \hat{H}$;
- (v) Compute the SVD of the $r' \times n$ matrix $B = \hat{U}\Sigma V^T$.

According to [13], for this algorithm to be robust, it is important to oversample a bit, namely to choose $r' > r$ if r is the desired rank. One can obtain a probabilistic error bound, i.e., that

$$\|\hat{H} - QB\|_2 \leq (1 + 11\sqrt{r'n})\delta_{r+1}, \quad (21)$$

holds with a probability at least $1 - 6(r' - r)^{-(r' - r)}$, under some very mild assumptions on r' [13]. Note that the diagonal entries of Σ are the eigenvalues of $\hat{H}$ and the columns of V are the eigenvectors. Since we obtain r' eigenpairs in the algorithm, we take $\{\delta_i, \mathbf{w}_i\}_{i=1}^r$ to be the r dominant ones among them. Above we only present the basic implementation of the randomized SVD method; further details and possible improvements of the method can be found in [13] and the references therein. Finally we also want to emphasize that our approximate EB method can use the eigenvalues/eigenvectors computed with any numerical approach; it is not tailored to or tied with the randomized SVD.

4.2. The Chebyshev spectral method for computing S_{pr}

Recall that $\hat{H} = S_{\text{pr}}^\top H S_{\text{pr}}$, and as a result the randomized SVD only requires evaluating $S_{\text{pr}}\Omega$ where Ω is a randomly generated $n \times r'$ matrix. Here we adopt the approximate method proposed in [15], which is based on the following lemma:

Lemma 4.1 (Lemma 2.1 in [15]) *Suppose that D is a real symmetric positive definite matrix. Then there exists a polynomial $p(\cdot)$ such that $\sqrt{D} = p(D)$, and the degree of p is equal to the number of distinct eigenvalues of D minus 1.*

Though $\sqrt{D}$ is exactly equal to a polynomial of D , the degree of the polynomial might be very large if D has a large number of distinct eigenvalues. Thus, instead of trying to find the exact polynomial p in D that equals $\sqrt{D}$, the aforementioned work computes a Chebyshev approximation to it.

It is clear that $D = O\Lambda O^T$, where $\Lambda = \text{diag}(\lambda_1 \cdots \lambda_n)$, and $\lambda_1 \geq \cdots \geq \lambda_n$. Now suppose that we have a polynomial $p(\cdot)$ such that $p(\lambda) \approx \sqrt{\lambda}$. Then we have

$$\begin{aligned} p(D) &= O p(\Lambda) O^T \\ &= O \text{diag}(p(\lambda_1) \cdots p(\lambda_n)) O^T \\ &\approx O \text{diag}(\sqrt{\lambda_1} \cdots \sqrt{\lambda_n}) O^T \\ &= \sqrt{D}. \end{aligned}$$

That is, once we have an approximation of $\sqrt{\lambda}$ in the interval $[\lambda_{\min}, \lambda_{\max}]$, where $\lambda_{\min}$ and $\lambda_{\max}$ are respectively lower and upper bounds on the eigenvalues of D , we immediately get an approximation of $\sqrt{D}$. A popular method to construct the approximation $p(\cdot)$ is the Chebyshev polynomial interpolation. The standard Chebyshev polynomials are defined in $[-1, 1]$ as [24],

$$T_k(x) = \cos(k \arccos x), \quad \forall k \in \mathbb{N}, \quad (22)$$

and the associated Chebyshev points are given by

$$x_j = \cos\left[\frac{(2j+1)\pi}{2k+2}\right], \quad j = 0, 1, \dots, k.$$

It is well known that the Chebyshev interpolant has spectral accuracy for analytical functions on $[-1, 1]$; see, e.g. [12].

Here since we intend to approximate the function $\sqrt{\lambda}$ over the range $[\lambda_{\min}, \lambda_{\max}]$ rather than $[-1, 1]$, we shall use the scaled and shifted Chebyshev polynomials:

$$\tilde{T}_k(x) = T_k(t_a x + t_b), \quad \forall n \in \mathbb{N}, \quad (23)$$

and the associated Chebyshev points become

$$x_j = \frac{\lambda_{\max} + \lambda_{\min}}{2} + \frac{\lambda_{\max} - \lambda_{\min}}{2} \cos\left(\frac{(2j+1)\pi}{2k+2}\right), \quad j = 0, 1, \dots, k. \quad (24)$$

The interpolant $p_k(x)$ can be expressed as [24],

$$p_k(x) = \sum_{i=0}^k c_i \tilde{T}_i(x) - \frac{c_0}{2}, \quad (25)$$

where the coefficients c_i are given by

$$c_i = \frac{2}{k+1} \sum_{j=1}^{k+1} \sqrt{x_j} \tilde{T}_i(x_j). \quad (26)$$

Algorithm 1: Chebyshev spectral approximation for computing $\sqrt{D}\Omega$

Data: Ω, k
Result: $B \approx \sqrt{D}\Omega$
 Compute the coefficients, t_a, t_b , and $c_0, \dots, c_k$;
 $\Omega_0 := \Omega$;
 $\Omega_1 := t_a D \Omega_0 + t_b \Omega_0$;
for $i = 1$ **to** $k - 1$ **do**
 $\Omega_{i+1} := 2(t_a D \Omega_i + t_b \Omega_i) - \Omega_{i-1}$;
end
 $B := \sum_{i=0}^k c_i \Omega_i - \frac{c_0}{2} \Omega_0$.

It is easy to verify that the scaled and shifted Chebyshev polynomials satisfy the following recurrence relation,

$$\begin{aligned} \tilde{T}_0 &= 1, \\ \tilde{T}_1(x) &= t_a x + t_b, \\ \tilde{T}_{i+1}(x) &= 2(t_a x + t_b) \tilde{T}_i(x) - \tilde{T}_{i-1}(x), \end{aligned} \tag{27}$$

with

$$t_a = \frac{2}{\lambda_{\max} - \lambda_{\min}}, \quad t_b = \frac{\lambda_{\max} + \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}}.$$

Now recall that we actually want to compute $\sqrt{D}\Omega$. Taking advantage of the recurrence relation (27), we obtain,

$$\sqrt{D}\Omega \approx p_k(D)\Omega = \sum_{i=0}^k c_i \tilde{T}_i(D)\Omega - \frac{c_0 \Omega}{2} = \sum_{i=0}^k c_i \Omega_i - \frac{c_0}{2} \Omega_0, \tag{28}$$

where

$$\begin{aligned} \Omega_0 &= \Omega, \\ \Omega_1 &= t_a D \Omega_0 + t_b \Omega_0, \\ \Omega_{i+1} &= 2(t_a D \Omega_i + t_b \Omega_i) - \Omega_{i-1}, \end{aligned}$$

for $i = 1, \dots, k - 1$. The complete procedure for computing $\sqrt{D}\Omega$ with the Chebyshev approximation is given in Algorithm 1.

Finally we provide some remarks regarding the implementation of the method:

- The original algorithm present in [15] is to compute the product of $\sqrt{D}$ and a vector, but as is shown in Algorithm 1, its extension to the computation of the product of $\sqrt{D}$ and a matrix is straightforward.

- The method requires upper and lower bounds on the eigenvalues of D , which are computed with an algorithm based on the safeguarded Lanczos method [15]. As is discussed in [15], these estimates need not be of high accuracy.
- The error bound of the proposed Chebyshev approximation is given by Theorems 3.3 and 3.4 in [15].
- If desired, all the matrix-vector multiplications can be performed with the fast multipole method to further improve efficiency [15].

4.3. The complete algorithm

Now we summarize the complete scheme for constructing the approximation $\hat{L}(\boldsymbol{\theta}, \mathbf{z})$ of (6) with a given rank r , using the methods discussed in Sections 4.1 and 4.2:

- Compute the first r eigenpairs, $\{(\delta_i^2, \mathbf{w}_i)\}_{i=1}^r$, of $\hat{H}(\boldsymbol{\theta}) = S_{\text{pr}} H S_{\text{pr}}$, using the randomized SVD method and the Chebyshev spectral method.
- Let $\hat{\mathbf{w}}_i = S_{\text{pr}} \mathbf{w}_i$ and evaluate

$$\hat{L}(\boldsymbol{\theta}, \mathbf{z}) = \mathbf{z}^\top \Gamma_{\text{pr}} \mathbf{z} - \mathbf{z}^\top \sum_{i=1}^r \frac{\delta_i^2}{1 + \delta_i^2} (\hat{\mathbf{w}}_i^\top \mathbf{z}) \hat{\mathbf{w}}_i + \sum_{i=1}^r \log(1 + \delta_i^2).$$

5. Examples

5.1. An image deblurring problem

We first test our method on an imaging deblurring problem, which involves recovering a latent image from noisy observations of a blurred version of the image [14]. In particular, it is assumed that the blurred image is obtained as a convolution of the latent image with a point spread function (PSF), and as a result the forward model is:

$$y(t_1, t_2) = \int \int_D f_{\text{PSF}}(t_1, t_2) x(\tau_1, \tau_2) d\tau_1 d\tau_2, \quad (29)$$

where $f_{\text{PSF}}(t_1, t_2)$ is the PSF and D is the domain of the image. In this example we take the image domain to be $D = [-1, 1]^2$ and the PSF to be

$$f_{\text{PSF}}(t_1, t_2) = \exp[-((t_1 - \tau_1)^2 + (t_2 - \tau_2)^2)/t],$$

where t is a parameter controlling the size of the spreading. Moreover, we assume that the data y are measured at $m = 64^2 = 4096$ observation locations evenly distributed in D , and that the observation errors are mutually independent and Gaussian with zero

mean and variance $(0.1)^2$. We represent the unknown x on 256×256 mesh grid points, and thus the dimensionality of the inverse problem is $n = 65536$. The prior on x is a Gaussian distribution with zero mean and covariance kernel [31, 27]:

$$K(\mathbf{t}, \mathbf{t}') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{d}{\rho} \right)^\nu B_\nu \left(\sqrt{2\nu} \frac{d}{\rho} \right), \quad (30)$$

where $d = \|\mathbf{t} - \mathbf{t}'\|_2$, $\Gamma(\cdot)$ is the Gamma function, and $B_\nu(\cdot)$ is the modified Bessel function. (30) is known as the Matérn covariance, and several authors have suggested that such covariances can often provide better models for many realistic physical processes than the popular squared exponential covariance [31, 27]. A random function with the Matérn covariance is $[\nu - 1]$ mean-square (MS) differentiable. In this example, we choose $\nu = 3$ implying second order MS differentiability. We set the standard deviation σ to one. The correlation length ρ is treated as a hyperparameter to be inferred in this example.

We now use the proposed empirical Bayesian method to solve the inverse problem. We first assume that the true correlation length is $\rho = 1$, and randomly generate a true image from the associated prior, shown in Fig. 1 (left). We then test two cases of the forward problem, with $t = 0.002$ and $t = 0.02$, where the latter yields an inverse problem that is much more ill-posed than the former. We assume that the data are observed on a 128×128 uniformly distributed mesh. We then apply the two convolution operators to the generated image and add observational noise to the results, producing the synthetic data also shown in Fig. 1.

Using the proposed approximate EB method, we evaluate the negative log marginal likelihood function L over a range of ρ values, for both cases of t , and plot the results in Fig 2. For $t = 0.002$ (left figure), we compute L with ranks $r = 300, 400$ and 500 , and 600 , and observe that the curves converge as the rank increases; indeed, the results with $r = 500$ and 600 appear identical. For $t = 0.02$ —because the problem is more ill-posed—we can implement the method with ranks $r = 50, 75, 100$, and 150 and observe convergence. In particular, while the results of $r = 50$ deviate from the others, the results with $r = 75, 100$, and 150 look nearly identical, implying that $r = 75$ is sufficient for an accurate approximation of the marginal likelihood in this case. For both cases, the optimal value of ρ is found to be 0.1 , which is actually the true hyperparameter value. We then compute the posterior mean of x , after fixing $\rho = 0.1$ in the Gaussian prior on x , and show the results in Fig. 3.

Note that the EB method is able to find the appropriate hyperparameter values in this example. Nevertheless, our intention is not to illustrate that the EB method can always identify the correct value of the hyperparameters; rather, the main purpose

Figure 1. From left to right: the ground truth, the simulated data at $t = 0.002$ and that at $t = 0.02$.

Figure 2. The negative log marginal posterior function plotted against ρ , for a series of ranks r indicated in the legend. Left: $t = 0.002$; right: $t = 0.02$.

Figure 3. Posterior mean for the $t = 0.002$ case. Left: result for $\rho = 0.1$ (optimal); right: result for $\rho = 0.2$.

Figure 4. Posterior mean for the $t = 0.02$ case. Left: result for $\rho = 0.1$ (optimal); right: result for $\rho = 0.02$.

of the example is to demonstrate that, should one *choose* to use the EB method, the proposed approximations can efficiently and accurately compute the marginal likelihood.

5.2. X-ray computed tomography

Our second example is an X-ray computed tomography (CT) problem. X-ray CT is a popular diagnostic tool for medical tomographic imaging of the human body. It provides images of the object by assigning an X-ray attenuation coefficient to each pixel [22]. Specifically let u denote the unknown image supported in the unit ball $B(\mathbf{0}, 1)$ in $\mathbb{R}^2$. In the case of two-dimensional parallel beam CT, the projection data (or sinogram) f for each $\varphi \in [0, 2\pi)$ and $s \in \mathbb{R}$ is obtained via the following Radon transform [25]:

$$f(\varphi, s) = \int_{-\infty}^{\infty} u(s\boldsymbol{\theta} + t\boldsymbol{\theta}^{\perp})dt \quad (31)$$

where $\boldsymbol{\theta} = (\cos \varphi, \sin \varphi)$ and $\boldsymbol{\theta}^{\perp} = (-\sin \varphi, \cos \varphi)$. We test the problem using a 256×256 ground truth image shown in Fig. 5 (left), which is taken from the Harvard Whole Brain Atlas [16]. The projection data, shown in Fig. 5 (right), is simulated by plugging the true image into the Radon transform and adding measurement noise to the result. The noise is taken to be Gaussian with zero mean and a signal-to-noise ratio of 1%. The size of the discrete Radon transform operator depends on both the size of the image and the number of projections. In our simulation, we used 60 projections

equispatially sampled from 0 to π . To avoid an inverse crime [18], we recover the image with a 128×128 resolution, and thus the dimensionality of the inverse problem is 128^2 . We take the prior distribution to be zero-mean Gaussian, with a covariance matrix given by an anisotropic Matérn kernel:

$$K(\mathbf{t}, \mathbf{t}') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} d(\mathbf{t}, \mathbf{t}') \right)^\nu B_\nu \left(\sqrt{2\nu} d(\mathbf{t}, \mathbf{t}') \right), \quad (32)$$

where

$$d(\mathbf{t}, \mathbf{t}') = \sqrt{\frac{(t_1 - t'_1)^2}{\rho_1^2} + \frac{(t_2 - t'_2)^2}{\rho_2^2}}.$$

Now we consider all four parameters ν , σ , ρ_1 and ρ_2 as hyperparameters to be determined, and we also assume that the variance ϵ^2 of the measurement noise is not available. We will identify the five parameters using the proposed low-rank approximate EB method.

We solve the optimization problems with 9 different ranks ($r = 500, 1000, 1500, 2000, 2500, 3000, 3500, 4000, 4500$), and plot the optimal values of the five hyperparameters and the associated value of L versus the rank r in Fig. 6. As one can see, the obtained optimal values of the hyperparameters converge as the rank increases, and in particular $r = 4000$ is sufficient to obtain accurate estimates of the hyperparameters. Next we shall illustrate that the obtained hyperparameter values can lead to rather good inference results. To show this, we compute the posterior mean and the variance of the image using the hyperparameter values computed with the various ranks, and we then compute the peak noise-to-signal ratio (PSNR) of the posterior means against the true image, which is a commonly used measure of quality for image reconstruction. Note that the low rank approximation is only used to obtain the hyperparameters; given these hyperparameter values, the posterior means and variances are computed directly without approximation. We plot the PSNR as a function of the rank r in Fig. 7. We also show the posterior mean and variance field associated with each rank in the plot. One can see from the figure that, as the rank increases, the resulting hyperparameters can recover images with higher quality in terms of PSNR. Finally, in Fig. 8, we plot the posterior mean and variance using the hyperparameters computed with $r = 4500$, which are taken to be the inference results of this problem.

These results demonstrate that the outcome of Bayesian inference can depend critically on the hyperparameters, and that, at least in this example, the proposed method can provide reasonable values for these parameters, provided that a sufficiently

Figure 5. Left: the ground truth image. Right: the projection data simulated from the ground truth.

Figure 6. The optimal value of the five hyperparameters and the associated value of L (bottom, right) plotted against the rank r used in the EB approximation.

large rank is used. We also note here that, to better recover images with sharp edges, one might resort to more complex Gaussian hypermodels [2, 4, 5] than the one considered here, or even non-Gaussian models such as the Besov [9, 19] and TV-Gaussian [34] priors. Nevertheless, if one has chosen to use a Gaussian hypermodel, our method can efficiently determine the values of the associated hyperparameters.

Figure 7. The PSNR of the posterior mean versus the rank r (red line). Above and below the line, we also show the posterior mean and variance obtained with the hyperparameters computed with each rank.

Figure 8. The posterior mean and the posterior variance computed using the optimized hyperparameter values with $r = 4500$.

6. Conclusions

This paper investigates empirical Bayesian methods for determining hyperparameters in linear inverse problems. We develop an efficient numerical method to approximately evaluate the marginal likelihood function of the hyperparameters for large-scale problems, based on a low-rank approximation of the update from the prior to the posterior covariance. The proposed method can achieve a computational complexity of $O(n^2r)$ for $r \ll n$, while a standard full-rank or direct evaluation approach requires computations of $O(n^3)$ complexity. We also show that this approximation of the marginal likelihood is optimal in a minimax sense. Two numerical examples illustrate that the proposed algorithm can accurately evaluate the marginal likelihood function required in the EB method. This approach may be useful in a wide range of applications where the unknown is of very high dimension, such as medical and geophysical image reconstruction. Finally, it is worth noting that, while the approximate EB method presented in this work can only be applied to linear inverse problems, there have been several efforts successfully extending low-rank posterior approximations to nonlinear problems [21, 8]; along these lines, we expect that our method could also be extended to EB computations for nonlinear inverse problems.

Acknowledgement

Li was partially supported by the NSFC under grant number 11771289. Marzouk was partially supported by the US Department of Energy, Office of Advanced Scientific Computing Research, under grant number de-sc0009297.

- [1] Alen Alexanderian, Philip J Gloor, Omar Ghattas, et al. On Bayesian A-and D-optimal experimental designs in infinite dimensions. *Bayesian Analysis*, 11(3):671–695, 2016.
- [2] Johnathan M Bardsley, Daniela Calvetti, and Erkki Somersalo. Hierarchical regularization for edge-preserving reconstruction of PET images. *Inverse Problems*, 26(3):035010, 2010.
- [3] Tan Bui-Thanh, Omar Ghattas, James Martin, and Georg Stadler. A computational framework for infinite-dimensional Bayesian inverse problems part i: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
- [4] Daniela Calvetti and Erkki Somersalo. A Gaussian hypermodel to recover blocky objects. *Inverse problems*, 23(2):733, 2007.
- [5] Daniela Calvetti and Erkki Somersalo. Hypermodels in the Bayesian imaging framework. *Inverse Problems*, 24(3):034013, 2008.
- [6] Bradley P Carlin and Thomas A Louis. *Bayes and empirical Bayes methods for data analysis*. Chapman & Hall/CRC Boca Raton, 2000.
- [7] George Casella. An introduction to empirical Bayes data analysis. *The American Statistician*, 39(2):83–87, 1985.

- [8] Tiangang Cui, James Martin, Youssef M Marzouk, Antti Solonen, and Alessio Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30(11):114015, 2014.
- [9] Masoumeh Dashti, Stephen Harris, and Andrew Stuart. Besov priors for Bayesian inverse problems. *Inverse Problems and Imaging*, 6(2):183–200, 2012.
- [10] Petros Drineas, Ravi Kannan, and Michael W Mahoney. Fast Monte Carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix. *SIAM Journal on Computing*, 36(1):158–183, 2006.
- [11] H Pearl Flath, Lucas C Wilcox, Volkan Akçelik, Judith Hill, Bart van Bloemen Waanders, and Omar Ghattas. Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial hessian approximations. *SIAM Journal on Scientific Computing*, 33(1):407–432, 2011.
- [12] Bengt Fornberg. *A practical guide to pseudospectral methods*, volume 1. Cambridge university press, 1998.
- [13] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [14] Per Christian Hansen, James G Nagy, and Dianne P O’leary. *Deblurring images: matrices, spectra, and filtering*. SIAM, 2006.
- [15] Shidong Jiang, Zhi Liang, and Jingfang Huang. A fast algorithm for Brownian dynamics simulation with hydrodynamic interactions. *Mathematics of Computation*, 82(283):1631–1645, 2013.
- [16] Keith A. Johnson and J. Alex Becker. Whole Brain Atlas, <http://www.med.harvard.edu/aanlib/>.
- [17] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160. Springer, 2005.
- [18] Jari Kaipio and Erkki Somersalo. Statistical inverse problems: discretization, model reduction and inverse crimes. *Journal of computational and applied mathematics*, 198(2):493–504, 2007.
- [19] Matti Lassas, Eero Saksman, and Samuli Siltanen. Discretization-invariant Bayesian inversion and besov space priors. *Inverse Problems and Imaging*, 3(1):87–122, 2009.
- [20] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, 2007.
- [21] James Martin, Lucas C Wilcox, Carsten Burstedde, and Omar Ghattas. A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.
- [22] Frank Natterer. *The mathematics of computerized tomography*. SIAM, 2001.
- [23] Sonia Petrone, Judith Rousseau, and Catia Scricciolo. Bayes and empirical Bayes: do they merge? *Biometrika*, page ast067, 2014.
- [24] William H Press. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- [25] Johann Radon. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Mathematisch-Physische Klasse*, 69:262–277, 1917.
- [26] Calyampudi Radhakrishna Rao. Information and accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc*, 37(3):81–91, 1945.

- [27] Carl Edward Rasmussen. *Gaussian processes for machine learning*. MIT Press, 2006.
- [28] Judith Rousseau and Botond Szabo. Asymptotic behaviour of the empirical bayes posteriors associated to maximum marginal likelihood estimator. *arXiv preprint arXiv:1504.04814*, 2015.
- [29] Alessio Spantini, Tiangang Cui, Karen Willcox, Luis Tenorio, and Youssef Marzouk. Goal-oriented optimal approximations of Bayesian linear inverse problems. *arXiv preprint arXiv:1607.01881*, 2016.
- [30] Alessio Spantini, Antti Solonen, Tiangang Cui, James Martin, Luis Tenorio, and Youssef Marzouk. Optimal low-rank approximations of Bayesian linear inverse problems. *SIAM Journal on Scientific Computing*, 37(6):A2451–A2487, 2015.
- [31] Michael L Stein. *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 2012.
- [32] Andrew M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- [33] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. siam, 2005.
- [34] Zhewei Yao, Zixi Hu, and Jinglai Li. A TV-Gaussian prior for infinite-dimensional Bayesian inverse problems and its numerical implementations. *Inverse Problems*, 32(7):075006, 2016.