

PAPER

Solving large-scale PDE-constrained Bayesian inverse problems with Riemann manifold Hamiltonian Monte Carlo

To cite this article: T Bui-Thanh and M Girolami 2014 *Inverse Problems* **30** 114014

View the [article online](#) for updates and enhancements.

Related content

- [Likelihood-informed dimension reduction for nonlinear inverse problems](#)
T Cui, J Martin, Y M Marzouk et al.
- [Well-posed Bayesian geometric inverse problems arising in subsurface flow](#)
Marco A Iglesias, Kui Lin and Andrew M Stuart
- [A data-scalable randomized misfit approach for solving large-scale PDE-constrained inverse problems](#)
E B Le, A Myers, T Bui-Thanh et al.

Recent citations

- [Local search methods for the solution of implicit inverse problems](#)
Eliás D. Nino-Ruiz *et al*
- [Geometry and Dynamics for Markov Chain Monte Carlo](#)
Alessandro Barp *et al*
- [A Randomized Maximum A Posteriori Method for Posterior Sampling of High Dimensional Nonlinear Bayesian Inverse Problems](#)
Kainan Wang *et al*



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Solving large-scale PDE-constrained Bayesian inverse problems with Riemann manifold Hamiltonian Monte Carlo

T Bui-Thanh¹ and M Girolami²

¹Department of Aerospace Engineering and Engineering Mechanics, Institute for Computational Engineering & Sciences, The University of Texas, Austin, TX 78712, USA

²Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK

E-mail: tanbui@ices.utexas.edu and m.girolami@warwick.ac.uk

Received 4 March 2014, revised 18 June 2014

Accepted for publication 30 June 2014

Published 28 October 2014

Abstract

We consider the Riemann manifold Hamiltonian Monte Carlo (RMHMC) method for solving statistical inverse problems governed by partial differential equations (PDEs). The Bayesian framework is employed to cast the inverse problem into the task of statistical inference whose solution is the posterior distribution in infinite dimensional parameter space conditional upon observation data and Gaussian prior measure. We discretize both the likelihood and the prior using the H^1 -conforming finite element method together with a matrix transfer technique. The power of the RMHMC method is that it exploits the geometric structure induced by the PDE constraints of the underlying inverse problem. Consequently, each RMHMC posterior sample is almost uncorrelated/independent from the others providing statistically efficient Markov chain simulation. However this statistical efficiency comes at a computational cost. This motivates us to consider computationally more efficient strategies for RMHMC. At the heart of our construction is the fact that for Gaussian error structures the Fisher information matrix coincides with the Gauss–Newton Hessian. We exploit this fact in considering a computationally simplified RMHMC method combining state-of-the-art adjoint techniques and the superiority of the RMHMC method. Specifically, we first form the Gauss–Newton Hessian at the maximum *a posteriori* point and then use it as a fixed constant metric tensor throughout RMHMC simulation. This eliminates the need for the computationally costly differential geometric Christoffel symbols, which in turn greatly reduces computational effort at a corresponding loss of sampling efficiency. We further reduce the cost of forming the Fisher information matrix by using a low rank approximation via a

randomized singular value decomposition technique. This is efficient since a small number of Hessian-vector products are required. The Hessian-vector product in turn requires only two extra PDE solves using the adjoint technique. Various numerical results up to 1025 parameters are presented to demonstrate the ability of the RMHMC method in exploring the geometric structure of the problem to propose (almost) uncorrelated/independent samples that are far away from each other, and yet the acceptance rate is almost unity. The results also suggest that for the PDE models considered the proposed fixed metric RMHMC can attain almost as high a quality performance as the original RMHMC, i.e. generating (almost) uncorrelated/independent samples, while being two orders of magnitude less computationally expensive.

Keywords: Riemann manifold Hamiltonian Monte Carlo, inverse problems, Bayesian, Gaussian measure

(Some figures may appear in colour only in the online journal)

1. Introduction

Inverse problems are ubiquitous in science and engineering. Perhaps the most popular family of inverse problems is to determine a set of parameters (or a function) given a set of indirect observations, which are in turn provided by a parameter-to-observable map plus observation uncertainties. For example, if one considers the problem of determining the heat conductivity of a thermal fin given measured temperature at a few locations on the thermal fin, then: (i) the desired unknown parameter is the distributed heat conductivity, (ii) the observations are the measured temperatures, (iii) the parameter-to-observable map is the mathematical model that describes the temperature on the thermal fin as a function of the heat conductivity; indeed the temperature distribution is a solution of an elliptic partial differential equation (PDE) whose coefficient is the heat conductivity, and (iv) the observation uncertainty is due to the imperfection of the measurement device and/or model inadequacy.

The Bayesian inversion framework refers to a mathematical method that allows one to solve statistical inverse problems taking into account all uncertainties in a systematic and coherent manner. The Bayesian approach does this by reformulating the inverse problem as a problem in statistical inference, incorporating uncertainties in the observations, the parameter-to-observable map, and prior information on the parameter. In particular, we seek a statistical description of all possible (set of) parameters that conform to the available prior knowledge and at the same time are consistent with the observations. The solution of the Bayesian framework is the so-called posterior measure that encodes the degree of confidence on each set of parameters as the solution to the inverse problem under consideration.

Mathematically the posterior is a surface in high dimensional parameter space. The task at hand is therefore to explore the posterior by, for example, characterizing the mean, the covariance, and/or higher moments. The nature of this task is to compute high dimensional integrals for which most contemporary methods are intractable. Perhaps the most general method to attack these problems is the Markov chain Monte Carlo (MCMC) method which shall be introduced in subsequent sections.

Let us now summarize the content of the paper. We start with the description of the statistical inverse problem under consideration in section 2. It is an inverse steady state heat conduction governed by elliptic PDEs. We postulate a Gaussian measure prior on the

parameter space to ensure that the inverse problem is well-defined. The prior itself is a well-defined object whose covariance operator is the inverse of an elliptic differential operator and with the mean function living in the Cameron–Martin space of the covariance. The posterior is given by its Radon–Nikodym derivative with respect to the prior measure, which is proportional to the likelihood. Since the Riemann manifold Hamiltonian Monte Carlo (RMHMC) simulation method requires the gradient, Hessian, and the derivative of the Fisher information operator, we discuss, in some depth, how to compute the derivatives of the potential function (the misfit functional) with PDE constraints efficiently using the adjoint technique in section 3. In particular, we define a Fisher information operator and show that it coincides with the well-known Gauss–Newton Hessian of the misfit. We next present a discretization scheme for the infinite Bayesian inverse problem in section 4. Specifically, we employ a standard continuous H^1 -conforming finite element method (FEM) to discretize both the likelihood and the Gaussian prior. We choose to numerically compute the truncated Karhunen–Loève (KL) expansion which requires one to solve an eigenvalue problem with fractional Laplacian. In order to accomplish this task, we use a matrix transfer technique (MTT) which leads to a natural discretization of the Gaussian prior measure. In section 5, we describe the RMHMC and its variants at length, and its application to our Bayesian inverse problem. Section 6 presents a low rank approach to approximate the Fisher information matrix and its inverse efficiently. This is possibly due to the fact that the Gauss–Newton Hessian, and hence the Fisher information operator, is a compact operator. Various numerical results supporting our proposed approach are presented in section 7. We begin this section with an extensive study and comparison of Riemannian manifold MCMC methods for problems with two parameters, and end the section with 1025-parameter problem. Finally, we conclude the paper in section 8 with a discussion on future work.

2. Problem statement

In order to clearly illustrate the challenges arising in PDE-constrained inverse problems for MCMC based Bayesian inference, we consider the following heat conduction problem governed by an elliptic PDE in the open and bounded domain $\Omega \subset \mathbb{R}^n$

$$\begin{aligned} -\nabla \cdot (e^u \nabla w) &= 0 && \text{in } \Omega, \\ -e^u \nabla w \cdot \mathbf{n} &= Bi u && \text{on } \partial\Omega \setminus \Gamma_R, \\ -e^u \nabla w \cdot \mathbf{n} &= -1 && \text{on } \Gamma_R, \end{aligned}$$

where w is the forward state, u the logarithm of distributed thermal conductivity on Ω , \mathbf{n} the unit outward normal on $\partial\Omega$, and Bi the Biot number.

In the forward problem, the task is to solve for the temperature distribution w given a description of distributed parameter u . In the inverse problem, the task is to reconstruct u given some available observations, e.g, temperature observed at some parts/locations of the domain Ω . We initially choose to cast the inverse problem in the framework of PDE-constrained optimization. To begin, let us consider the following additive noise-corrupted pointwise observation model³

$$d_j := w(\mathbf{x}_j) + \eta_j, \quad j = 1, \dots, K, \quad (1)$$

³ We assume the forward state w is sufficiently regular, i.e. $w \in H^s$, $s > n/2$, so that w is, by virtue of the Sobolev embedding theorem, continuous, and therefore it is meaningful to measure w pointwise.

where K is the total number of observation locations, $\{\mathbf{x}_j\}_{j=1}^K$ the set of points at which w is observed, η_j the additive noise, and d_j the actual noise-corrupted observations. In this paper we work with synthetic observations and hence there is no model inadequacy in (1). Concatenating all the observations, one can rewrite (1) as

$$\mathbf{d} := \mathcal{G}(u) + \eta, \quad (2)$$

with $\mathcal{G} := [w(\mathbf{x}_1), \dots, w(\mathbf{x}_K)]^T$ denoting the map from the distributed parameter u to the noise-free observables, η being random numbers normally distributed by $\mathcal{N}(0, \mathbf{L})$ with bounded covariance matrix \mathbf{L} , and $\mathbf{d} = [d_1, \dots, d_K]^T$. For simplicity, we take $\mathbf{L} = \sigma^2 \mathbf{I}$, where \mathbf{I} is the identity matrix.

Our inverse problem can be now formulated as

$$\min_u \mathcal{J}(u, \mathbf{d}) := \frac{1}{2} \|\mathbf{d} - \mathcal{G}(u)\|_{\mathbf{L}}^2 = \frac{1}{2\sigma^2} \sum_{j=1}^K (w(\mathbf{x}_j) - d_j)^2, \quad (3)$$

subject to

$$-\nabla \cdot (e^u \nabla w) = 0 \quad \text{in } \Omega, \quad (4a)$$

$$-e^u \nabla w \cdot \mathbf{n} = Bi u \quad \text{on } \partial\Omega \setminus \Gamma_R, \quad (4b)$$

$$-e^u \nabla w \cdot \mathbf{n} = -1 \quad \text{on } \Gamma_R, \quad (4c)$$

where $\|\cdot\|_{\mathbf{L}} := \|\mathbf{L}^{-\frac{1}{2}} \cdot\|$ denotes the weighted Euclidean norm induced by the canonical inner product (\cdot, \cdot) in \mathbb{R}^K . This optimization problem is however ill-posed. An intuitive reason is that the dimension of observations \mathbf{d} is much smaller than that of the parameter u , and hence they provide limited information about the distributed parameter u . As a result, the null space of the Jacobian of the parameter-to-observation map \mathcal{G} is non-empty. Indeed, we have shown that the Gauss–Newton approximation of the Hessian (which is the square of this Jacobian, and is also equal to the full Hessian of the data misfit \mathcal{J} evaluated at the optimal parameter) is a compact operator [1–3], and hence its range space is effectively finite-dimensional.

One way to overcome the ill-posedness is to use *Tikhonov regularization* (see, e.g., [4]), which proposes to augment the cost functional (3) with a quadratic term, i.e.

$$\tilde{\mathcal{J}} := \frac{1}{2} \|\mathbf{d} - \mathcal{G}(u)\|_{\mathbf{L}}^2 + \frac{\kappa}{2} \|R^{1/2}u\|^2, \quad (5)$$

where κ is a regularization parameter, R some regularization operator, and $\|\cdot\|$ some appropriate norm. This method is a representative of deterministic inverse solution techniques that typically do not take into account the randomness due to measurements and other sources, though one can equip the deterministic solution with a confidence region by post-processing (see, e.g., [5] and references therein). It should be pointed out that if the regularization term is replaced by the Cameron–Martin norm of u (the second term in (7)), the Tikhonov solution is in fact identical to the maximum *a posteriori* (MAP) point in (7). However, such a point estimate is insufficient for the purpose of fully taking the randomness into account.

In this paper, we choose to tackle the ill-posedness using a *Bayesian* framework [6–10]. We seek a statistical description of all possible u that conform to some prior knowledge and at the same time are consistent with the observations. The Bayesian approach does this by reformulating the inverse problem as a problem in *statistical inference*, incorporating uncertainties in the observations, the forward map \mathcal{G} , and prior information. This approach is appealing since it can incorporate most, if not all, kinds of randomness in a systematic

manner. To begin, we postulate a Gaussian measure $\mu := \mathcal{N}(u_0, \alpha^{-1}C)$ on u in $L^2(\Omega)$ where

$$C := (I - \Delta)^{-s} =: \mathcal{A}^{-s},$$

with the domain of definition

$$D(\mathcal{A}) := \left\{ u \in H^2(\Omega) : \frac{\partial u}{\partial \mathbf{n}} = 0 \text{ on } \partial\Omega \right\},$$

where $H^2(\Omega)$ is the usual Sobolev space. Assume that the mean function u_0 lives in the Cameron–Martin space of C , then one can show (see [9]) that the measure μ is well-defined when $s > n/2$ (d is the spatial dimension), and in that case, any realization from the prior distribution μ is almost surely in the Hölder space $X := C^{0, \beta}(\Omega)$ with $0 < \beta < s/2$. That is, $\mu(X) = 1$, and the Bayesian posterior measure ν satisfies the Radon–Nikodym derivative

$$\frac{\partial \nu}{\partial \mu}(u|\mathbf{d}) \sim \exp(-\mathcal{J}(u, \mathbf{d})) = \exp\left(-\frac{1}{2} \|\mathbf{d} - \mathcal{G}(u)\|_{\mathbb{L}}^2\right), \quad (6)$$

if \mathcal{G} is a continuous map from X to \mathbb{R}^K . Note that the Radon–Nikodym derivative is proportional to the likelihood defined by

$$\pi_{\text{like}}(\mathbf{d}|u) \sim \exp(-\mathcal{J}(u, \mathbf{d})).$$

The MAP point is defined as

$$u^{\text{MAP}} := \arg \min_u \mathcal{J}(u, \mathbf{d}) := \frac{1}{2} \|\mathbf{d} - \mathcal{G}(u)\|_{\mathbb{L}}^2 + \frac{\alpha}{2} \|u\|_C^2, \quad (7)$$

where $\|\cdot\|_C := \|\cdot\|_{C^{-\frac{1}{2}}}$ denotes the weighted $L^2(\Omega)$ norm induced by the $L^2(\Omega)$ inner product $\langle \cdot, \cdot \rangle$.

3. Adjoint computation of gradient, Hessian, and the third derivative tensor

In this section, we briefly present the adjoint method to efficiently compute the gradient, Hessian, and the third derivative of the cost functional (3). We start by considering the weak form of the (first order) forward equation (4):

$$\int_{\Omega} e^u \nabla w \cdot \nabla \hat{\lambda} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} B_i w \hat{\lambda} \, ds = \int_{\Gamma_R} \hat{\lambda} \, ds, \quad (8)$$

with $\hat{\lambda}$ as the test function. Using the standard reduced space approach (see, e.g., a general discussion in [11] and a detailed derivation in [12]) one can show that the gradient $\nabla \mathcal{J}(u)$, namely the Fréchet derivative of the cost functional \mathcal{J} , acting in any direction \tilde{u} is given by

$$\langle \nabla \mathcal{J}(u), \tilde{u} \rangle = \int_{\Omega} \tilde{u} e^u \nabla w \cdot \nabla \lambda \, d\Omega, \quad (9)$$

where the (first order) adjoint state λ satisfies the adjoint equation

$$\int_{\Omega} e^u \nabla \lambda \cdot \nabla \hat{w} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} B_i \lambda \hat{w} \, ds = -\frac{1}{\sigma^2} \sum_{j=1}^K (w(\mathbf{x}_j) - d_j) \hat{w}(\mathbf{x}_j), \quad (10)$$

with \hat{w} as the test function. On the other hand, the Hessian, the Fréchet derivative of the gradient, acting in directions \tilde{u} and u^2 (superscript ‘2’ means the second variation direction) reads

$$\begin{aligned} \left\langle \left\langle \nabla^2 \mathcal{J}(u), \tilde{u} \right\rangle, u^2 \right\rangle &= \int_{\Omega} \tilde{u} e^u \nabla w \cdot \nabla \lambda^2 \, d\Omega + \int_{\Omega} \tilde{u} e^u \nabla w^2 \cdot \nabla \lambda \, d\Omega \\ &\quad + \int_{\Omega} \tilde{u} u^2 e^u \nabla w \cdot \nabla \lambda \, d\Omega, \end{aligned} \quad (11)$$

where the second order forward state w^2 obeys the second order forward equation

$$\int_{\Omega} e^u \nabla w^2 \cdot \nabla \hat{\lambda} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, w^2 \hat{\lambda} \, ds = - \int_{\Omega} u^2 e^u \nabla w \cdot \nabla \hat{\lambda} \, d\Omega, \quad (12)$$

and the second order adjoint state λ^2 is governed by the second order adjoint equation

$$\int_{\Omega} e^u \nabla \lambda^2 \cdot \nabla \hat{w} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, \lambda^2 \hat{w} \, ds = \frac{1}{\sigma^2} \sum_{j=1}^K w^2(\mathbf{x}_j) \hat{w}(\mathbf{x}_j) - \int_{\Omega} u^2 e^u \nabla \lambda \cdot \nabla \hat{w} \, d\Omega. \quad (13)$$

We define the generalized Fisher information operator⁴ acting in directions \tilde{u} and u^2 as

$$\left\langle \left\langle G(u), \tilde{u} \right\rangle, u^2 \right\rangle := \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} \left[\left\langle \left\langle \nabla^2 \mathcal{J}(u), \tilde{u} \right\rangle, u^2 \right\rangle \right], \quad (14)$$

where the expectation is taken with respect to the likelihood—the distribution of the observation \mathbf{d} . Now, substituting (11) into (14) and assuming that the integrals/derivatives can be interchanged we obtain

$$\begin{aligned} \left\langle \left\langle G(u), \tilde{u} \right\rangle, u^2 \right\rangle &= \int_{\Omega} \tilde{u} e^u \nabla w \cdot \nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} \left[\lambda^2 \right] \, d\Omega + \int_{\Omega} \tilde{u} e^u \nabla w^2 \cdot \nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda] \, d\Omega \\ &\quad + \int_{\Omega} \tilde{u} u^2 e^u \nabla w \cdot \nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda] \, d\Omega, \end{aligned}$$

where we have used the assumption that the parameter u is independent of observation \mathbf{d} and the fact that w and w^2 do not depend on \mathbf{d} . The next step is to compute $\nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda]$ and $\mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda^2]$. To begin, let us take the expectation the first order adjoint equation (10) with respect to $\pi_{\text{like}}(\mathbf{d}|u)$ to arrive at

$$\begin{aligned} \int_{\Omega} e^u \nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda] \cdot \nabla \hat{w} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda] \hat{w} \, ds = \\ - \frac{1}{\sigma^2} \sum_{j=1}^K \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} \left[w(\mathbf{x}_j) - d_j \right] \hat{w}(\mathbf{x}_j) = 0, \end{aligned}$$

where the second equality is obtained from (1) and the assumption $\eta_j \sim \mathcal{N}(0, \sigma^2)$. We conclude that

$$\nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} [\lambda] = 0. \quad (15)$$

On the other hand, if we take the expectation of the second order adjoint equation (13) and use (15) we have

$$\begin{aligned} \int_{\Omega} e^u \nabla \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} \left[\lambda^2 \right] \cdot \nabla \hat{w} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} \left[\lambda^2 \right] \hat{w} \, ds \\ = - \frac{1}{\sigma^2} \sum_{j=1}^K w^2(\mathbf{x}_j) \hat{w}(\mathbf{x}_j). \end{aligned} \quad (16)$$

Let us define

$$\tilde{\lambda}^2 := \mathbb{E}_{\pi_{\text{like}}(\mathbf{d}|u)} \left[\lambda^2 \right],$$

⁴ Note that the Fisher information operator is typically defined for finite dimensional settings in which it is a matrix.

then (16) becomes

$$\int_{\Omega} e^u \nabla \tilde{\lambda}^2 \cdot \nabla \hat{w} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, \tilde{\lambda}^2 \hat{w} \, ds = -\frac{1}{\sigma^2} \sum_{j=1}^K w^2(\mathbf{x}_j) \hat{w}(\mathbf{x}_j). \quad (17)$$

As a result, the Fisher information operator acting along directions \tilde{u} and u^2 reads

$$\langle \langle G(u), \tilde{u} \rangle, u^2 \rangle = \int_{\Omega} \tilde{u} e^u \nabla w \cdot \nabla \tilde{\lambda}^2 \, d\Omega, \quad (18)$$

where $\tilde{\lambda}^2$ is the solution of (17), a variant of the second order adjoint equation (13). The Fisher information operator therefore coincides with the Gauss–Newton Hessian of the cost functional (3).

The procedure for computing the gradient acting on an arbitrary direction is clear. One first solves the first order forward equation (8) for w , then the first order adjoint (10) for λ , and finally evaluate (9). Similarly, one can compute the Hessian (or the Fisher information operators) acting on two arbitrary directions by first solving the second order forward equation (12) for w^2 , then the second order adjoint equation (13) (or its variant (17)) for λ^2 (or λ^2), and finally evaluating (11) (or (18)).

One of the main goals of the paper is to study the RMHMC method in the context of Bayesian inverse problems governed by PDEs. It is therefore essential to compute the derivative of the Fisher information operator. This task is obvious for problems with available closed form expressions of the likelihood and the prior, but it is not so for those governed by PDEs. Nevertheless, using the adjoint technique we can compute the third order derivative tensor acting on three arbitrary directions with three extra PDE solves, as we now show. To that end, recall that the Fisher information operator acting on directions \tilde{u} and u^2 is given by (18). The Fréchet derivative of the Fisher information operator along the additional direction u^3 (superscript ‘3’ means the third variation direction) is given by

$$\begin{aligned} \langle \langle \langle T(u), \tilde{u} \rangle, u^2 \rangle, u^3 \rangle &:= \langle \nabla \langle \langle G(u), \tilde{u} \rangle, u^2 \rangle, u^3 \rangle \\ &= \int_{\Omega} \tilde{u} u^3 e^u \nabla w \cdot \nabla \tilde{\lambda}^2 \, d\Omega + \int_{\Omega} \tilde{u} e^u \nabla w^3 \cdot \nabla \tilde{\lambda}^2 \, d\Omega + \int_{\Omega} \tilde{u} e^u \nabla w \cdot \nabla \lambda^{2,3} \, d\Omega, \end{aligned} \quad (19)$$

where $w^3, \lambda^{2,3}$ are the variation of w and λ^2 in the direction u^3 , respectively. One can show that u^3 satisfies another second order forward equation

$$\int_{\Omega} e^u \nabla w^3 \cdot \nabla \hat{\lambda} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, w^3 \hat{\lambda} \, ds = - \int_{\Omega} u^3 e^u \nabla w \cdot \nabla \hat{\lambda} \, d\Omega. \quad (20)$$

Similarly, $\lambda^{2,3}$ is the solution of the third order adjoint equation

$$\int_{\Omega} e^u \nabla \lambda^{2,3} \cdot \nabla \hat{w} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, \lambda^{2,3} \hat{w} \, ds = \frac{1}{\sigma^2} \sum_{j=1}^K w^{2,3}(\mathbf{x}_j) \hat{w}(\mathbf{x}_j) - \int_{\Omega} u^3 e^u \nabla \tilde{\lambda}^2 \cdot \nabla \hat{w} \, d\Omega, \quad (21)$$

and $w^{2,3}$, the variation of u^2 in direction u^3 , satisfies the following third order forward equation

$$\begin{aligned} \int_{\Omega} e^u \nabla w^{2,3} \cdot \nabla \hat{\lambda} \, d\Omega + \int_{\partial\Omega \setminus \Gamma_R} Bi \, w^{2,3} \hat{\lambda} \, ds = \\ - \int_{\Omega} u^3 e^u \nabla w^2 \cdot \nabla \hat{\lambda} \, d\Omega - \int_{\Omega} u^3 u^2 e^u \nabla w \cdot \nabla \hat{\lambda} \, d\Omega - \int_{\Omega} u^2 e^u \nabla w^3 \cdot \nabla \hat{\lambda} \, d\Omega. \end{aligned} \quad (22)$$

Note that it would have required four extra PDE solves if one computes the third derivative of the full Hessian (11).

It is important to point out that the operator T is only symmetric with respect to \tilde{u} and u^2 since the Fisher information is symmetric, but not with respect to \tilde{u} and u^3 or u^2 and u^3 . The full symmetry only holds for the derivative of the full Hessian, that is, the true third derivative of the cost functional.

4. Discretization

As presented in section 2, we view our inverse problem from an infinite dimensional point of view. As such, to implement our approach on computers, we need to discretize the prior, the likelihood and hence the posterior. We choose to use the FEM. In particular, we employ the standard $H^1(\Omega)$ FEM to discretize the forwards and adjoints (the likelihood), and the operator \mathcal{A} (the prior). It should be pointed out that the Cameron–Martin space can be shown (see, e.g., [9]) to be a subspace of the usual fractional Sobolev space $H^s(\Omega)$, which is in turn a subspace of $H^1(\Omega)$. Thus, we are using a non-conforming FEM approach (outer approximation). For convenience, we further assume that the discretized state and parameter live on the same finite element mesh. Since FEM approximation of elliptic operators is standard (see, e.g., [13]), we will not discuss it here. Instead, we describe the MTT (see, e.g., [14] and the references therein) to discretize the prior.

Define $\mathcal{Q} := C^{1/2} = \mathcal{A}^{-s/2}$, then the eigenpairs (λ_i, v_i) of \mathcal{Q} define the KL expansion of the prior distribution as

$$u = u_0 + \frac{1}{\sqrt{\alpha}} \sum_{i=1}^{\infty} a_i \lambda_i v_i,$$

where $a_i \sim \mathcal{N}(0, 1)$. We need to solve

$$\mathcal{Q}v_i = \lambda_i v_i,$$

or equivalently

$$\mathcal{A}^{s/2}v_i = \frac{1}{\lambda_i}v_i. \quad (23)$$

To solve (23) using the MTT, let us denote by \mathbf{M} the mass matrix, and \mathbf{K} the stiffness matrix resulting from the discretization of the Laplacian Δ . The representation of \mathcal{A} in the finite element space (see, e.g., [15] and the references therein) is given by

$$\mathbf{A} := \mathbf{M}^{-1}\mathbf{K} + \mathbf{I}.$$

Let bold symbols denote the corresponding vector of FEM nodal values, e.g., \mathbf{u} is the vector containing all FEM nodal values of u . If we define (σ_i, \mathbf{v}_i) as eigenpairs for \mathbf{A} , i.e.

$$\mathbf{A}\mathbf{v}_i = \sigma_i \mathbf{v}_i, \quad \text{or } \mathbf{A}\mathbf{V} = \Sigma\mathbf{V}, \quad (24)$$

where $\mathbf{v}_i^T \mathbf{M} \mathbf{v}_j = \delta_{ij}$, and hence $\mathbf{V}^{-1} = \mathbf{V}^T \mathbf{M}$, δ_{ij} is the Kronecker delta function, and Σ is the diagonal matrix with entries σ_i . Since \mathbf{A} is similar to $\mathbf{M}^{-\frac{1}{2}}(\mathbf{K} + \mathbf{M})\mathbf{M}^{-\frac{1}{2}}$, a symmetric positive definite matrix, \mathbf{A} has positive eigenvalues. Using MTT method, the matrix representation of (23) reads

$$\mathbf{A}^{s/2}\mathbf{v}_i = \frac{1}{\lambda_i}\mathbf{v}_i,$$

where

$$\mathbf{A}^{s/2} := \mathbf{V}\Sigma^{s/2}\mathbf{V}^{-1}.$$

It follows that

$$\lambda_i = \sigma_i^{-s/2}.$$

The Galerkin FEM approximation of the prior via truncated KL expansion reads

$$\mathbf{u} = \mathbf{u}_0 + \frac{1}{\sqrt{\alpha}} \sum_{i=1}^N a_i \lambda_i \mathbf{v}_i, \quad (25)$$

with \mathbf{u} as the FEM nodal value of the approximate prior sample u and N as the number of FEM nodal points. Note that for ease in writing, we have used the same notation u for both infinite dimensional prior sample and its FEM approximation. Since $u \in L^2(\Omega)$, \mathbf{u} naturally lives in \mathbb{R}_M^N , the Euclidean space with weighted inner product $(\cdot, \cdot)_M := (\cdot, \mathbf{M}\cdot)$.

A question arises: What is the distribution of \mathbf{u} ? Clearly \mathbf{u} is a Gaussian with mean \mathbf{u}_0 since a_i are. The covariance matrix \mathbf{C} for \mathbf{u} is defined by

$$(\mathbf{z}, \mathbf{C}\mathbf{y})_M := \mathbb{E}[(\mathbf{u} - \mathbf{u}_0, \mathbf{M}\mathbf{z})(\mathbf{u} - \mathbf{u}_0, \mathbf{M}\mathbf{y})] = \frac{1}{\alpha} \mathbf{z}^T \mathbf{M} \mathbf{V} \Lambda^2 \mathbf{V}^T \mathbf{M} \mathbf{y},$$

where we have used (25) to obtain the second equality and Λ is the diagonal matrix with entries $\Lambda_{ii} = \lambda_i^{-1}$. It follows that

$$\mathbf{C} = \frac{1}{\alpha} \mathbf{V} \Lambda^2 \mathbf{V}^T \mathbf{M}, \quad (26)$$

as a map from \mathbb{R}_M^N to \mathbb{R}_M^N , and its inverse can be shown to be

$$\mathbf{C}^{-1} = \alpha \mathbf{V} \Lambda^{-2} \mathbf{V}^T \mathbf{M},$$

whence the distribution of \mathbf{u} is

$$\mathbf{u} \sim \mathcal{N}(\mathbf{u}_0, \alpha \mathbf{V} \Lambda^{-2} \mathbf{V}^T \mathbf{M}) \sim \exp \left[-\frac{\alpha}{2} (\mathbf{u} - \mathbf{u}_0)^T \mathbf{M} \mathbf{V} \Lambda^{-2} \mathbf{V}^T \mathbf{M} (\mathbf{u} - \mathbf{u}_0) \right]. \quad (27)$$

As a result, the FEM discretization of the prior can be written as

$$\frac{\alpha}{2} \|u - u_0\|_C^2 := \frac{\alpha}{2} \|\mathcal{A}^{s/2}(u - u_0)\|^2 \stackrel{\text{MTT}}{\approx} \frac{\alpha}{2} (\mathbf{u} - \mathbf{u}_0)^T \mathbf{M} \mathbf{V} \Lambda^{-2} \mathbf{V}^T \mathbf{M} (\mathbf{u} - \mathbf{u}_0).$$

Thus, the FEM approximation of the posterior is given by

$$\pi(\mathbf{u}|\mathbf{d}) \sim \exp \left(-\frac{1}{2} |\mathbf{d} - \mathcal{G}(u)|_{\mathbb{L}}^2 \right) \times \exp \left(-\frac{\alpha}{2} (\mathbf{u} - \mathbf{u}_0)^T \mathbf{M} \mathbf{V} \Lambda^{-2} \mathbf{V}^T \mathbf{M} (\mathbf{u} - \mathbf{u}_0) \right).$$

The detailed derivation of the FEM approximation of infinite Bayesian inverse problems in general and the prior in particular will be presented elsewhere [16].

5. Riemannian manifold Langevin and Hamiltonian Monte Carlo methods

In this section we give a brief overview of the MCMC algorithms that we consider in this work. Some familiarity with the concepts of MCMC is required by the reader since an introduction to the subject is out of the scope of this paper.

5.1. Metropolis–Hastings

For a random vector $\mathbf{u} \in \mathbb{R}^N$ with density $\pi(\mathbf{u})$ the Metropolis–Hastings algorithm employs a proposal mechanism $q(\mathbf{u}^*|\mathbf{u}^{t-1})$ and proposed moves are accepted with probability $\min\{1, \pi(\mathbf{u}^*)q(\mathbf{u}^{t-1}|\mathbf{u}^*)/\pi(\mathbf{u}^{t-1})q(\mathbf{u}^*|\mathbf{u}^{t-1})\}$. Tuning the Metropolis–Hastings algorithm involves selecting an appropriate proposal mechanism. A common choice is to use a Gaussian proposal of the form $q(\mathbf{u}^*|\mathbf{u}^{t-1}) = \mathcal{N}(\mathbf{u}^*|\mathbf{u}^{t-1}, \Sigma)$, where $\mathcal{N}(\cdot|\mu, \Sigma)$ denotes the multivariate normal density with mean μ and covariance matrix Σ .

Selecting the covariance matrix however, is far from trivial in most of cases since knowledge about the target density is required. Therefore a more simplified proposal mechanism is often considered where the covariance matrix is replaced with a diagonal matrix such as $\Sigma = \epsilon\mathbf{I}$ where the value of the scale parameter ϵ has to be tuned in order to achieve fast convergence and good mixing. Small values of ϵ imply small transitions and result in high acceptance rates while the mixing of the Markov Chain is poor. Large values on the other hand, allow for large transitions but they result in most of the samples being rejected. Tuning the scale parameter becomes even more difficult in problems where the standard deviations of the marginal posteriors differ substantially, since different scales are required for each dimension, and when correlations between different variables exist. In the case of PDE-constrained inverse problems in very high dimensions with strong nonlinear interactions inducing complex non-convex structures in the target posterior this tuning procedure is typically doomed to failure of convergence and mixing.

There have been many subsequent developments of this basic algorithm however the most important with regard to inverse problems, arguably, is the formal definition of Metropolis Hastings in an infinite dimensional functional space. One of the main failings of Metropolis Hastings is the drop-off in acceptance probability as the dimension of the problem increases. By defining the Metropolis acceptance probability in the appropriate Hilbert space the acceptance probability should then be invariant to the dimension of the problem and this is indeed the case as is described in a number of scenarios by [17]. Furthermore the definition of a Markov chain transition kernel directly in the Hilbert space which exploits Hamiltonian dynamics in the proposal mechanism followed in [18].

These are important methodological advances for MCMC applied to Inverse Problems. As the infinite dimensional nature of the problem is a fundamental aspect of the problem it is sensible that this characteristic is embedded in the MCMC scheme. In a similar vein by noting that the statistical model associated with the specific inverse problem is generated from an underlying PDE or system of ordinary differential equations a natural geometric structure on the space of probability distributions is induced. This structure provides a rich source of model specific information that can be exploited in devising MCMC schemes that are informed by the underlying structure of the model itself.

In [19] a way around this situation was provided by accepting that the statistical model can itself be considered as an object with an underlying geometric structure that could be embedded into the proposal mechanism. A class of MCMC methods were developed based on the differential geometric concepts underlying Riemannian manifolds.

5.2. Riemann manifold Metropolis adjusted Langevin algorithm (RMMLA)

Denoting the log of the target density as $\mathcal{L}(\mathbf{u}) = \log \pi(\mathbf{u})$, the manifold Metropolis Adjusted Langevin Algorithm method [19], defines a Langevin diffusion with stationary distribution $\pi(\mathbf{u})$ on the Riemann manifold of density functions with metric tensor $\mathbf{G}(\mathbf{u})$. By employing a first order Euler integrator for discretising the stochastic differential equation a proposal

mechanism with density $q(\mathbf{u}^*|\mathbf{u}^{t-1}) = \mathcal{N}(\mathbf{u}^*|\mu(\mathbf{u}^{t-1}, \epsilon), \epsilon^2 \mathbf{G}^{-1}(\mathbf{u}^{t-1}))$ is defined, where ϵ is the integration step size, a parameter which needs to be tuned, and the k th component of the mean function $\mu(\mathbf{u}, \epsilon)_k$ is

$$\mu(\mathbf{u}, \epsilon)_k = \mathbf{u}_k + \frac{\epsilon^2}{2} \left(\mathbf{G}^{-1}(\mathbf{u}) \nabla_{\mathbf{u}} \mathcal{L}(\mathbf{u}) \right)_k - \epsilon^2 \sum_{i=1}^N \sum_{j=1}^N \mathbf{G}(\mathbf{u})_{i,j}^{-1} \Gamma_{i,j}^k, \quad (28)$$

where $\Gamma_{i,j}^k$ are the Christoffel symbols of the metric in local coordinates. Note that we have used the Christoffel symbols to express the derivatives of the metric tensor, and they are computed using the adjoint method presented in section 3.

Due to the discretisation error introduced by the first order approximation convergence to the stationary distribution is not guaranteed anymore and thus the Metropolis–Hastings ratio is employed to correct for this bias. In [19] a number of examples are provided illustrating the potential of such a scheme for challenging inference problems.

One can interpret the proposal mechanism of RMMALA as a local Gaussian approximation to the target density where the effective covariance matrix in RMMALA is the inverse of the metric tensor evaluated at the current position. Furthermore a simplified version of the RMMALA algorithm, termed sRMMALA, can also be derived by assuming a manifold with constant curvature thus cancelling the last term in equation (28) which depends on the Christoffel symbols. Whilst this is a step forward in that much information about the target density is now embedded in the proposal mechanism it is still driven by a random walk. The next approach to be taken goes beyond the direct and scaled random walk by defining proposals which follow the geodesic flows on the manifold of densities and thus presents a potentially really powerful scheme to explore posterior distributions.

5.3. RMHMC

The RMHMC method defines a Hamiltonian on the Riemann manifold of probability density functions by introducing the auxiliary variables $\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbf{G}(\mathbf{u}))$ which are interpreted as the momentum at a particular position \mathbf{u} and by considering the negative log of the target density as a potential function. More formally the Hamiltonian defined on the Riemann manifold is

$$\mathcal{H}(\mathbf{u}, \mathbf{p}) = -\mathcal{L}(\mathbf{u}) + \frac{1}{2} \log(2\pi|\mathbf{G}(\mathbf{u})|) + \frac{1}{2} \mathbf{p}^T \mathbf{G}(\mathbf{u})^{-1} \mathbf{p}, \quad (29)$$

where the terms $-\mathcal{L}(\mathbf{u}) + \frac{1}{2} \log(2\pi|\mathbf{G}(\mathbf{u})|)$ and $\frac{1}{2} \mathbf{p}^T \mathbf{G}(\mathbf{u})^{-1} \mathbf{p}$ are the potential energy and kinetic energy terms respectively. and the dynamics given by Hamiltons equations are

$$\frac{d\mathbf{u}_k}{dt} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}_k} = \left(\mathbf{G}(\mathbf{u})^{-1} \mathbf{p} \right)_k, \quad (30)$$

$$\frac{d\mathbf{p}_k}{dt} = -\frac{\partial \mathcal{H}}{\partial \mathbf{u}_k} = \frac{\partial \mathcal{L}(\mathbf{u})}{\partial \mathbf{u}_k} - \frac{1}{2} \text{Tr} \left[\mathbf{G}(\mathbf{u})^{-1} \frac{\partial \mathbf{G}(\mathbf{u})}{\partial \mathbf{u}_k} \right] + \frac{1}{2} \mathbf{p}^T \mathbf{G}(\mathbf{u})^{-1} \frac{\partial \mathbf{G}(\mathbf{u})}{\partial \mathbf{u}_k} \mathbf{G}(\mathbf{u})^{-1} \mathbf{p}. \quad (31)$$

These dynamics define geodesic flows at a particular energy level and as such make proposals which follow deterministically the most efficient path across the manifold from the current density to the proposed one. Simulating the Hamiltonian requires a time-reversible and volume preserving numerical integrator. For this purpose the generalized leapfrog algorithm can be employed and provides a deterministic proposal mechanism for simulating from the conditional distribution, i.e. $\mathbf{u}^*|\mathbf{p} \sim \pi(\mathbf{u}^*|\mathbf{p})$. More details about the generalized leapfrog integrator can be found in [19]. To simulate a path (which turns out to be a local geodesic) across the manifold, the Leapfrog integrator is iterated L times which along with the

integration step size ϵ are parameters requiring tuning. Again due to the discrete integration errors on simulating the Hamiltonian in order to ensure convergence to the stationary distribution the Metropolis–Hastings acceptance ratio is applied.

The RMHMC method has been shown to be highly effective in sampling from posteriors induced by complex statistical models and offers the means to efficiently explore the hugely complex and high dimensional posteriors associated with PDE-constrained inverse problems.

6. Low rank approximation of the Fisher information matrix

As presented in section 5, we use the Fisher information matrix at the MAP point augmented with the Hessian of the prior as the metric tensor in our HMC simulations. It is therefore necessary to compute the augmented Fisher matrix and its inverse. In [1–3], we have shown that the Gauss–Newton Hessian of the cost functional (3), also known as the data misfit, is a compact operator, and that for smooth u its eigenvalues decay exponentially to zero. Thus, the range space of the Gauss–Newton Hessian is effectively finite-dimensional even before discretization, i.e., it is independent of the mesh. In other words, the Fisher information matrix admits accurate low rank approximations and the accuracy can be improved as desired by simply increasing the rank of the approximation. We shall exploit this fact to compute the augmented Fisher information matrix and its inverse efficiently. We start with the augmented Fisher information matrix in \mathbb{R}_M^N

$$\mathbf{G} := \mathbf{M}^{-1}\mathbf{H} + \alpha\mathbf{V}\Lambda^{-2}\mathbf{V}^T\mathbf{M} = \alpha\mathbf{V}\Lambda^{-1}\left(\frac{1}{\alpha}\Lambda\mathbf{V}^T\mathbf{H}\mathbf{V}\Lambda + \mathbf{I}\right)\Lambda^{-1}\mathbf{V}^{-1},$$

where \mathbf{H} is the Fisher information matrix obtained from (18) by taking \tilde{u} and u^2 as FEM basis functions.

Assume that \mathbf{H} is compact (see, e.g., [1, 2]), together with the fact that Λ_{ii} decays to zero, we conclude that the prior-preconditioned Fisher information matrix

$$\tilde{\mathbf{H}} := \frac{1}{\alpha}\Lambda\mathbf{V}^T\mathbf{H}\mathbf{V}\Lambda,$$

also has eigenvalues decaying to zero. Therefore it is expected that the eigenvalues of the prior-preconditioned matrix decays faster than those of the original matrix \mathbf{H} . Indeed, the numerical results in section 7 will confirm this observation. It follows that (see, e.g., [20, 21] for similar decomposition) $\tilde{\mathbf{H}}$ admits a r -rank approximation of the form

$$\tilde{\mathbf{H}} = \frac{1}{\alpha}\Lambda\mathbf{V}^T\mathbf{H}\mathbf{V}\Lambda \approx \mathbf{V}_r\mathbf{S}\mathbf{V}_r^T,$$

where \mathbf{V}_r and \mathbf{S} (diagonal matrix) contain the first r dominant eigenvectors and eigenvalues of $\tilde{\mathbf{H}}$, respectively. In this work, similar to [21], we use the one-pass randomized algorithm in [22] to compute the low rank approximation. Consequently, the augmented Fisher information matrix becomes

$$\mathbf{G} \approx \alpha\mathbf{V}\Lambda^{-1}\left(\mathbf{V}_r\mathbf{S}\mathbf{V}_r^T + \mathbf{I}\right)\Lambda^{-1}\mathbf{V}^{-1},$$

from which we obtain the inverse, by using the Woodbury formula [23]

$$\mathbf{G}^{-1} \approx \frac{1}{\alpha}\mathbf{V}\Lambda\left(\mathbf{I} - \mathbf{V}_r\mathbf{D}\mathbf{V}_r^T\right)\Lambda\mathbf{V}^{-1},$$

where \mathbf{D} is a diagonal matrix with $\mathbf{D}_{ii} = \mathbf{S}_{ii}/(\mathbf{S}_{ii} + 1)$.

In the RMHMC method, we need to randomly draw the momentum variable as $\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbf{G})$. If one considers

$$\mathbf{p} = \sqrt{\alpha} \mathbf{V} \Lambda^{-1} \mathbf{b} + \sqrt{\alpha} \mathbf{V} \Lambda^{-1} \mathbf{V}_r \mathbf{S}^{1/2} \mathbf{c},$$

where $\mathbf{b}_i, \mathbf{c}_i \sim \mathcal{N}(0, 1)$, then one can show, by inspection, that \mathbf{p} is distributed by $\mathcal{N}(\mathbf{0}, \mathbf{G})$.

7. Numerical results

For convenience, let us recall that the FEM approximation of the posterior is given as

$$\pi(\mathbf{u}|\mathbf{d}) \sim \exp\left(-\frac{1}{2} \|\mathbf{d} - \mathcal{G}(u)\|_{\mathbf{L}}^2\right) \times \exp\left(-\frac{\alpha}{2} (\mathbf{u} - \mathbf{u}_0)^T \mathbf{M} \mathbf{V} \Lambda^2 \mathbf{V}^T \mathbf{M} (\mathbf{u} - \mathbf{u}_0)\right), \quad (32)$$

\mathbf{u}_0 is the FEM nodal value of the prior mean function u_0 , \mathbf{M} is the mass matrix, \mathbf{V} the matrix of eigenvectors defined in (24), Λ the diagonal matrix introduced in (26), $\mathbf{L} = \sigma^2 \mathbf{I}$, \mathbf{d} vector of observation data, and $\mathcal{G}(u)$ the forward map given by the forward equation

$$\begin{aligned} -\nabla \cdot (e^u \nabla w) &= 0 && \text{in } \Omega, \\ -e^u \nabla w \cdot \mathbf{n} &= Bi u && \text{on } \partial\Omega \setminus \Gamma_R, \\ -e^u \nabla w \cdot \mathbf{n} &= -1 && \text{on } \Gamma_R, \end{aligned}$$

that is discretized by the H^1 -conforming FEM method.

In this section, we study Riemann manifold Monte Carlo methods and their variations to explore the posterior (32). In particular, we compare the performance of four methods: (i) sRMMALA obtained by ignoring the third derivative in RMMALA, (ii) RMMALA, (iii) sRMHMC obtained by first computing the augmented Fisher metric tensor at the MAP point and then using it as the constant metric tensor, (iv) RMHMC. For all methods, we form the augmented Fisher information matrix exactly using (18) with \tilde{u} and u^2 as finite element basis vectors. For RMMALA and RMHMC we also need the derivative of the metric tensor which is a third order tensor. It can be constructed exactly using (19) with \tilde{u} , u^2 and u^3 as finite element basis vectors. We also need extra work for the RMHMC method since each Stormer–Verlet step requires an implicit solve for both the first half of momentum and full position. For inverse problems such as those considered in this paper, the fixed point approach proposed in [19] does not seem to converge. We therefore have to resort to a full Newton method. Since we explicitly construct the metric tensor and its derivative, it is straightforward for us to develop the Newton scheme. For all problems considered in this section, we have observed that it takes at most five Newton iterations to converge.

Note that we limit ourselves in comparing these four methods in the Riemannian manifold MCMC sampling family. Clearly, other methods are available, we avoid ‘unmatched comparison’ in terms of cost and the level of exploiting the structure of the problem. Even in this limited family, RMHMC is most expensive since it requires not only third derivatives but also implicit solves, but the ability in generating almost independent samples is attractive and appealing as we shall show.

Though our proposed approach described in previous sections are valid for any spatial dimension d , we restrict ourselves to a one dimensional problem, i.e. $d = 1$, to clearly illustrate our points and findings. In particular, we take $\Omega = [0, 1]$, $\Gamma_R = \{1\}$. We set $Bi = 0.1$ for all examples. As discussed in section 2, for the Gaussian prior to be well-defined, we take $s = 0.6 > n/2 = 1/2$.

7.1. Two-parameter examples

We start our numerical experiments with two parameters. This will help demonstrate various aspects of RMHMC which are otherwise too computationally expensive for high dimensional problems. In particular, two-parameter example allows us to compute the complete third derivative tensor and perform the Newton method for each Stormer–Verlet step. This in turn allows us to show the capability of the full RMHMC over its simplified variants in tackling challenging posterior densities in which the metric tensor changes rapidly.

In order to construct the case with two parameters we consider FEM with one finite element. We assume that there is one observation point, i.e. $K = 1$, and it is placed at the left boundary $x = 0$. In the first example, we first take $s = 0.6$, $\sigma = 0.1$, and $\alpha = 0.1$. The posterior in this case is shown in figure 1(a). We start by taking a time step of 0.02 with 100 Stormer–Verlet steps for both sRMHMC and RMHMC. The acceptance rate for both methods is 1. One would take a time step of 2 for both sRMMALA and RMMALA to be comparable with sRMHMC and RMHMC, but the acceptance would be zero. Instead we take time step of 1 so that the acceptance rate is about 0.5 for sRMMALA and 0.3 for RMMALA. The MAP point is chosen as the initial state for all the chains with 5000 sample excluding the first 100 burn-ins. The result is shown in figure 2.

As can be seen, the RMHMC chain is the best in terms of mixing by comparing the second column (the trace plot) and the third column (the autocorrelation function (ACF)). Each RMHMC sample is almost uncorrelated to the previous ones. The sRMHMC is the second best, but the samples are strongly correlated compared to those of RMHMC, e.g. one uncorrelated sample for every 40. It is interesting to observe that the full RMMALA and sRMMALA have performance in terms of auto-correlation length that is qualitatively similar at least in the first 5000 samples. This is due to the RMMALA schemes being driven by a single step random walk that cannot exploit fully the curvature information available to the geodesic flows of RMHMC, see rejoinder of [19].

Note that it is not our goal to compare the behavior of the chains when they converge. Rather we would like to qualitatively study how fast the chains are well-mixed (mixing time). This is important for large-scale problems governed by PDEs since ‘unpredicted’ mixing time implies a lot of costly waste in PDE solves which one must avoid. Though RMHMC is expensive in generating a sample, the cost of generating an uncorrelated/independent sample seems to be comparable to sRMHMC for this example. In fact, if we measure the cost in terms of the number of PDE solves, the total number of PDE solves for RMHMC is 42 476 480 while it is 1 020 002 for sRMHMC, a factor of 40 more expensive. However, the cost in generating an almost uncorrelated/independent sample is the same since sRMHMC generates one uncorrelated sample out of 40 while it is one out of one for RMHMC.

To see how each method distributes the samples we plot one for every five samples in figure 3. All methods seem to explore the high probability density region very well. This explains why the sample mean and the 95% credibility region are similar for all methods in the first column of figure 2.

In the second example we consider the combination $s = 0.6$, $\sigma = 0.01$, and $\alpha = 1$ which leads to the posterior shown in figure 1(b). For sRMHMC and RMHMC, we take time step $\varepsilon = 0.04$ with $L = 100$ time steps, while it is 1 for both sRMMALA and RMMALA. Again, the acceptance rate is unity for both sRMHMC and RMHMC while it is 0.65 for sRMMALA and 0.55 for RMMALA, respectively. The result for four methods is shown in figure 4. As can be seen, this example seems to be easier than the first one since even though the time step is larger, the trace plot and the ACF looks better. It is interesting to observe that sRMHMC is comparable with RMHMC (in fact the ACF seems to be a bit better) for this example. As a

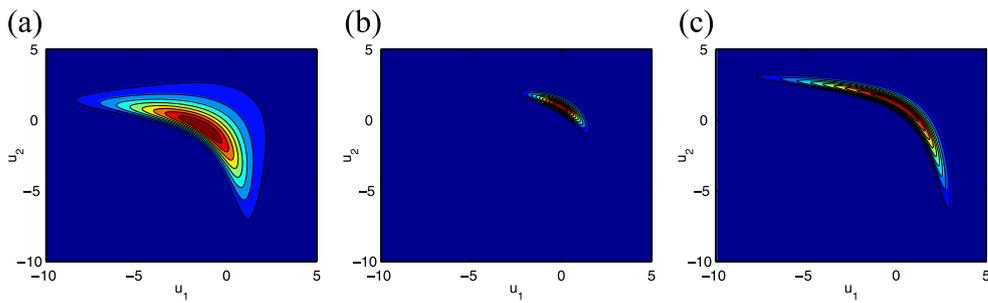


Figure 1. The contours of the posterior for three combinations of s (the prior smoothness), α (the ‘amount’ of the prior) and σ (the noise standard deviation). (a) $s = 0.6$, $\alpha = 0.1$ and $\sigma = 0.1$, (b) $s = 0.6$, $\alpha = 1$ and $\sigma = 0.01$ (c) $s = 0.6$, $\alpha = 0.1$ and $\sigma = 0.01$.

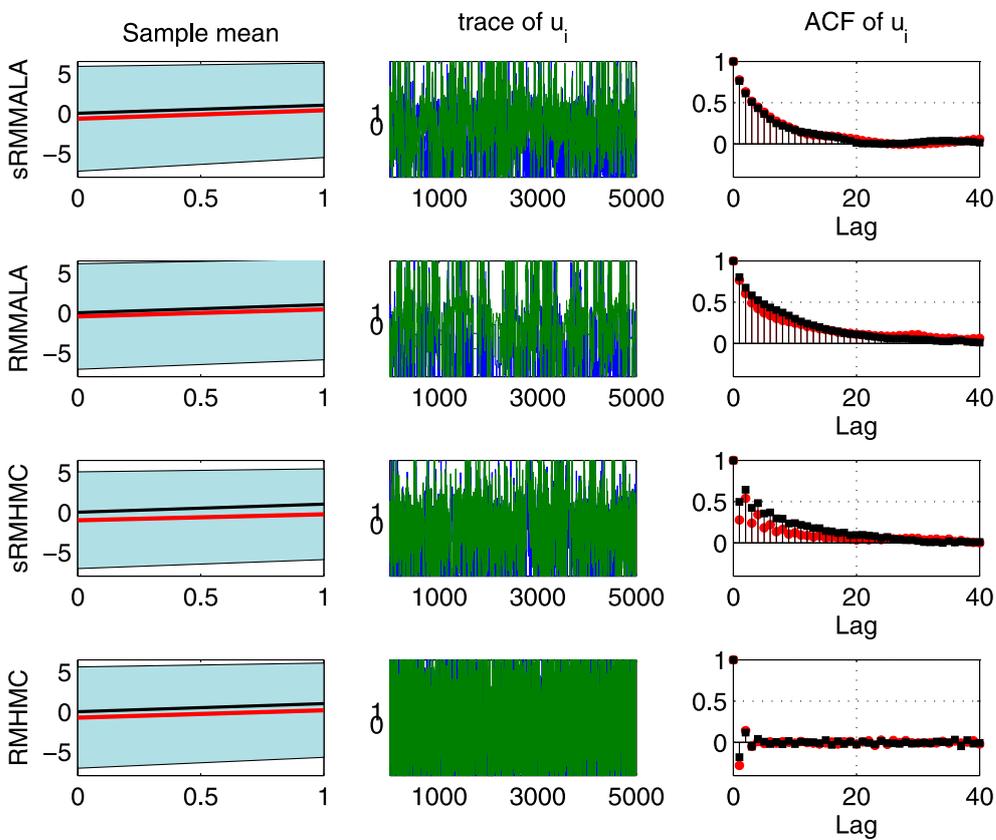


Figure 2. Comparison of simRMMALA, RMMALA, simRMHMC and RMHMC: chains with 5000 samples, burn-in of 100, starting at the MAP point. In this example, $s = 0.6$, $\alpha = 0.1$ and $\sigma = 0.1$. Time step is $\varepsilon = 1$ for simRMMALA and RMMALA, and $\varepsilon = 0.02$ with the number of time steps $L = 100$ for simRMHMC and RMHMC. In the left column: the exact synthetic solution is black, the sample mean is red, and the shaded region is the 95% credibility region. In the middle column: blue is the trace plot for u_1 while green is for u_2 . In the right column: red and black are the autocorrelation functions for u_1 and u_2 , respectively.

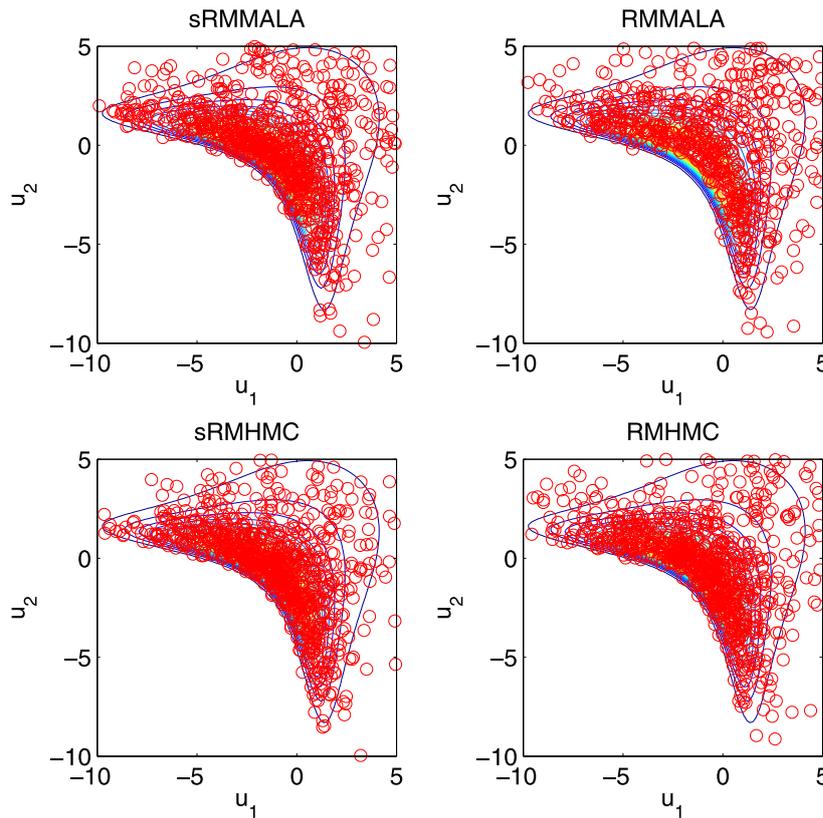


Figure 3. Comparison MCMC trajectories (plot one for each five samples) among simRMMALA, RMMALA, simRMHMC and RMHMC: chains with 5000 samples, burn-in of 100, starting at the MAP point. In this example, $s = 0.6$, $\alpha = 0.1$ and $\sigma = 0.1$. Time step is $\varepsilon = 1$ for simRMMALA and RMMALA, and $\varepsilon = 0.02$ with the number of time steps $L = 100$ for simRMHMC and RMHMC.

result, RMHMC is more expensive than sRMHMC for less challenging posterior in figure 1(b). Here, by less challenging we mean that the posterior is quite well approximated by a Gaussian at the MAP point, e.g. the metric tensor is almost constant. This is true for the posterior in figure 1(b) in which the Gaussian prior contribution is significant, i.e., $\alpha = 1$ instead of $\alpha = 0.1$. Conversely, the posterior is challenging if the metric tensor changes rapidly. Similar to the first example, one also see that the sample mean and the 95% credibility region are almost the same for all methods.

In the third example we consider the combination $s = 0.6$, $\sigma = 0.01$, and $\alpha = 0.1$ which leads to a skinny posterior with a long ridge as shown in figure 1(c). For sRMHMC and RMHMC, we take time step $\varepsilon = 0.02$ with $L = 100$ time steps, while it is 1 for both sRMMALA and RMMALA. Again, the acceptance rate is unity for both sRMHMC and RMHMC while it is 0.45 for sRMMALA and RMMALA. The result for four methods is shown in figure 5. For this example, the RMHMC is more desirable than sRMHMC since the cost to generate an uncorrelated/independent sample is smaller for the former than the latter. The reason is that the total number of PDEs solves for the former is 40 times more than the latter, but one out of very sixty samples is uncorrelated/independent.

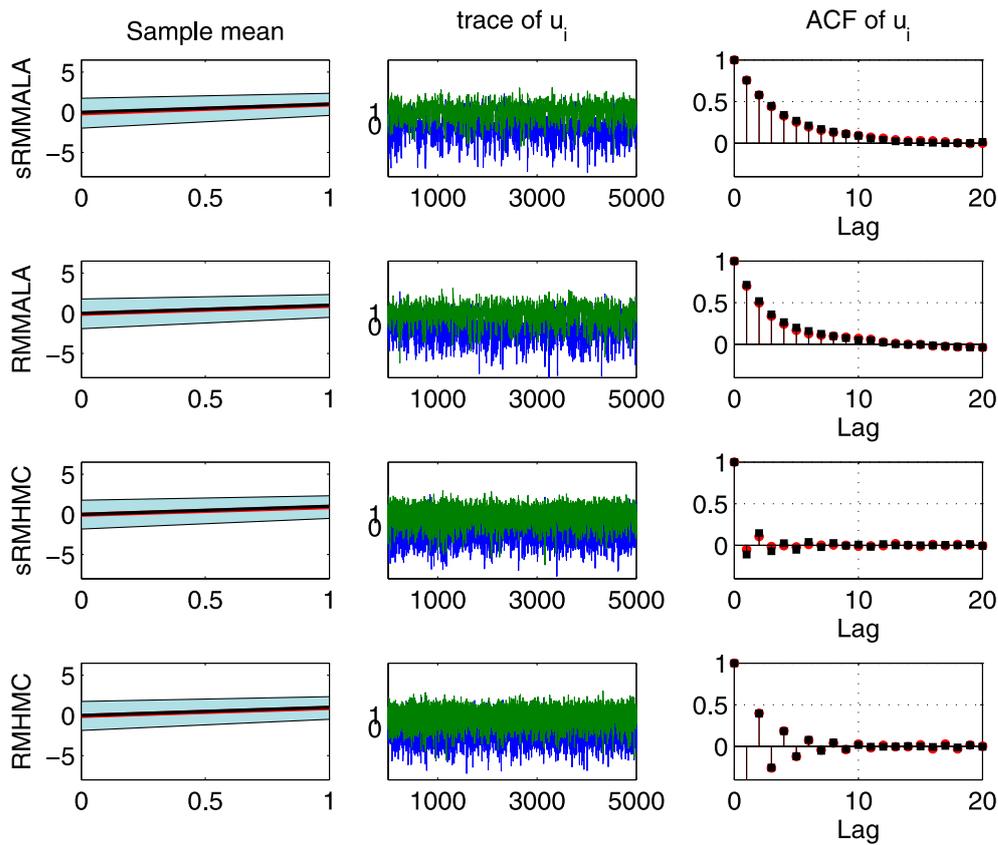


Figure 4. Comparison among simRMMALA, RMMALA, simRMHMC and RMHMC: chains with 5000 samples, burn-in of 100, starting at the MAP point. In this example, $s = 0.6$, $\alpha = 1$ and $\sigma = 0.01$. Time step is $\epsilon = 1$ for simRMMALA and RMMALA, and $\epsilon = 0.04$ with the number of time steps $L = 100$ for simRMHMC and RMHMC. In the left column: the exact synthetic solution is black, the sample mean is red and the shaded region is the 95% credibility region. In the middle column: blue is the trace plot for u_1 while green is for u_2 . In the right column: red and black are the autocorrelation functions for u_1 and u_2 , respectively.

7.2. Multi-parameter examples

In this section we choose to discretize $\Omega = [0, 1]$ with $2^{10} = 1024$ elements, and hence the number of parameters is 1025. For all simulations in this section, we choose $s = 0.6$, $\alpha = 10$, and $\sigma = 0.01$. For synthetic observations, we take $K = 64$ observations at $x_j = (j - 1)/2^6$, $j = 1, \dots, K$. Clearly, using the full blown RMHMC is out of the question since it is too expensive to construct the third derivative tensor and Newton method for each Stomer–Verlet step. For that reason, the sRMHMC becomes the viable choice. As studied in section 7.1, though sRMHMC loses the ability to efficiently sample from highly nonlinear posterior surfaces compared to the full RMHMC it is much less expensive to generate a sample since it does not require the derivative of the Fisher information matrix. In fact sRMHMC requires to (approximately) compute the Fisher information at the MAP point and then uses it as the fixed constant metric tensor throughout all leap-frog steps for all samples. Clearly, the gradient (9)

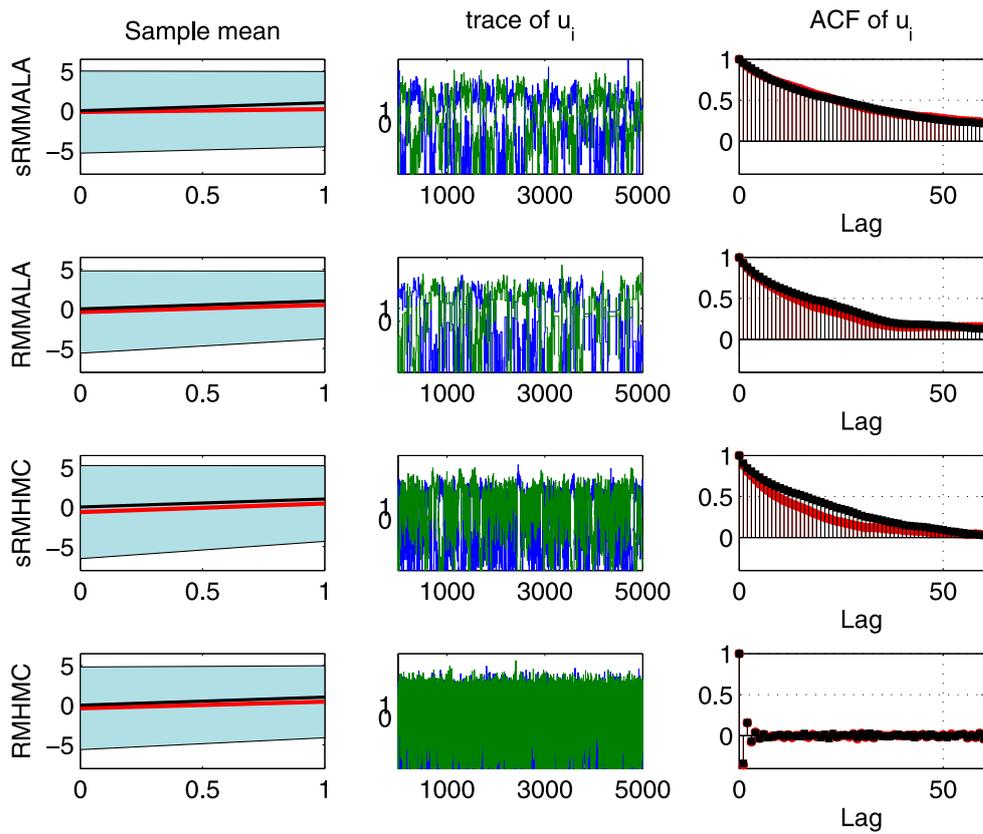


Figure 5. Comparison among simRMMALA, RMMALA, simRMHMC and RMHMC: chains with 5000 samples, burn-in of 100, starting at the MAP point. In this example, $s = 0.6$, $\alpha = 0.1$ and $\sigma = 0.01$. Time step is $\varepsilon = 0.7$ for simRMMALA and RMMALA, and $\varepsilon = 0.02$ with the number of time steps $L = 100$ for simRMHMC and RMHMC. In the left column: the exact synthetic solution is black, the sample mean is red and the shaded region is the 95% credibility region. In the middle column: blue is the trace plot for u_1 while green is for u_2 . In the right column: red and black are the autocorrelation functions for u_1 and u_2 , respectively.

has to be evaluated at each leap-frog step, but it can be computed efficiently using the adjoint method presented in section 3.

Nevertheless, constructing the exact Fisher information matrix requires 2×1025 PDEs solves. This is impractical if the dimension of the finite element space increases, e.g. by refining the mesh. Alternatively, due to the compactness of the Hessian of the prior-preconditioned misfit as discussed in section 6, we can use the randomized singular value decomposition (RSVD) technique [22] to compute its low rank approximations. Shown in figure 6 are the first 35 dominant eigenvalues of the Fisher information matrix and its prior-preconditioned counterpart. We also plot 20 approximate eigenvalues of the prior-preconditioned Fisher information matrix obtained from the RSVD method. As can be seen, the eigen spectrum of the prior-preconditioned Fisher information matrix decays faster than that of the original one. This is not surprising since the prior-preconditioned Fisher operator is a composition of the prior covariance, a compact operator, and the Fisher information operator,

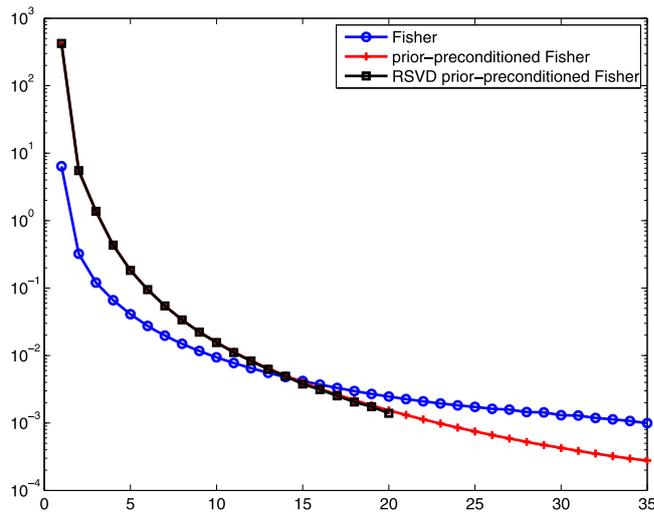


Figure 6. The eigen spectrum of the Fisher information matrix, the prior-preconditioned Fisher matrix, and the first 20 eigenvalues approximated using RSVD. Here, $s = 0.6$, $\alpha = 10$ and $\sigma = 0.01$.

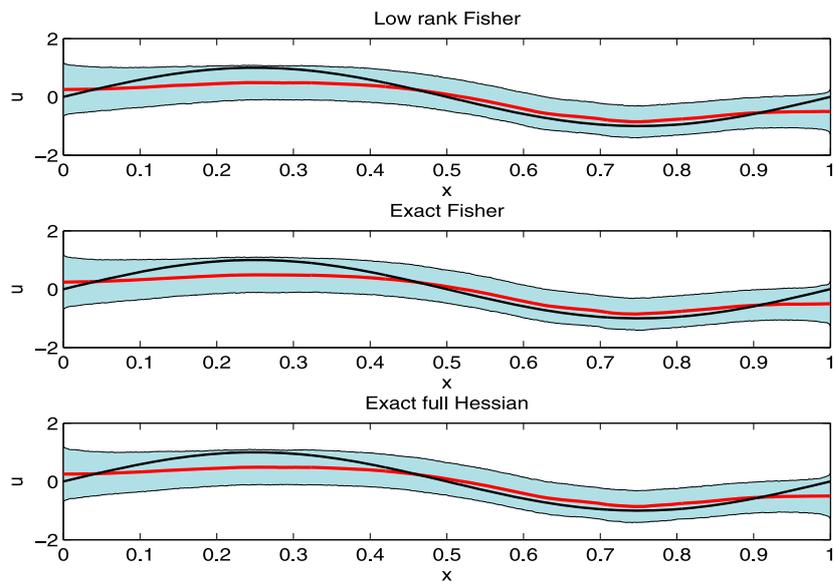


Figure 7. MCMC results of three sRMHMC methods with (i) the low rank Gauss–Newton Hessian, (ii) the exact Gauss–Newton Hessian and (iii) the full Hessian. In the figure are the empirical mean (red line), the exact distributed parameter used to generate the observation (black line) and 95% credibility (shaded region).

also a compact operator. The power of the RSVD is clearly demonstrated as the RSVD result for the first 20 eigenvalues is very accurate.

Next, we perform the sRMHMC method using three different constant metric tensors: (i) the low rank Gauss–Newton Hessian, (ii) the exact Gauss–Newton Hessian, and (iii) the full

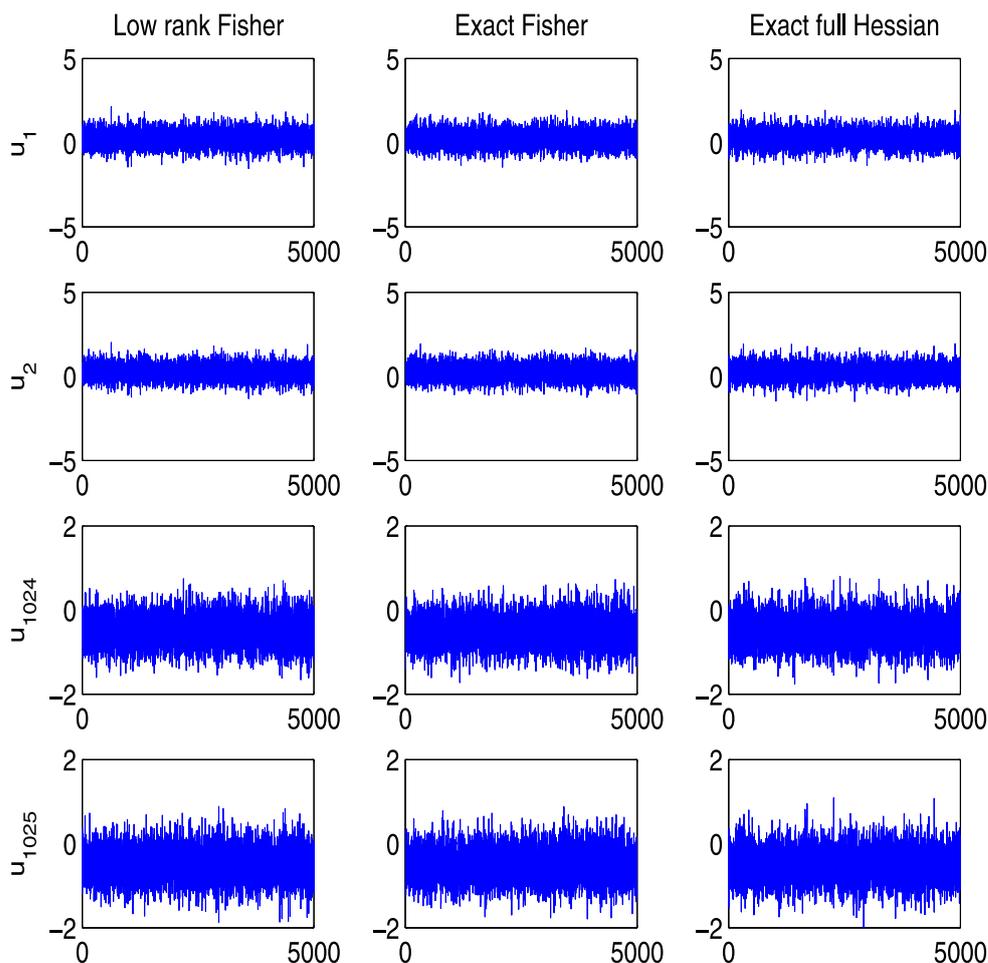


Figure 8. MCMC results of three sRMHMC methods with (i) the low rank Gauss–Newton Hessian (left column), (ii) the exact Gauss–Newton Hessian (middle column) and (iii) the full Hessian at the MAP point (right column). In the figure are the trace plot of the first two (1 and 2) and last two (1024 and 1025) parameters.

Hessian. For each case, we start the Markov chain at the MAP point and compute 5100 samples, the first 100 of which is then discarded as burn-ins. The empirical mean (red line), the exact distributed parameter used to generate the observation (black line), and 95% credibility region are shown in figure 7. As can be seen, the results from the three methods are indistinguishable. The first sRMHMC is the most appealing since it requires $2 \times 20 = 40$ PDE solves to construct the low rank Fisher information while the others need 2×1025 PDE solves. For large-scale problems with computationally expensive PDE solves, the first approach is the method of choice.

To further compare the three methods we record the trace plot of the first two (1 and 2) and the last two (1024 and 1025) parameters in figure 8. As can be observed, the chains from the three methods seem to be well-mixed and it is hard to see the difference among them. We also plot the ACF for these four parameters in figure 9. Again, results for the three sRMHMC methods are almost identical, namely, they generate almost uncorrelated samples. We

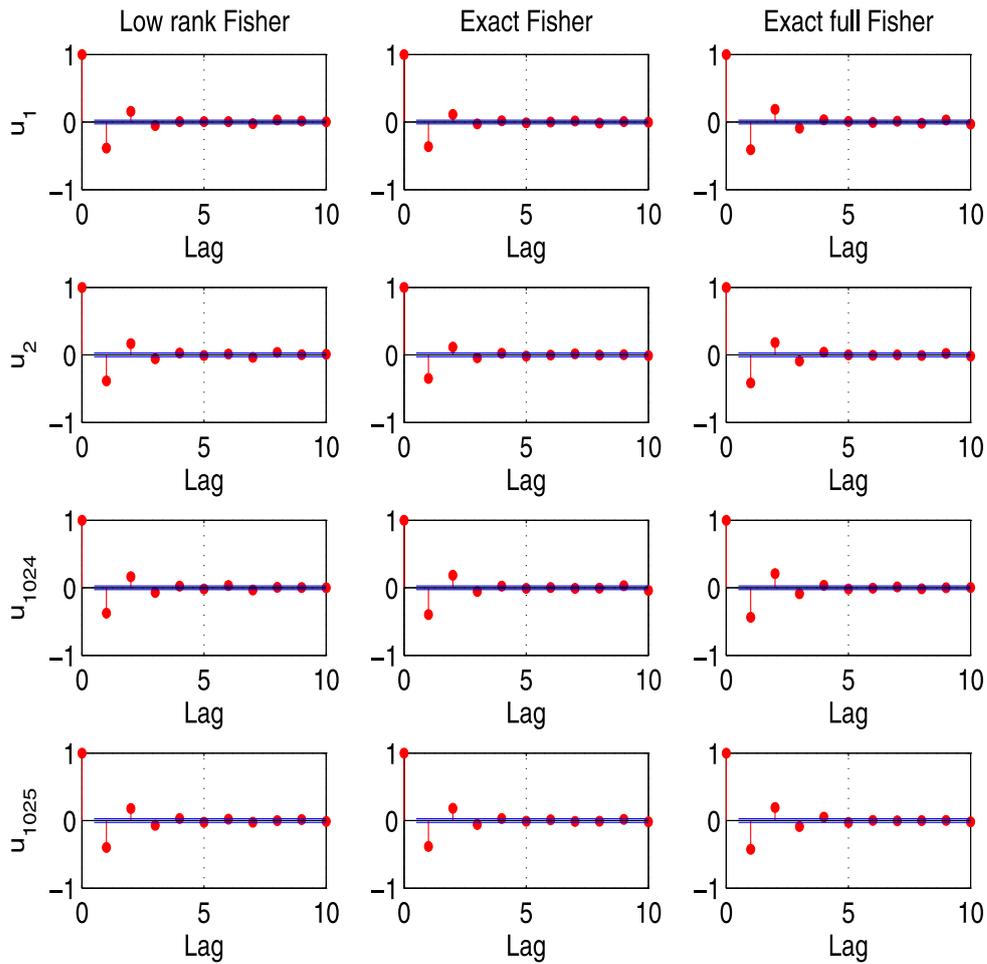


Figure 9. MCMC results of three sRMHMC methods with (i) the low rank Gauss–Newton Hessian (left column), (ii) the exact Gauss–Newton Hessian (middle column) and (iii) the full Hessian at the MAP point (right column). In the figure are the autocorrelation function plot of the first two (1 and 2) and last two (1024 and 1025) parameters.

therefore conclude that low rank approach is the least computational extensive, yet it maintains the attractive features of the original RMHMC. As such, it is the most suitable method for large-scale Bayesian inverse problems with costly PDE solves.

8. Conclusions and future work

We have proposed the adoption of a computationally inexpensive Riemann manifold Hamiltonian Monte method to explore the posterior of large-scale Bayesian inverse problems governed by PDEs in a highly efficient manner. We first adopt an infinite dimensional Bayesian framework to guarantee that the inverse formulation is well-defined. In particular, we postulate a Gaussian prior measure on the parameter space and assume regularity for the

likelihood. This leads to a well-defined posterior distribution. Then, we discretize the posterior using the standard FEM and a MTT, and apply the RMHMC method on the resulting discretized posterior in finite dimensional parameter space. We present an adjoint technique to efficiently compute the gradient, the Hessian, and the third derivative of the potential function that are required in the RMHMC context. This is at the expense of solving a few extra PDEs: one for the gradient, two for a Hessian-vector product, and four for the product of third order derivative with a matrix. For large-scale problems, repeatedly computing the action of the Hessian and third order derivative is too computationally expensive and this motivates us to design a simplified RMHMC in which the Fisher information matrix is computed once at the MAP point. We further reduce the effort by constructing low rank approximation of the Fisher information using a RSVD technique. The effectiveness of the proposed approach is demonstrated on a number of numerical results up to 1025 parameters in which the computational gain is about two orders of magnitude while maintaining the quality of the original RMHMC method in generating (almost) uncorrelated/independent samples.

For more challenging inverse problems with significant change of metric tensor across the parameter space, we expect that sRMHMC with constant metric tensor is inefficient. In that case, RMHMC seems to be a better option, but it is too computationally extensive for large-scale problems. Ongoing work is to explore approximation of the RMHMC methods in which we approximate the trace and the third derivative in (31) using adjoint and randomized techniques.

Acknowledgments

The first author is grateful for the support from DOE grants DE-SC0010518 and DE-SC0011118.

References

- [1] Bui-Thanh T and Ghattas O 2012 Analysis of the Hessian for inverse scattering problems. Part I: Inverse shape scattering of acoustic waves *Inverse Problems* **28** 055001
- [2] Bui-Thanh T and Ghattas O 2012 Analysis of the Hessian for inverse scattering problems. Part II: Inverse medium scattering of acoustic waves *Inverse Problems* **28** 055002
- [3] Bui-Thanh T and Ghattas O 2013 Analysis of the Hessian for inverse scattering problems. Part III: Inverse medium scattering of electromagnetic waves *Inverse Probl. Imaging* **7** 1139–55
- [4] Vogel C R 2002 *Computational Methods for Inverse Problems (Frontiers in Applied Mathematics)* (Philadelphia, PA: SIAM)
- [5] Vexler B 2004 Adaptive finite element methods for parameter identification problems *PhD Thesis* University of Heidelberg
- [6] Franklin J N 1970 Well-posed stochastic extensions of ill-posed linear problems *J. Math. Anal. Appl.* **31** 682–716
- [7] Lehtinen M S, Päivärinta L and Somersalo E 1989 Linear inverse problems for generalized random variables *Inverse Problems* **5** 599–612
- [8] Lasanen S 2002 Discretizations of generalized random variables with applications to inverse problems *PhD Thesis* University of Oulu
- [9] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559
- [10] Piiroinen P 2005 Statistical measurements, experiments, and applications *PhD Thesis* Department of Mathematics and Statistics, University of Helsinki
- [11] Nocedal J and Wright S J 2006 *Numerical Optimization* 2nd edn (Berlin: Springer)
- [12] Bui-Thanh T and Ghattas O 2014 Randomized maximum likelihood sampling for large-scale Bayesian inverse problems In preparation
- [13] Ciarlet P G 1978 *The Finite Element Method for Elliptic Problems* (Amsterdam: North-Holland)

- [14] Ilić M, Liu F, Turner I and Anh V 2005 Numerical approximation of a fractional-in-space diffusion equation *Frac. Calc. Appl. Anal.* **8** 323–41
- [15] Peter Simpson D 2008 Krylov subspace methods for approximating functions of symmetric positive definite matrices with applications to applied statistics and anomalous diffusion *PhD Thesis* School of Mathematical Sciences, Queensland University of Technology
- [16] Bui-Thanh T 2014 On the finite element approximation of infinite dimensional Bayesian inverse problems unpublished
- [17] Cotter S L, Roberts G O, Stuart A M and White D 2013 MCMC methods for functions: modifying old algorithms to make them faster *Stat. Sci.* **28** 424–46
- [18] Ottobre M, Pillai N S, Pinski F J and Stuart A M 2014 A function space HMC algorithm with second order Langevin diffusion limit arXiv:1308.0543
- [19] Girolami M and Calderhead B 2011 Riemann manifold Langevin and Hamiltonian Monte Carlo methods *J. R. Stat. Soc. B* **73** 123–214
- [20] Pearl Flath H, Wilcox L C, Akçelik V, Hill J, van Bloemen Waanders B and Ghattas O 2011 Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations *SIAM J. Sci. Comput.* **33** 407–32
- [21] Bui-Thanh T, Ghattas O, Martin J and Stadler G 2013 A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion *SIAM J. Sci. Comput.* **35** 2494–523
- [22] Halko N, Martinsson P-G and Tropp J A 2011 Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions *SIAM Rev.* **53** 217–88
- [23] Golub G H and van Loan C F 1996 *Matrix Computations (Johns Hopkins Studies in the Mathematical Sciences)* 3rd edn (Baltimore, MD: John Hopkins University Press)