

MCMC METHODS FOR FUNCTIONS: MODIFYING OLD ALGORITHMS TO MAKE THEM FASTER

S.L. COTTER¹, G.O. ROBERTS², A.M. STUART³ AND D. WHITE³ *

Abstract. Many problems arising in applications result in the need to probe a probability distribution for functions. Examples include Bayesian nonparametric statistics and conditioned diffusion processes. Standard MCMC algorithms typically become arbitrarily slow under the mesh refinement dictated by nonparametric description of the unknown function. We describe an approach to modifying a whole range of MCMC methods which ensures that their speed of convergence is robust under mesh refinement. In the applications of interest the data is often sparse and the prior specification is an essential part of the overall modeling strategy. The algorithmic approach that we describe is applicable whenever the desired probability measure has density with respect to a Gaussian process or Gaussian random field prior, and to some useful non-Gaussian priors constructed through random truncation. Applications are shown in density estimation, data assimilation in fluid mechanics, subsurface geophysics and image registration. The key design principle is to formulate the MCMC method for functions. This leads to algorithms which can be implemented via minor modification of existing algorithms, yet which show enormous speed-up on a wide range of applied problems.

1. Introduction. The use of Gaussian process (or field) priors is widespread in statistical applications (geostatistics [58], non-parametric regression [30], Bayesian emulator modeling [43], density estimation [2] and inverse quantum theory [34] to name but a few substantial areas where they are commonplace). The success of using Gaussian priors to model an unknown function stems largely from the model flexibility they afford, together with recent advances in computational methodology (particularly MCMC for exact likelihood-based methods). In this paper we describe a wide class of statistical problems, and an algorithmic approach to their study, which adds to the growing literature concerning the use of Gaussian process priors. To be concrete, we consider a process $\{u(x); x \in D\}$ for $D \subseteq \mathbb{R}^d$ for some d . In most of the examples we consider here u is not directly observed: it is hidden (or latent) and some complicated nonlinear function of it generates the data at our disposal.

Gaussian processes [3] can be characterized by either the covariance or inverse covariance (precision) operator. In most statistical applications, the covariance is specified. This has the major advantage that the distribution can be readily marginalized to suit a prescribed statistical use. For instance in geostatistics, it is often enough to consider the joint distribution of the process at locations where data is present. However the inverse covariance specification has particular advantages in the interpretability of parameters when there is information about the local structure of u . (For example hence the advantages of using Markov random field models in image analysis.) In the context where x varies over a continuum (such as ours) this creates particular computational difficulties since we can no longer work with a projected prior chosen to reflect available data and quantities of interest (eg $\{u(x_i); 1 \leq i \leq m\}$ say. Instead it is necessary to consider the entire distribution of $\{u(x); x \in D\}$. This poses major computational challenges, particularly in avoiding unsatisfactory compromises between model approximation (discretization in x typically) and computational cost.

There is a growing need in many parts of applied mathematics to blend data with sophisticated models involving partial and/or stochastic nonlinear differential equations (PDEs/SDEs). In particular, credible mathematical models must respect physical laws and/or Markov conditional independence relationships, which are typ-

*¹ Mathematics Institute University of Oxford; ² Statistics and ³ Mathematics Department, University of Warwick, Coventry CV4 7AL, UK

ically expressed through differential equations. Gaussian priors arises naturally in this context for several reasons. In particular: (i) they allow for straightforward enforcement of differentiability properties, adapted to the model setting; and (ii) they allow for specification of prior information in a manner which is well-adapted to the computational tools routinely used to solve the differential equations themselves. Regarding (ii) it is notable that in many applications, it may be desirable to adopt an inverse covariance specification, or to work with the Karhunen-Loève representation of the covariance, rather than through the covariance function; both precision and Karhunen-Loève representations link naturally to efficient computational tools that have been developed in numerical analysis. Exploration of links between numerical analysis and statistics is by no means new, both from the theoretical standpoint [19] and the algorithmic viewpoint [55]. However the particular links that we develop in this paper are not yet fully exploited in applications and our aim is to highlight the possibility of doing so.

This paper will consider MCMC based computational methods for simulating from posterior distributions of the type described above. Although this is motivated by Bayesian statistical applications, our approach can be applied to any simulation in which the target distribution is absolutely continuous with respect to a reference Gaussian field law. However our focus is on the commonly arising situation where the reference is a prior arising in a non-parametric Bayesian statistical application and is a Gaussian process or Gaussian random field. We also study some generalizations of these priors which arise from truncation of the Karhunen-Loève expansion to a random number of terms; these can be useful to prevent over-fitting and allow the data to automatically determine the scales about which it is informative. Our purpose is to demonstrate that small, but significant, modifications of a variety of standard Markov chain-Monte Carlo (MCMC) methods lead to substantial algorithmic speed-up when tackling Bayesian estimation problems for functions. Furthermore we show that the framework adopted encompasses a range of interesting applications.

1.1. A Description of the Problem. Crucial to our algorithm construction will be a detailed understanding of the dominating reference Gaussian measure. Although prior specification might be Gaussian, it is likely that the posterior distribution is not. However the posterior will at least be absolutely continuous with respect to an appropriate Gaussian density. Typically the dominating Gaussian measure can be chosen to be the prior, with the corresponding Radon-Nikodym derivative just being a re-expression of Bayes' formula

$$\frac{d\mu}{d\mu_0}(u) \propto L(u)$$

for likelihood L and Gaussian dominating measure (prior in this case) μ_0 . This framework extends in a natural way to the case where the prior distribution is not Gaussian, but is absolutely continuous with respect to an appropriate Gaussian distribution.

It is often the case that the dimension d_y of the space in which the data is found is considerably smaller than the dimension d_u of the object u to be estimated. A frequently occurring situation of this type is where the data y is a fixed *finite-dimensional* set of observations, whilst the unknown is a *function* u [47]. If the function is modelled non-parametrically then the ratio d_y/d_u can hence be arbitrarily small and the prior specification plays a significant role in the overall modeling strategy.

When the dimension d_u of the unknown is fixed then data augmentation methods can be very effective: the data space is enlarged to $d_{y,z}$ and the hidden variable z is

sampled [61] or integrated out [18, 38]. For these problems algorithms may be very slow in the limit where $d_u/d_{y,z}$ tends to zero, when u and z are strongly dependent. In this case, understanding the problem in the limit where the augmented data (y, z) is a *function* is key to developing sampling algorithms which do not slow down as the amount of augmentation is increased [53]: naive algorithms for the sampling of the parameter u and latent function z are not ergodic because of dependence between u and z ; a different choice of variables overcomes this problem. In this paper we study the completely different asymptotic $d_y/d_u \rightarrow 0$, but show that viewing u as a *function* is the key observation underlying construction of the faster algorithms we present: standard algorithms do not generalize to the setting for functions as the Metropolis-Hastings ratio is undefined; modification of standard algorithms leads to new methods for which the Metropolis-Hastings ratio is defined.

MCMC methods form an invaluable family of algorithms for sampling probability distributions, such as those arising from Bayesian statistics [35, 49]. In particular the Metropolis [39] and Metropolis-Hastings [28] variants are widely used, and the hybrid Monte Carlo (HMC) algorithm [20] is also starting to be used. Random walk-type MCMC methods, in which the proposed moves are localized, are particularly popular because of their wide applicability. The most commonly used methodology for application of random walk-type MCMC algorithms to Bayesian inference problems for functions is to first discretize the problem, resulting in a measure on \mathbb{R}^{d_u} ; the MCMC technology is then applied to sample in finite dimensions. For random walk methods this requires constraining the proposal variance δ to be of $\mathcal{O}(1/d_u)$ [50, 52, 9]. The number of steps required to sample is then inversely proportional to the proposal variance and is hence of $\mathcal{O}(d_u)$. We will show, in contrast, that defining an appropriate modification of the random walk for functions, and then discretizing, leads to algorithms with no restriction on δ in terms of the dimension n of the finite dimensional approximation space and hence that the number of steps required to sample the measure is $\mathcal{O}(1)$. In practice this means that larger moves can be made in state space and hence more rapid sampling. In particular, refining the mesh size for the non-parametric representation of the functions of interest does not increase the number of steps required to explore the probability measure. We will show that similar ideas apply to other standard MCMC methods which, when modified to apply to functions, give substantial speed-up. In particular we will highlight functional versions of the Langevin (MALA) and HMC methods. Analyzing the algorithmic innovations which are the basis for this work constitutes a significant challenge for the theory of MCMC. The paper [27] shows that the use of Wasserstein spectral gaps is well-adapted to the study of infinite dimensional sampling problems, and their finite dimensional approximations. That work is focussed on random walk algorithms and their modifications, but the methodology will extend to other MCMC algorithms.

Our setting is to consider measures on function spaces which possess a density with respect to a Gaussian random field measure, or some related non-Gaussian measures. This setting arises in many applications including the Bayesian approach to inverse problems [59] and conditioned diffusion processes (SDEs) [23]. Our goals in the paper are then fourfold:

- to show that a wide range of problems may be cast in a common framework requiring samples to be drawn from a measure known via its density with respect to a Gaussian random field, or related, prior;
- to explain the principles underlying the derivation of these new MCMC algorithms for functions, leading to desirable d_u -independent mixing properties;

- to illustrate the new methods in action on some non-trivial problems, all drawn from Bayesian non-parametric models where inference is made concerning a function;
- to develop some simple theoretical ideas which give deeper understanding of the benefits of the new methods.

Section 2 describes the common framework into which many applications fit, and shows a range of examples which are used throughout the paper. In section 3 we illustrate the potential impact of the ideas presented in this paper by showing numerical results comparing the performance of a standard random walk method and the vastly improved modification of the random walk method that we introduce here. Section 4 is concerned with the reference (prior) measure μ_0 and the assumptions that we make about it; these assumptions form an important part of the model specification and are guided by both modeling and implementation issues. In section 5 we detail the derivation of a range of MCMC methods on function space, including generalizations of the random walk, MALA, independence samplers, Metropolis-within-Gibbs' samplers and the HMC method. We use a variety of problems to demonstrate the new random walk method in action: subsections 6.1, 6.2, 6.3 and 6.4 include examples arising from density estimation, two inverse problems arising in oceanography and groundwater flow, and the shape registration problem. Section 7 contains a brief analysis of these methods. We make some concluding remarks in section 8.

Throughout we denote by $\langle \cdot, \cdot \rangle$ the standard Euclidean scalar product on \mathbb{R}^m , which induces the standard Euclidean norm $|\cdot|$. We also define $\langle \cdot, \cdot \rangle_C := \langle C^{-\frac{1}{2}} \cdot, C^{-\frac{1}{2}} \cdot \rangle$ for any positive definite symmetric matrix C ; this induces the norm $|\cdot|_C := |C^{-\frac{1}{2}} \cdot|$. Given a positive-definite self-adjoint operator \mathcal{C} on a Hilbert space with inner-product $\langle \cdot, \cdot \rangle$ we will also define the new inner-product $\langle \cdot, \cdot \rangle_{\mathcal{C}} = \langle \mathcal{C}^{-\frac{1}{2}} \cdot, \mathcal{C}^{-\frac{1}{2}} \cdot \rangle$, with resulting norm denoted by $\|\cdot\|_{\mathcal{C}}$ or $|\cdot|_{\mathcal{C}}$.

Acknowledgements SLC is supported by EPSRC, ERC (*FP7/2007-2013* and grant number *239870*) and St. Cross College. GOR is supported by EPSRC-CRISM. AMS is grateful to EPSRC, ERC and ONR for financial support. DW is supported by ERC.

2. Common Structure. We will now describe a wide-ranging set of examples which fit a common mathematical framework giving rise to a probability measure $\mu(du)$ on a Hilbert space X ,¹ when given its density with respect to a random field measure μ_0 , also on X . Thus we have

$$\frac{d\mu}{d\mu_0}(u) \propto \exp(-\Phi(u)) \quad (2.1)$$

for some *potential* $\Phi : X \rightarrow \mathbb{R}$. We assume that Φ can be evaluated to any desired accuracy, by means of a numerical method. *Mesh-refinement* refers to increasing the resolution of this numerical evaluation to obtain a desired accuracy and is tied to the number d_u of basis functions or points used in a finite dimensional representation of the target function u . For many problems of interest Φ satisfies certain common properties which are detailed in Assumptions 7.1 below. These properties underlie much of the algorithmic development in this paper.

A situation where (2.1) arises frequently is nonparametric density estimation (see subsection 2.1) where μ_0 is a random process prior for the unnormalized log density, and μ the posterior. There are also many inverse problems in differential equations

¹Extension to Banach space is also possible.

which have this form (see subsections 2.2, 2.3 and 2.4). For these inverse problems we assume that the data $y \in \mathbb{R}^{d_y}$ is obtained by applying an operator² \mathcal{G} to the unknown function u and adding a realization of a mean zero random variable with density ρ supported on \mathbb{R}^{d_y} , thereby determining $\mathbb{P}(y|u)$. That is,

$$y = \mathcal{G}(u) + \eta, \quad \eta \sim \rho. \quad (2.2)$$

After specifying $\mu_0(du) = \mathbb{P}(du)$, Bayes' theorem gives $\mu(dy) = \mathbb{P}(u|y)$ with $\Phi(u) = -\ln \rho(y - \mathcal{G}(u))$. We will work mainly with Gaussian random field priors $\mathcal{N}(0, \mathcal{C})$, although we will also consider generalizations of this setting found by random truncation of the Karhunen-Loève expansion of a Gaussian random field. This leads to non-Gaussian priors, but much of the methodology for the Gaussian case can be usefully extended, as we will show.

2.1. Density Estimation. Consider the problem of estimating the probability density function $\rho(x)$ of a random variable supported on $[-\ell, \ell]$, given d_y i.i.d observations y_i . To ensure positivity and normalization we may write

$$\rho(x) = \frac{\exp(u(x))}{\int_{-\ell}^{\ell} \exp(u(s)) ds}. \quad (2.3)$$

If we place a Gaussian process prior μ_0 on u and apply Bayes theorem then we obtain formula (2.1) with $\Phi(u) = -\sum_{i=1}^{d_y} \ln \rho(y_i)$ and ρ given by (2.3).

2.2. Data Assimilation in Fluid Mechanics. In weather forecasting and oceanography it is frequently of interest to determine the initial condition u for a PDE dynamical system modeling a fluid, given observations [4, 32]. To gain insight into such problems we consider a model of incompressible fluid flow, namely either the Stokes ($\gamma = 0$) or Navier-Stokes equation ($\gamma = 1$), on a two dimensional unit torus \mathbb{T}^2 :

$$\begin{aligned} \partial_t v - \nu \Delta v + \gamma v \cdot \nabla v + \nabla p &= \psi, \quad \forall (x, t) \in \mathbb{T}^2 \times (0, \infty) \\ \nabla \cdot v &= 0, \quad \forall t \in (0, \infty) \\ v(x, 0) &= u(x), \quad x \in \mathbb{T}^2. \end{aligned} \quad (2.4)$$

A simple model of the situation arising in weather forecasting is to determine v from *Eulerian data* of the form $y = \{y_{j,k}\}_{j,k=1}^{N,M}$ where

$$y_{j,k} \sim \mathcal{N}(v(x_j, t_k), \Gamma). \quad (2.5)$$

Thus the inverse problem is to find u from y of the form (2.2) with $\mathcal{G}_{j,k}(u) = v(x_j, t_k)$.

In oceanography *Lagrangian data* is often encountered: data is gathered from the trajectories $z_j(t)$ advected by the velocity field of interest, and thus satisfy the integral equation

$$z_j(t) = z_{j,0} + \int_0^t v(z_j(s), s) ds. \quad (2.6)$$

Data is of the form

$$y_{j,k} \sim \mathcal{N}(z_j(t_k), \Gamma). \quad (2.7)$$

Thus the inverse problem is to find u from y of the form (2.2) with $\mathcal{G}_{j,k}(u) = z_j(t_k)$.

²This operator, mapping the unknown function to the measurement space, is sometimes termed the observation operator in the applied literature; however we do not use that terminology in the paper.

2.3. Groundwater Flow. In study of groundwater flow an important inverse problem is to determine the permeability k of the subsurface rock from measurements of the head (water table height) p [36]. To ensure the (physically required) positivity of k we write $k(x) = \exp(u(x))$ and recast the inverse problem as one for the function u . The head p solves the PDE

$$\begin{aligned} -\nabla \cdot (\exp(u) \nabla p) &= g, & x \in D \\ p &= h, & x \in \partial D. \end{aligned} \quad (2.8)$$

Here D is a domain containing the measurement points x_i and ∂D its boundary; in the simplest case g and h are known. The forward solution operator is $\mathcal{G}(u)_j = p(x_j)$. The inverse problem is to find u , given y of the form (2.2).

2.4. Image Registration. In many applications arising in medicine and security it is of interest to calculate the distance between a curve Γ_{obs} given only through a finite set of noisy observations, and a curve Γ_{db} from a database of known outcomes. As we demonstrate below, this may be recast as an inverse problem for two functions, the first, η , representing reparameterization of the database curve Γ_{db} and the second, p , representing a momentum variable, normal to the curve Γ_{db} , which initiates a dynamical evolution of the reparameterized curve in an attempt to match observations of the curve Γ_{obs} . This approach to inversion is described in [11] and developed in the Bayesian context in [12]. Here we outline the methodology.

Suppose for a moment that we know the entire observed curve Γ_{obs} and that it is noise free. We parameterize Γ_{db} by q_{db} and Γ_{obs} by q_{obs} , $s \in [0, 1]$. We wish to find a path $q(s, t)$, $t \in [0, 1]$, between Γ_{db} and Γ_{obs} , satisfying

$$q(s, 0) = q_{\text{db}}(\eta(s)), \quad q(s, 1) = q_{\text{obs}}(s), \quad (2.9)$$

where η is an orientation-preserving reparameterization. Following the methodology of [40, 21, 63] we constrain the motion of the curve $q(s, t)$ by asking that the evolution between the two curves results from the differential equation

$$\frac{\partial}{\partial t} q(s, t) = v(q(s, t), t). \quad (2.10)$$

Here $v(x, t)$ is a time-parameterized family of vector fields on \mathbb{R}^2 chosen as follows. We define a metric on the “length” of paths as

$$\int_0^1 \frac{1}{2} \|v\|_B^2 dt, \quad (2.11)$$

where B is some appropriately chosen Hilbert space. The dynamics (2.10) is defined by choosing an appropriate v which minimizes this metric, subject to the end point constraints (2.9).

In [11] it is shown that this minimization problem can be solved via a dynamical system obtained from the Euler-Lagrange equation. This dynamical system yields $q(s, 1) = G(p, \eta, s)$, where p is an initial momentum variable normal to Γ_{db} , and η is the reparameterization. In the perfectly observed scenario the optimal values of $u = (p, \eta)$ solve the equation $G(u, s) := G(p, \eta, s) = q_{\text{obs}}(s)$.

In the partially and noisily observed scenario we are given observations

$$\begin{aligned} y_j &= q_{\text{obs}}(s_j) + \eta_j \\ &= G(u, s_j) + \eta_j \end{aligned}$$

for $j = 1, \dots, J$; the η_j represent noise. Thus we have data in the form (2.2) with $\mathcal{G}_j(u) = G(u, s_j)$. The inverse problem is to find the distributions on p and η , given a prior distribution on them, a distribution on η and the data y .

2.5. Conditioned Diffusions. There are many interesting applications involving diffusion processes, conditioned in one way or another [26]. The objective is to find $u(t)$ solving the equation

$$du(t) = f(u(t))dt + \gamma dB(t),$$

where B is a Brownian motion, and where u is conditioned on, for example (i) end-point constraints (bridge diffusions, arising in econometrics and chemical reactions); (ii) observation of a single sample path $y(t)$ given by

$$dy(t) = g(u(t))dt + \sigma dW(t)$$

for some Brownian motion W (continuous time signal processing); or (iii) discrete observations of the path given by

$$y_j = h(u(t_j)) + \eta_j.$$

For all three problems use of the Girsanov formula enables them to be written in the form (2.1).

3. Faster Algorithms. One of the aims of this paper is to describe how a number of widely used MCMC methods can be modified so that they are much more efficient at sampling the measure $\mu(du)$ for functions given by (2.1). In order to illustrate this idea we present the results obtained from sampling the posterior by two different methods; we use the fluid mechanics data assimilation problem of subsection 2.2, with simulated Lagrangian data. In the first, (a), we discretize the space X and then apply a standard random walk method; in the second, (b), we apply a modification of the random walk method on X , introduced as the proposal pCN in section 5, and then discretize. In both cases the discretization introduces a mesh size Δx and the finite dimensional sampling problem is of dimension $d_u \propto \Delta x^{-2}$.

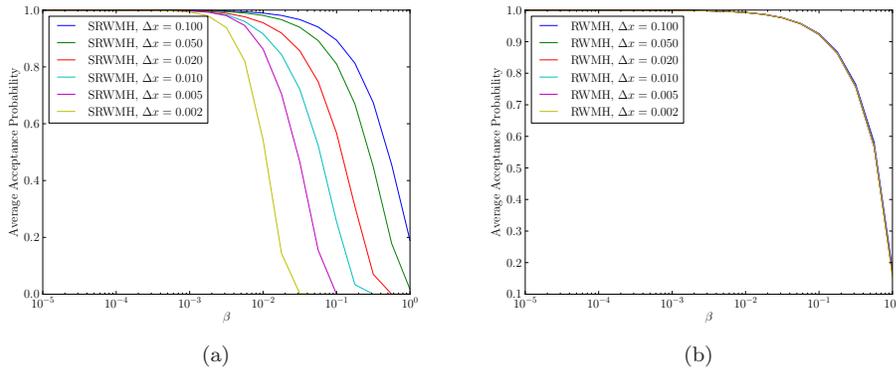


FIG. 3.1. Acceptance probabilities versus mesh-spacing, with (a) Standard Random Walk (b) Modified Random Walk (pCN)

Figures 3.1(a) and 3.1(b) show the average acceptance probability curves, as a function of a parameter β measuring the proposal variance, computed by the standard

and the modified random walk methods. It is instructive to imagine running the algorithms when tuned to obtain an average acceptance probability of, say, 0.25. Note that for the standard method, Figure 3.1(a), the acceptance probability curves shift to the left as the mesh is refined, meaning that smaller proposal variances are required to obtain the same acceptance probability as the mesh is refined. However for the new method shown in Figure 3.1(b), the acceptance probability curves have a limit as the mesh is refined and hence as the random field model is represented more accurately; thus a fixed proposal variance can be used to obtain the same acceptance probability at all levels of mesh refinement. The practical implications of this difference in acceptance probability curves is that the number of steps required by the new method is independent of the number of mesh points d_u used to represent the function, whilst of the old random walk method it is not. The new method thus mixes more rapidly than the standard method and, furthermore, the disparity in mixing rates becomes greater as the mesh is refined. Theoretical results substantiating this have recently been obtained: the paper [27] studies the spectral gap of the standard random walk, showing that it shrinks to zero as d_u grows, and the spectral gap of the pCN method, showing that it stabilizes to a limiting value as d_u grows. Furthermore, the papers [37, 45] identify diffusion limits of the standard random walk and of the pCN method; the scalings required for the proposal variance β reflect the dimension-dependence of the standard random walk, and the dimension-independent properties of the pCN method.

Similar acceptance probability curves can also be seen for other MCMC methods, suggesting small modifications of standard algorithms which lead to significant computational savings. For example, modifications of standard Langevin based proposals to the conditioned diffusions problems described in subsection 2.5 can be found in [8, 23] and modifications of the standard Hybrid Monte Carlo (HMC) algorithm can be found in [6]. In the remainder of the paper we show how such new methods can be designed, study their properties theoretically, and provide numerical illustrations of their applicability.

4. Specification of the Reference Measure. The class of algorithms that we describe are primarily based on measures defined through density with respect to a Gaussian random field model $\mu_0 = \mathcal{N}(0, \mathcal{C})$. To be able to implement the algorithms in this paper in an efficient way it is necessary to make assumptions about this Gaussian reference measure. We assume that information about μ_0 can be obtained in at least one of the following three ways:

1. the eigenpairs (ϕ_i, λ_i^2) of \mathcal{C} are known so that exact draws from μ_0 can be made from truncation of the Karhunen-Loève expansion (KL, below) and that, furthermore, efficient methods exist for evaluation of the resulting sum (such as the FFT);
2. exact draws from μ_0 can be made on a mesh, for example by building on exact sampling methods for Brownian motion or the stationary Ornstein-Uhlenbeck (OU) process or other simple Gaussian process priors;
3. the precision operator $\mathcal{L} = \mathcal{C}^{-1}$ is known and efficient numerical methods exist for the inversion of $(I + \zeta\mathcal{L})$ for $\zeta > 0$.

These assumptions are not mutually exclusive and for many problems two or more of these will be possible. Note that in many applications the precision operator \mathcal{L} will be a differential operator. The book [55] describes the literature concerning methods for sampling from Gaussian random fields, and links with efficient numerical methods for inversion of differential operators such as $(I + \zeta\mathcal{L})$.

4.1. The Karhunen-Loève Expansion. The book [3] introduces the Karhunen-Loève expansion and its' properties. Let $\mu_0 = \mathcal{N}(0, \mathcal{C})$ denote a Gaussian measure on a Hilbert space X . Recall that the orthonormalized eigenvalue/eigenfunction pairs of \mathcal{C} form an orthonormal basis for X and solve the problem

$$\mathcal{C}\phi_i = \lambda_i^2 \phi_i, \quad i = 1, 2, \dots$$

Furthermore we assume that the operator is *trace-class*:

$$\sum_{i=1}^{\infty} \lambda_i^2 < \infty. \quad (4.1)$$

Draws from the centered Gaussian measure μ_0 can then be made as follows. Let $\{\xi_i\}_{i=1}^{\infty}$ denote an independent sequence of normal random variables with distribution $\mathcal{N}(0, \lambda_i^2)$ and consider the random function

$$u(x) = \sum_{i=1}^{\infty} \xi_i \phi_i(x). \quad (4.2)$$

This series converges in $L^2(\Omega; X)$ under the trace-class condition (4.1). It is sometimes useful, both conceptually and for purposes of implementation, to think of the unknown function u as being the infinite sequence $\{\xi_i\}_{i=1}^{\infty}$, rather than the function with these expansion coefficients.

We let \mathcal{P}_d denotes projection onto the first d modes of the Karhunen-Loève basis:

$$\mathcal{P}^{d_u} u(x) = \sum_{i=1}^{d_u} \xi_i \phi_i(x). \quad (4.3)$$

If the series (4.3) can be summed quickly on a grid then this provides an efficient method for computing exact samples from truncation of μ_0 to a finite dimensional space. When we refer to *mesh-refinement* then, in the context of the prior, this refers to increasing the number of terms d_u used to represent the target function u .

4.2. Random Truncation and Sieve Priors. Non-Gaussian priors can be constructed from the KL expansion (4.3) by allowing d_u itself to be a random variable supported on \mathbb{N} ; we let $p(i) = \mathbb{P}(d_u = i)$. Much of the methodology in this paper can be extended to these priors. A draw from such a prior measure can be written as

$$u(x) = \sum_{i=1}^{\infty} \mathbb{I}(i \leq d_u) \xi_i \phi_i(x) \quad (4.4)$$

where $\mathbb{I}(i \in E)$ is the indicator function. We refer to this as *random truncation prior*. Functions drawn from this prior are non-Gaussian and almost surely C^∞ . However, expectations with respect to d_u will be Gaussian and can be less regular: they are given by the formula

$$\mathbb{E}^{d_u} u(x) = \sum_{i=1}^{\infty} \alpha_i \xi_i \phi_i(x) \quad (4.5)$$

where $\alpha_i = \mathbb{P}(d_u \geq i)$. As in the Gaussian case it can be useful, both conceptually and for purposes of implementation, to think of the unknown function u as being the infinite vector $(\{\xi_i\}_{i=1}^{\infty}, d_u)$ rather than the function with these expansion coefficients.

Making d_u a random variable has the effect of switching on (non-zero) and off (zero) coefficients in the expansion of the target function. This formulation switches the basis functions on and off in a fixed order. Random truncation as expressed by equation (4.4) is not the only variable dimension formulation. In dimension greater than one we will employ the *sieve prior* which allows every basis function to have an individual on/off switch. This prior relaxes the constraint imposed on the order in which the basis functions are switched on and off and we write

$$u(x) = \sum_{i=1}^{\infty} \chi_i \xi_i \phi_i(x) \quad (4.6)$$

where $\{\chi_i\}_{i=1}^{\infty} \in \{0, 1\}$. We define the distribution on $\chi = \{\chi_i\}_{i=1}^{\infty}$ as follows. Let ν_0 denote a reference measure formed from considering an i.i.d sequence of Bernoulli random variables with success probability one half. Then define the prior measure ν on χ to have density

$$\frac{d\nu}{d\nu_0}(\chi) \propto \exp\left(-\lambda \sum_{i=1}^{\infty} \chi_i\right)$$

where $\lambda \in \mathbb{R}^+$. As for the random truncation method it is both conceptually and practically valuable to think of the unknown function as being the pair of random infinite vectors $\{\xi_i\}_{i=1}^{\infty}$ and $\{\chi_i\}_{i=1}^{\infty}$.

5. MCMC Methods for Functions. The transferable idea in this section is that design of MCMC methods on function spaces leads, after discretization, to algorithms which are robust under mesh refinement $d_u \rightarrow \infty$. We demonstrate this idea for a number of algorithms, generalizing random walk and Langevin-based Metropolis-Hastings methods, the independence sampler, the Gibbs sampler and the HMC method; we anticipate that many other generalizations are possible.

Subsection 5.1 gives the framework for MCMC methods on a general state space. In subsection 5.2 we state and derive the new *Crank-Nicolson* proposals, arising from discretization of an OU process. In subsection 5.3 we generalize these proposals to the Langevin setting where steepest descent information is incorporated: *MALA proposals*. Subsection 5.4 is concerned with *Independence Samplers* which may be derived from particular parameter choices in the Random Walk algorithm. Subsection 5.5 introduces the idea of randomizing the choice of δ as part of the proposal which is effective for the Random Walk methods. In subsection 5.6 we introduce *Gibbs samplers* based on the Karhunen-Loève expansion (4.2). In subsection 5.7 we work with non-Gaussian priors specified through random truncation of the Karhunen-Loève expansion as in (4.4), showing how Gibbs samplers can again be used in this situation. Subsection 5.8 briefly describes the HMC method and its generalization to sampling functions.

5.1. Set-Up. We are interested in defining MCMC methods for measures μ on a Hilbert space $(X, \langle \cdot, \cdot \rangle)$, with induced norm $\|\cdot\|$, given by (2.1) where $\mu_0 = \mathcal{N}(0, \mathcal{C})$. The setting we adopt is that given in [62] where Metropolis-Hastings methods are developed in a general state space. Let $q(u, \cdot)$ denote the transition kernel on X and $\eta(du, dv)$ denote the measure on $X \times X$ found by taking $u \sim \mu$ and then $v|u \sim q(u, \cdot)$. We use $\eta^\perp(u, v)$ to denote the measure found by reversing the roles of u and v in the preceding construction of η . If $\eta^\perp(u, v)$ is equivalent (in the sense of measures)

to $\eta(u, v)$ then the Radon-Nikodym derivative $\frac{d\eta^\perp}{d\eta}(u, v)$ is well-defined and we may define the *acceptance probability*

$$a(u, v) = \min \left\{ 1, \frac{d\eta^\perp}{d\eta}(u, v) \right\}. \quad (5.1)$$

We accept the proposed move from u to v with this probability. The resulting Markov chain is μ -reversible.

A key idea underlying the new variants on random walk and Langevin based Metropolis-Hastings algorithms derived below is to use discretizations of stochastic partial differential equations (SPDEs) which are invariant for either the reference or the target measure. These SPDEs have the form, for $\mathcal{L} = \mathcal{C}^{-1}$ the precision operator for μ_0 ,

$$\frac{du}{ds} = -\mathcal{K}(\mathcal{L}u + \gamma D\Phi(u)) + \sqrt{2\mathcal{K}} \frac{db}{ds}, \quad (5.2)$$

where b is a Brownian motion in X with covariance operator the identity and $\mathcal{K} = \mathcal{C}$ or I . We refer to it as an SPDE because in many applications \mathcal{L} is a differential operator. The SPDE has invariant measure μ_0 for $\gamma = 0$ (when it is an infinite dimensional OU process) and μ for $\gamma = 1$ [25, 22, 16]. The target measure μ will behave like the reference measure μ_0 on high frequency (rapidly oscillating) functions. Intuitively this is because the data, which is finite, is not informative about the function on small scales; mathematically this is manifest in the absolute continuity of μ with respect to μ_0 given by formula (2.1). Thus discretizations of equation (5.2) with either $\gamma = 0$ or $\gamma = 1$ form sensible candidate proposal distributions. This idea was developed in the specific context of conditioned diffusions with $\gamma = 1$ in [60, 8].

5.2. Vanilla Local Proposals. The **standard random walk** proposal for $v|u$ takes the form

$$v = u + \sqrt{2\delta\mathcal{K}}\xi_0, \quad (5.3)$$

for any $\delta \in [0, \infty)$, $\xi_0 \sim \mathcal{N}(0, I)$ and $\mathcal{K} = I$ or $\mathcal{K} = \mathcal{C}$. This can be seen as a discrete skeleton of (5.2) after ignoring the drift terms. Therefore such a proposal leads to an infinite-dimensional version of the well-known Random Walk Metropolis algorithm.

The random walk proposal in finite-dimensional problems always leads to a well-defined algorithm and rarely encounters any reducibility problems ([56]). Therefore, this method can certainly be applied for arbitrarily fine mesh size. However, taking this approach does not lead to a well-defined MCMC method for *functions*. This is because η^\perp is singular with respect to η so that all proposed moves are rejected with probability 1. (We prove this in Theorem 7.3 below). Returning to the finite mesh case, algorithm mixing time therefore increase to ∞ as $d_u \rightarrow \infty$.

To define methods with convergence properties robust to increasing d_u , alternative approaches leading to well-defined and irreducible algorithms on the Hilbert space need to be considered. We consider two possibilities here, both based on Crank-Nicolson approximations [48] of the linear part of the drift. In particular, we consider discretization of equation (5.2) with the form

$$v = u - \frac{1}{2}\delta\mathcal{K}\mathcal{L}(u + v) - \delta\gamma\mathcal{K}D\Phi(u) + \sqrt{2\mathcal{K}\delta}\xi_0 \quad (5.4)$$

for a (spatial) white noise ξ_0 .

Firstly consider the discretization (5.4) with $\gamma = 0$ and $\mathcal{K} = I$. Rearranging shows that the resulting **Crank-Nicolson proposal** (CN) for $v|u$ is found by solving

$$(I + \frac{1}{2}\delta\mathcal{L})v = (I - \frac{1}{2}\delta\mathcal{L})u + \sqrt{2\delta}\xi_0. \quad (5.5)$$

It is this form that the proposal is best implemented whenever the prior/reference measure μ_0 is specified via the precision operator \mathcal{L} and when efficient algorithms exist for inversion of the identity plus a multiple of \mathcal{L} . However, for the purposes of analysis it is also useful to write this equation in the form

$$(2\mathcal{C} + \delta I)v = (2\mathcal{C} - \delta I)u + \sqrt{8\delta\mathcal{C}}w, \quad (5.6)$$

where $w \sim \mathcal{N}(0, \mathcal{C})$, found by applying the operator $2\mathcal{C}$ to equation (5.5).

A well-established principle in finite-dimensional sampling algorithms advises that proposal variance should be approximately a scalar multiple of that of the target (see for example [52]). The variance in the prior, \mathcal{C} , can provide a good approximation. The CN algorithm violates this principle: the proposal variance operator is proportional to $(2\mathcal{C} + \delta I)^{-2}\mathcal{C}^2$, suggesting that algorithm efficiency might be improved still further by obtaining a proposal variance of \mathcal{C} . In the familiar finite-dimensional case, this can be achieved by a standard *reparameterization* argument which has its origins in [29] if not before. This motivates our final local proposal in this subsection.

The **preconditioned CN** proposal (pCN) for $v|u$ is obtained from (5.4) with $\gamma = 0$ and $\mathcal{K} = \mathcal{C}$ giving the proposal

$$(2 + \delta)v = (2 - \delta)u + \sqrt{8\delta}w \quad (5.7)$$

where, again, $w \sim \mathcal{N}(0, \mathcal{C})$. As discussed after (5.5), and in section 4, there are many different ways in which the prior Gaussian may be specified. If the specification is via the precision \mathcal{L} and if there are numerical methods for which $(I + \zeta\mathcal{L})$ can be efficiently inverted then (5.5) is a natural proposal. If, however, sampling from \mathcal{C} is straightforward (via the Karhunen-Loève expansion or directly) then it is natural to use the proposal (5.7), which requires only that it is possible to draw from μ_0 efficiently. For $\delta \in [0, 2]$ the proposal (5.7) can be written as

$$v = (1 - \beta^2)^{1/2}u + \beta w, \quad (5.8)$$

where $w \sim \mathcal{N}(0, \mathcal{C})$, and $\beta \in [0, 1]$; in fact $\beta^2 = 8\delta/(2 + \delta)^2$. In this form we see very clearly a simple generalization of the finite dimensional random walk given by (5.3) with $\mathcal{K} = \mathcal{C}$.

The numerical experiments described in section 3 show that the pCN proposal significantly improves upon the naive random walk method (5.3), and similar positive results can be obtained for the CN method. Furthermore, for both the proposals (5.5) and (5.7) we show in Theorem 7.2 that η^\perp and η are equivalent (as measures) by showing that they are both equivalent to the same Gaussian reference measure η_0 , whilst in Theorem 7.3 we show that the proposal (5.3) leads to mutually singular measures η^\perp and η . This theory explains the numerical observations, and motivates the importance of designing algorithms directly on function space.

The accept-reject formula for CN and pCN is very simple. If, for some $\rho : X \times X \rightarrow$

\mathbb{R} , and some reference measure η_0 ,

$$\begin{aligned}\frac{d\eta}{d\eta_0}(u, v) &= Z \exp(-\rho(u, v)), \\ \frac{d\eta^\perp}{d\eta_0}(u, v) &= Z \exp(-\rho(v, u)),\end{aligned}\tag{5.9}$$

it then follows that

$$\frac{d\eta^\perp}{d\eta}(u, v) = \exp(\rho(u, v) - \rho(v, u)).\tag{5.10}$$

For both CN proposals (5.5) and (5.7) we show in Theorem 7.2 below that, for appropriately defined η_0 , we have $\rho(u, v) = \Phi(u)$ so that the acceptance probability is given by

$$a(u, v) = \min\{1, \exp(\Phi(u) - \Phi(v))\}.\tag{5.11}$$

In this sense the CN and pCN proposals may be seen as the *natural generalizations of random walks* to the setting where the target measure is defined via density with respect to a Gaussian, as in (2.1). This point of view may be understood by noting that the accept/reject formula is defined entirely through differences in this log density, as happens in finite dimensions for the standard random walk, if the density is specified with respect to Lebesgue measure.

5.3. MALA Proposal Distributions. The CN proposals (5.5), (5.7) contain no information about the potential Φ given by (2.1); they contain only information about the reference measure μ_0 . Indeed they are derived by discretizing the SDE (5.2) in the case $\gamma = 0$, for which μ_0 is an invariant measure. The idea behind the Metropolis-adjusted Langevin (MALA) proposals (see [54], [49] and the references therein) is to discretize an equation which is invariant for the measure μ . Thus to construct such proposals in the function space setting we discretize the SPDE (5.2) with $\gamma = 1$. Taking $\mathcal{K} = I$ and $\mathcal{K} = \mathcal{C}$ then gives the following two proposals.

The **Crank-Nicolson Langevin proposal** (CNL) is given by

$$(2\mathcal{C} + \delta)v = (2\mathcal{C} - \delta)u - 2\delta\mathcal{C}\mathcal{D}\Phi(u) + \sqrt{8\delta\mathcal{C}}w\tag{5.12}$$

where, as before $w \sim \mu_0 = \mathcal{N}(0, \mathcal{C})$. If we define

$$\rho(u, v) = \Phi(u) + \frac{1}{2}\langle v - u, \mathcal{D}\Phi(u) \rangle + \frac{\delta}{4}\langle \mathcal{C}^{-1}(u + v), \mathcal{D}\Phi(u) \rangle + \frac{\delta}{4}\|\mathcal{D}\Phi(u)\|^2,$$

then the acceptance probability is given by (5.1) and (5.10). Implementation of this proposal simply requires inversion of $(I + \zeta\mathcal{L})$, as for (5.5). The CNL method is the special case $\theta = \frac{1}{2}$ for the IA algorithm introduced in [8].

The **preconditioned Crank Nicolson Langevin proposal** (pCNL) is given by

$$(2 + \delta)v = (2 - \delta)u - 2\delta\mathcal{C}\mathcal{D}\Phi(u) + \sqrt{8\delta}w\tag{5.13}$$

where w is again a draw from μ_0 . Defining

$$\rho(u, v) = \Phi(u) + \frac{1}{2}\langle v - u, \mathcal{D}\Phi(u) \rangle + \frac{\delta}{4}\langle u + v, \mathcal{D}\Phi(u) \rangle + \frac{\delta}{4}\|\mathcal{C}^{1/2}\mathcal{D}\Phi(u)\|^2,$$

the acceptance probability is given by (5.1) and (5.10). Implementation of this proposal requires draws from the reference measure μ_0 to be made, as for (5.7). The pCNL method is the special case $\theta = \frac{1}{2}$ for the PIA algorithm introduced in [8].

5.4. Independence Sampler. Making the choice $\delta = 2$ in the pCN proposal (5.7) gives an **independence sampler**. The proposal is then simply a draw from the prior: $v = w$. The acceptance probability remains (5.11). An interesting generalization of the independence sampler is to take $\delta = 2$ in the MALA proposal (5.13) giving the proposal

$$v = -\mathcal{C}\mathcal{D}\Phi(u) + w, \quad (5.14)$$

with resulting acceptance probability given by (5.1) and (5.10) with

$$\rho(u, v) = \Phi(u) + \langle v, \mathcal{D}\Phi(u) \rangle + \frac{1}{2} \|\mathcal{C}^{1/2} \mathcal{D}\Phi(u)\|^2.$$

5.5. Random Proposal Variance. It is sometimes useful to randomize the proposal variance δ in order to obtain better mixing. We discuss this idea in the context of the pCN proposal (5.7). To emphasize the dependence of the proposal kernel on δ we denote it by $q(u, dv; \delta)$. We show in subsection 7.1 that the measure $\eta_0(du, dv) = q(u, dv; \delta) \mu_0(du)$ is well-defined and symmetric in u, v for every $\delta \in [0, \infty)$. If we choose δ at random from any probability distribution ν on $[0, \infty)$, independently from w , then the resulting proposal has kernel

$$q(u, dv) = \int_0^\infty q(u, dv; \delta) \nu(d\delta).$$

Furthermore the measure $q(u, dv) \mu_0(du)$ may be written as

$$\int_0^\infty q(u, dv; \delta) \mu_0(du) \nu(d\delta)$$

and is hence also symmetric in u, v . Hence both the CN and pCN proposals (5.5) and (5.7) may be generalized to allow for δ chosen at random independently of u and w , according to some measure ν on $[0, \infty)$. The acceptance probability remains (5.11), as for fixed δ .

5.6. Metropolis-Within-Gibbs: Blocking in Karhunen-Loève Coordinates. Any function $u \in X$ can be expanded in the Karhunen-Loève basis and hence written as

$$u(x) = \sum_{i=1}^{\infty} \xi_i \phi_i(x). \quad (5.15)$$

Thus we may view the probability measure μ given by (2.1) as a measure on the coefficients $u = \{\xi_i\}_{i=1}^{\infty}$. For any index set $I \subset \mathbb{N}$ we write $\xi^I = \{\xi_i\}_{i \in I}$ and $\xi_-^I = \{\xi_i\}_{i \notin I}$. Both ξ^I and ξ_-^I are independent and Gaussian under the prior μ_0 with diagonal covariance operators \mathcal{C}^I \mathcal{C}_-^I respectively. If we let μ_0^I denote the Gaussian $\mathcal{N}(0, \mathcal{C}^I)$ then (2.1) gives

$$\frac{d\mu}{d\mu_0^I}(\xi^I | \xi_-^I) \propto \exp(-\Phi(\xi^I, \xi_-^I)) \quad (5.16)$$

where we now view Φ as a function on the coefficients in the expansion (5.15). This formula may be used as the basis for Metropolis-within-Gibbs samplers using blocking with respect to a set of partitions $\{I_j\}_{j=1, \dots, J}$ with the property $\bigcup_{j=1}^J I_j = \mathbb{N}$. Because

the formula is defined for functions this will give rise to methods which are robust under mesh refinement when implemented in practice. We have found it useful to use the partitions $I_j = \{j\}$ for $j = 1, \dots, J-1$ and $I_J = \{J, J+1, \dots\}$. On the other hand, standard Gibbs and Metropolis-within-Gibbs samplers are based on partitioning via $I_j = \{j\}$, and do not behave well under mesh-refinement, as we will demonstrate.

5.7. Metropolis-Within-Gibbs: Random Truncation and Sieve Priors.

We will also use Metropolis-Within-Gibbs to construct sampling algorithms which alternate between updating the coefficients $\xi = \{\xi_i\}_{i=1}^\infty$ in (4.4) or (4.6), and the integer d_u , for (4.4), or the infinite sequence $\chi = \{\chi_i\}_{i=1}^\infty$ for (4.6). In words we alternate between the coefficients in the expansion of a function, and the parameters determining which parameters are active.

If we employ the non-Gaussian prior with draws given by (4.4) then the negative log likelihood Φ can be viewed as a function of (ξ, d_u) and it is natural to consider Metropolis-within-Gibbs methods which are based on the conditional distributions for $\xi|d_u$ and $d_u|\xi$. Note that, under the prior, ξ and d_u are independent with $\xi \sim \mu_{0,\xi} := \mathcal{N}(0, \mathcal{C})$ and $d_u \sim \mu_{0,d_u}$, the latter being supported on \mathbb{N} with $p(i) = \mathbb{P}(d_u = i)$. For fixed d_u we have

$$\frac{d\mu}{d\mu_{0,\xi}}(\xi|d_u) \propto \exp(-\Phi(\xi, d_u)), \quad (5.17)$$

with $\Phi(u)$ rewritten as a function of ξ and d_u via the expansion (4.4). This measure can be sampled by any of the preceding Metropolis-Hastings methods designed in the case with Gaussian μ_0 . For fixed ξ we have

$$\frac{d\mu}{d\mu_{0,d_u}}(d_u|\xi) \propto \exp(-\Phi(\xi, d_u)). \quad (5.18)$$

A natural biased random walk for $d_u|\xi$ arises by proposing moves from a random walk on \mathbb{N} which satisfies detailed balance with respect to the distribution $p(i)$. The acceptance probability is then

$$a(u, v) = \min\{1, \exp(\Phi(\xi, d_u) - \Phi(\xi, d_v))\}.$$

Variants on this are possible and, if $p(i)$ is monotonic decreasing, a simple random walk proposal on the integers, with local moves $d_u \rightarrow d_v = d_u \pm 1$, is straightforward to implement. Of course different proposal stencils can give improved mixing properties, but we employ this particular random walk for expository purposes.

If, instead of (4.4), we use the non-Gaussian sieve prior defined by equation (4.6), the prior and posterior measures may be viewed as measures on $u = \left(\{\xi_i\}_{i=1}^\infty, \{\chi_j\}_{j=1}^\infty\right)$. These variables may be modified as stated above via Metropolis-within-Gibbs for sampling the conditional distributions $\xi|\chi$ and $\chi|\xi$. If, for example, the proposal for $\chi|\xi$ is reversible with respect to the prior on ξ then the acceptance probability for this move is given by

$$a(u, v) = \min\{1, \exp(\Phi(\xi_u, \chi_u) - \Phi(\xi_v, \chi_v))\}.$$

In section 6.3 we implement a slightly different proposal in which, with probability $\frac{1}{2}$ a non-active mode is switched on with the remaining probability an active mode is switched off. If we define $N_{\text{on}} = \sum_{i=1}^N \chi_i$ the the probability of moving from ξ_u to a

state ξ_v in which an extra mode is switched on is

$$a(u, v) = \min \left\{ 1, \exp \left(\Phi(\xi_u, \chi_u) - \Phi(\xi_v, \chi_v) + \frac{N - N_{\text{on}}}{N_{\text{on}}} \right) \right\}.$$

Similarly, the probability of moving to a situation in which a mode is switched off is

$$a(u, v) = \min \left\{ 1, \exp \left(\Phi(\xi_u, \chi_u) - \Phi(\xi_v, \chi_v) + \frac{N_{\text{on}}}{N - N_{\text{on}}} \right) \right\}.$$

5.8. Hybrid Monte Carlo Methods. The algorithms discussed above have been based on proposals which can be motivated through discretization of an SPDE which is invariant for either the prior measure μ_0 , or for the posterior μ itself. HMC methods are based on a different idea which is to consider a Hamiltonian flow in a state space found from introducing extra ‘momentum’ or ‘velocity’ variables to complement the variable u in (2.1). If the momentum/velocity is chosen randomly from an appropriate Gaussian distribution at regular intervals then the resulting Markov chain in u is invariant under μ . Discretizing the flow, and adding an accept/reject step, results in a method which remains invariant for μ [20]. These methods can break random-walk type behaviour of methods based on local proposal [42, 41]. It is hence of interest to generalize these methods to the function sampling setting dictated by (2.1) and this is undertaken in [6]. The key novel idea required to design this algorithm is the development of a new integrator for the Hamiltonian flow underlying the method; this integrator is exact in the Gaussian case $\Phi \equiv 0$ and for this reason behaves well in infinite dimensions.

6. Computational Illustrations. In this section we describe numerical experiments designed to illustrate various properties of the sampling algorithms overviewed in this paper. We employ the examples introduced in section 2.

6.1. Density Estimation. Section 3 shows an example which illustrates the advantage of using the function-space algorithms highlighted in this paper, in comparison with standard techniques; there we compared pCN with a standard random walk. The first goal of the experiments in this subsection is to further illustrate the advantage of the function-space algorithms over standard algorithms. Specifically we compare the Metropolis-within-Gibbs method from subsection 5.6, based on the partition $I_j = \{j\}$ and labelled MwG here, with the pCN sampler from subsection 5.2. The second goal is to study the effect of prior modeling on algorithmic performance; to do this we study a third algorithm, RTM, based on sampling the randomly truncated Gaussian prior (4.4) using the Gibbs method from subsection 5.7, with the pCN sampler for the coefficient update.

6.1.1. Target Distributions. We will use the two true densities

$$\rho_1 \propto \mathcal{N}(-3, 1) \mathbb{I}(x \in (-\ell, +\ell)) + \mathcal{N}(+3, 1) \mathbb{I}(x \in (-\ell, +\ell))$$

and

$$\rho_2 \propto \exp \left(\sin \left(\frac{15\pi x}{\ell} \right) \right) \mathbb{I}(x \in (-\ell, +\ell))$$

where $\ell = 10$. (Recall that $\mathbb{I}(\cdot)$ denotes the indicator function of a set.) These densities were chosen to illustrate (a) a multi-modal distribution; and (b) a density with strong frequency dependence. These one dimensional problems are sufficient to expose the advantages of the function spaces samplers pCN and RTM over MwG.

Algorithm	δ	$\mathbb{E}[\hat{\alpha}]$
MwG	NA	0.777
pCN	0.065	0.241
RTM	0.07	0.224

TABLE 6.1

Average Acceptance Probabilities and Proposal Variance for Target ρ_2

6.1.2. Prior. Two priors are used for these experiments: the Gaussian prior given by (4.2), and the randomly truncated Gaussian given by (4.4). We apply the MwG and pCN schemes in the former case, and the RTM scheme for the latter. Both priors use the same Gaussian covariance structure for the independent ξ , namely $\xi_i \sim \mathcal{N}(0, \lambda_i^2)$, where $\lambda_i \propto i^{-2}$ for ρ_1 and $\lambda_i \propto i^{-1.001}$ for ρ_2 . Note that the eigenvalues are summable, as required for draws from the Gaussian measure to be square integrable and to be continuous. The prior for the number of active terms d_u is an exponential distribution with rate $\lambda = 0.01$.

6.1.3. Numerical Implementation. Two numerical experiments are undertaken in this section, the first involves ρ_1 and the second involves ρ_2 . In order to facilitate a fair comparison we tuned the value of δ in the pCN and RTM proposals to obtain similar average acceptance probabilities of around 0.234. This average acceptance only applies to the proposals involving $\{\xi\}_{i=1}^\infty$ updates as the d_u updates are independent of δ .

In the first experiment, involving ρ_1 , both the pCN and RTM with $\delta \approx 0.27$ led to an acceptance probability of around 0.234. We note that with the value $\delta = 2$ we obtain the independence sampler for pCN; however this sampler only accepted 12 proposals out of 10^6 MCMC steps, indicating the importance of tuning δ correctly. For MwG there is no tunable parameter, and we obtain an acceptance of around 0.99.

For the second experiment, involving ρ_2 , a similar procedure of tuning the step size to attain an average acceptance of 0.234 was undertaken. On this heavily frequency dependent problem the RTM permitted a slightly larger step size, for the same acceptance probability, than pCN, as shown in Table 6.1. It is shown below that this larger step size results in smaller correlations between MCMC samples.

6.1.4. Numerical Results. In order to compare the performance of pCN, MwG and RTM we study trace plots, correlation functions and integrated auto-correlation times. Figure 6.1 shows both a trace plot over 10^5 samples, and a truncated auto-correlation function calculated using 10^6 samples for the ρ_1 experiment. It is clear from the trace plot that the Metropolis within Gibbs proposal mixes slower than the other two samplers. This conclusion is substantiated numerically in Table 6.2 where we display the integrated autocorrelation times for the three methods. Here, and throughout, the integrated autocorrelation times are shown only to three significant figures and should only be treated as indicators of their true values – computing them accurately is notoriously difficult [57].

These integrated auto-correlation times clearly show that the pCN and RTM outperform MwG by an order of magnitude. This reflects the fact that pCN and RTM are function space samplers, designed to mix independently of the mesh-size. In contrast, the MwG method is heavily mesh-dependent, since updates are made one Fourier component at a time.

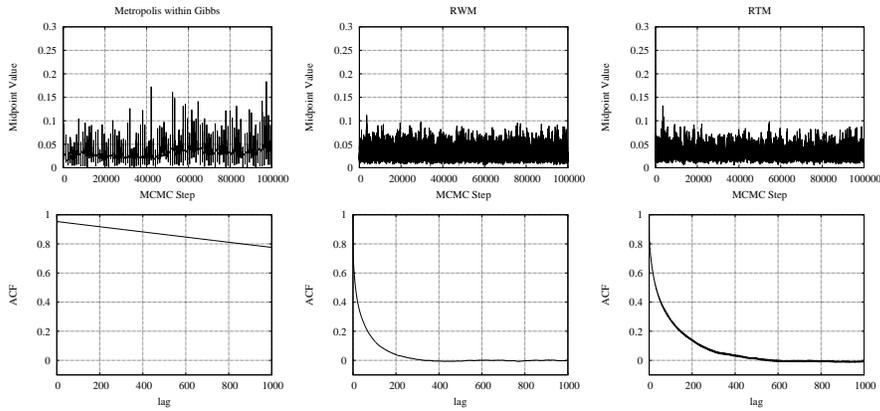


FIG. 6.1. Trace and auto-correlation plots for sampling posterior measure with true density ρ_1 using MwG, pCN and RTM methods.

Algorithm	IACT
MwG	894
pCN	73.2
RTM	143

TABLE 6.2

Approximate Integrated Autocorrelation Times for Target ρ_1

It is also interesting to compare the effect of the different priors. Note that the resulting asymptotic variance for the RTM is approximately double that of pCN. This suggests that the choice of a randomly truncated prior does not improve mixing when measured in terms of correlations. However RTM can have a reduced runtime, per unit error, when compared with pCN, as Table 6.3 shows. These results were obtained on a laptop computer (quad core M520 cpu, core freq 1197MHz with 3MB shared cache) and the calculated time taken to draw an independent sample from the posterior. The improvement of RTM over pCN is due mainly to the reduction in the number of random number generations.

The same three algorithms were also applied to the strongly frequency dependent distribution ρ_2 . Figure 6.2 shows the resulting trace and auto-correlation plots in the same format as above. These plots clearly show that the RTM outperforms the pCN method for this example. The Table 6.1.4 shows the resulting IACTs for the respective methods and these figures show that the random truncation method gives a speedup factor of 2.76.

As above the run time and IACT can be combined to express the saving in terms of time taken to draw an independent sample. These values are shown in Table 6.5:

We note that the observations concerning the target measure based on ρ_2 , and the role of the randomly truncated prior, are driven by the strong frequency dependence of the target and may not transfer to other examples, as experiments based on the target measure based on ρ_1 shows.

Algorithm	Time for 10^6 steps (s)	Time to draw an Indep Sample (s)
MwG	262	0.234
pCN	451	0.0331
RTM	278	0.0398

TABLE 6.3
Comparison of Computational Timings for Target ρ_2

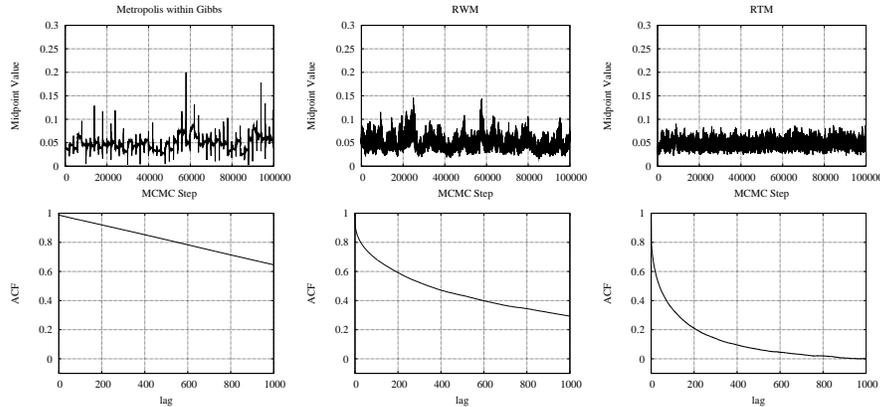


FIG. 6.2. Trace plots and correlation functions for sampling posterior measure of ρ_2 .

6.2. Data Assimilation in Fluid Mechanics. We now proceed to a more complex problem and describe numerical results which demonstrate that the function space samplers successfully sample nontrivial problems arising in applications. We study both the Eulerian and Lagrangian data assimilation problems from subsection 2.2, for the Stokes flow forward model $\gamma = 0$. It has already been demonstrated that the pCN can successfully sample from the posterior distribution for the data assimilation problems of interest [13, 14]. In this subsection we will show three features of the methods: convergence of the algorithm from different starting states, convergence with different proposal step sizes, and behaviour with random distributions for the proposal step size, as discussed in subsection 5.5.

6.2.1. Target Distributions. In this application we aim to characterize the posterior distribution on the initial condition of the two-dimensional velocity field u_0 for Stokes flow (equation 2.4 with $\gamma = 0$), given a set of either Eulerian (2.5) or Lagrangian (2.7) observations. In both cases, the posterior is of the form (2.1) with $\Phi(u) = \frac{1}{2}\|\mathcal{G}(u) - y\|_{\Gamma}^2$ with \mathcal{G} mapping u to the observation space. We choose that the observational noise covariance to be $\Gamma = \sigma^2 I$ with $\sigma = 10^{-2}$.

6.2.2. Prior. We let A be the Stokes operator defined by writing (2.4) as $dv/dt + Av = 0, v(0) = u$ in the case $\gamma = 1$ and $\psi = 0$. Thus A is ν times the negative Laplacian, restricted to a divergence free space; we also work on the space of functions whose spatial average is zero and then A is invertible. It is important to note that, in the periodic setting adopted here, A is diagonalized in the basis of divergence free Fourier series. Thus fractional powers of A are easily calculated. The prior measure is then chosen as

$$\mu_0 = \mathcal{N}(0, \delta A^{-\alpha}), \quad (6.1)$$

Algorithm	IACT
MwG	874
pCN	547
RTM	198

TABLE 6.4

Approximate Integrated Autocorrelation Times for Target ρ_2

Algorithm	Time for 10^6 steps (s)	Time to draw an Indep Sample (s)
MwG	350	0.306
pCN	518	0.283
RTM	344	0.0681

TABLE 6.5

Approximate Algorithm Run-Times for Target ρ_2

in both the Eulerian and Lagrangian data scenarios. We require $\alpha > 1$ to ensure that the covariance is trace-class (a necessary and sufficient condition for draws from it to be continuous functions). In the numerics that follow, the parameters of the prior were chosen to be $\delta = 400$ and $\alpha = 2$.

6.2.3. Numerical Implementation. The figures that follow in this section are taken from what are termed *identical twin* experiments in the data assimilation community: the same approximation of the model described above to simulate the data is also used for evaluation of Φ in the statistical algorithm in the calculation of the likelihood of u_0 given the data, with the same assumed covariance structure of the observational noise as was used to simulate the data.

Since the domain is the two dimensional torus, the evolution of the velocity field can be solved exactly for a truncated Fourier series, and in the numerics that follow we truncate this to 100 unknowns. In the case of the Lagrangian data, we integrate the trajectories (2.6) using an Euler scheme with time step $\Delta t = 0.01$. In each case we will give the values of N (number of spatial observations, or particles) and M (number of temporal observations) that were used. The observation stations (Eulerian data) or initial positions of the particles (Lagrangian data) are evenly spaced on a grid. The M observation times are evenly spaced, with the final observation time given by $T_M = 1$. The true initial condition u is chosen randomly from the prior distribution.

6.2.4. Convergence from different initial states. We consider a posterior distribution found from data comprised of 900 Lagrangian tracers observed at 100 evenly spaced times on $[0, 1]$. The data volume is high and a form of posterior consistency is observed for low Fourier modes, meaning that the posterior is approximately a Dirac mass at the truth. Observations were made of each of these tracers up to a final time $T = 1$. Figure 6.3 shows traces of the value of one particular Fourier mode of the true initial conditions. Different starting values are used for pCN and all converge to the same distribution. The proposal variance β was chosen in order to give an average acceptance probability of approximately 25%.

6.2.5. Convergence with different step size β . Here we study the effect of varying the proposal variance. Eulerian data is used with 900 observations in space and 100 observation times on $[0, 1]$. Figure 6.4 shows the different rates of convergence of the algorithm with different values of β . The value labelled β_{opt} here is chosen to give an acceptance rate of approximately 50%. This value of β is obtained by using an

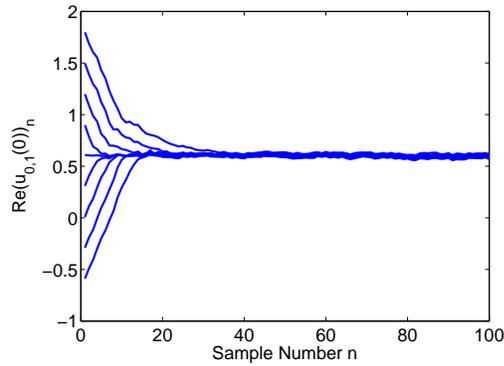


FIG. 6.3. Convergence of value of one Fourier mode of the initial condition u_0 in the pCN Markov chains with different initial states, with Lagrangian data.

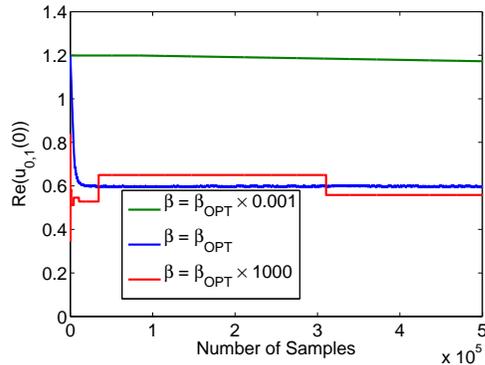


FIG. 6.4. Convergence of one of the Fourier modes of the initial condition in the pCN Markov chains with different proposal variances, with Eulerian data.

adaptive burn-in, in which the acceptance probability is estimated over short bursts and the step size β adapted accordingly. With β too small, the algorithm accepts proposed states often, but these changes in state are too small so the algorithm does not explore the state space efficiently. This can be seen in the diagram in the trace which is approximately a straight line with a small fixed negative gradient. In contrast, with β too big, larger jumps are possible. However, having large proposals often means that the algorithm is attempting to move to a state that has much smaller probability density, and so the algorithm will often reject proposals, and therefore will not explore the state space efficiently. This can be seen in the diagram in the trace which looks like a step function. The final plot in this diagram shows an optimal choice of proposal variance which compromises between these two extremes.

6.2.6. Convergence with random step size β . Here we illustrate the possibility of using a random proposal variance β , as introduced in subsection 5.5 (expressed in terms of δ and (5.7) rather than β and (5.8)). Such methods have the potential advantage of including the possibility of large and small steps in the proposal. In this example we use Eulerian data once again, this time with only 9 observation stations, with only one observation time at $T = 0.1$. Two instances of the sampler

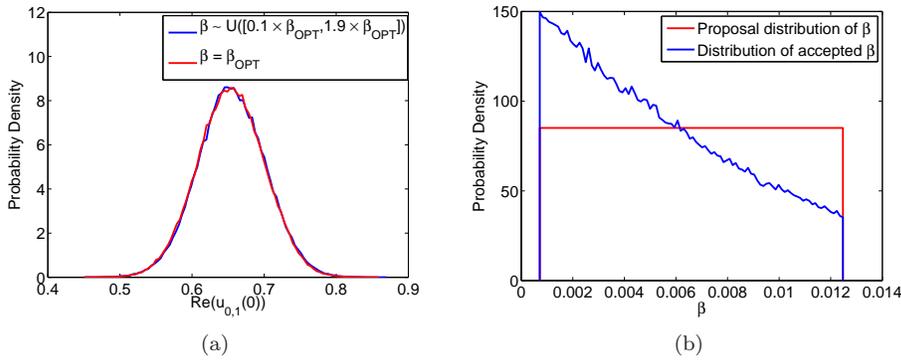


FIG. 6.5. *Eulerian Data Assimilation Example.* (a) *Empirical marginal distributions estimated using the pCN with and without random β .* (b) *Plots of the proposal distribution for β and the distribution of values for which the pCN proposal was accepted.*

were run with the same data, one with a static value of $\beta = \beta_{\text{opt}}$, and one with $\beta \sim U([0.1 \times \beta_{\text{opt}}, 1.9 \times \beta_{\text{opt}}])$. The marginal distributions for both Markov chains are shown in Figure 6.5(a). The two computed distributions are very close indeed, verifying that randomness in the proposal variance scale gives rise to (empirically) ergodic Markov chains.

Figure 6.5(b) shows the distribution of the β for which the proposed state was accepted. As expected, the initial uniform distribution is skewed as proposals with smaller jumps are more likely to be accepted.

The convergence of the method with these two choices for β were roughly comparable in this simple experiment. However it is conceivable that when attempting to explore multimodal posterior distributions it may be advantageous to have a mix of both large proposal steps which may allow large leaps between different areas of high probability density, and smaller proposal steps in order to explore a more localized region.

6.3. Subsurface Geophysics. The purpose of this section is twofold: we demonstrate another nontrivial application where function space sampling is potentially useful; and we demonstrate the use of sieve priors in this context. Key to understanding what follows in this problem is to appreciate that, for the data volume we employ, the posterior distribution can be very diffuse and expensive to explore unless severe prior modeling is imposed. This is because the homogenizing property of the elliptic PDE means that a whole range of different length-scale solutions can explain the same data. To combat this we choose very restrictive priors, either through the form of Gaussian covariance, or through the sieve mechanism, which favour a small number of active Fourier modes.

6.3.1. Target Distribution. We consider equation (2.8) in the case $D = [0, 1]^2$. Recall that the objective in this problem is to recover the permeability $\kappa = \exp(u)$. The sampling algorithms discussed here are applied to the log permeability u . In order to test the algorithms, two true permeability functions were used. These are shown in Figure 6.6 and are given by (Case 1)

$$\kappa_1(x) = \exp(u_1(x)) = \frac{1}{10} \quad (6.2)$$

and (Case 2)

$$\begin{aligned}\kappa_2(x) &= \exp(u_2(x)) \\ &= \exp\left(\frac{1}{10} + \sin(2\pi x_1) + 0.5 \cos(4\pi x_2) + 0.5 \sin(6\pi x_1) \cos(6\pi x_2)\right).\end{aligned}\quad (6.3)$$

The pressure measurement data is $y_j = p(x_j) + \sigma\eta_j$ with the η_j i.i.d. standard unit Gaussians. with measurement locations shown in Figure 6.7.

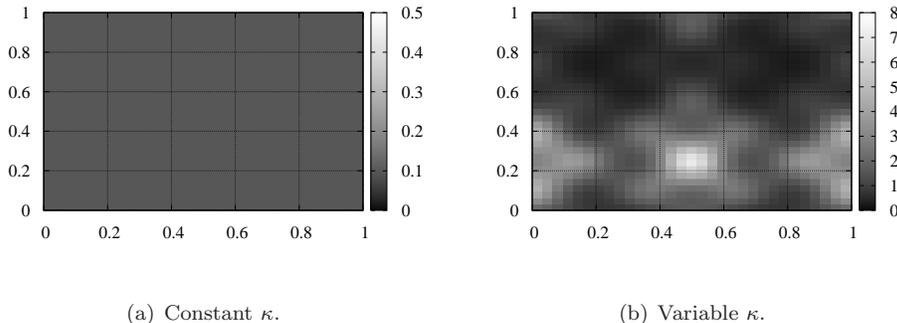


FIG. 6.6. True permeability functions used to create target distributions in subsurface geophysics application.

6.3.2. Prior. The priors will either be Gaussian, or a sieve prior based on a Gaussian. In both cases the Gaussian structure is defined via a Karhunen-Loève expansion of the form

$$u(x) = \zeta_{0,0}\varphi^{(0,0)} + \sum_{(p,q) \in \mathbb{Z}^2 \setminus \{0,0\}} \frac{\zeta_{p,q}\varphi^{(p,q)}}{(p^2 + q^2)^\alpha}, \quad (6.4)$$

where $\varphi^{(p,q)}$ are two dimensional Fourier basis functions and the $\zeta_{p,q}$ are independent random variables with distribution $\zeta_{p,q} \sim \mathcal{N}(0, 1)$ and $a \in \mathbb{R}$. To make the covariance operator trace-class (a necessary and sufficient condition for draws from it to be continuous functions) we require that $\alpha > 1$. We will consider (i) for target defined via κ_1 we take $\alpha = 1.001$ and (ii) for target defined via κ_2 we take $\alpha = 2$.

6.3.3. Numerical Implementation. The forward model is evaluated by solving equation 2.8 on the two dimensional domain $D = [0, 1]^2$ using a finite difference method with mesh of size $J \times J$. This results in a $J^2 \times J^2$ banded matrix with bandwidth J which may be solved, using a banded matrix solver, in $\mathcal{O}(J^4)$ floating point operations (see p. 171 [31]). As drawing a sample is a $\mathcal{O}(J^4)$ operation, the grid sizes used within these experiments was kept deliberately low: for target defined via κ_1 we take $J = 64$, and for target defined via κ_2 the value $J = 32$ we used. This allowed a sample to be drawn in less than 100ms and therefore 10^6 samples to be drawn in around a day. For Case 1 we use 9 measurement points, and for Case 2 we use 5 measurement points. Recall that the measurements are distributed as shown in Figure 6.7.

As with other numerics presented within this paper, the calculations were preceded by a tuning stage. During this stage the step size of the pCN method and the sieve method was selected to ensure an average acceptance of around 0.234.

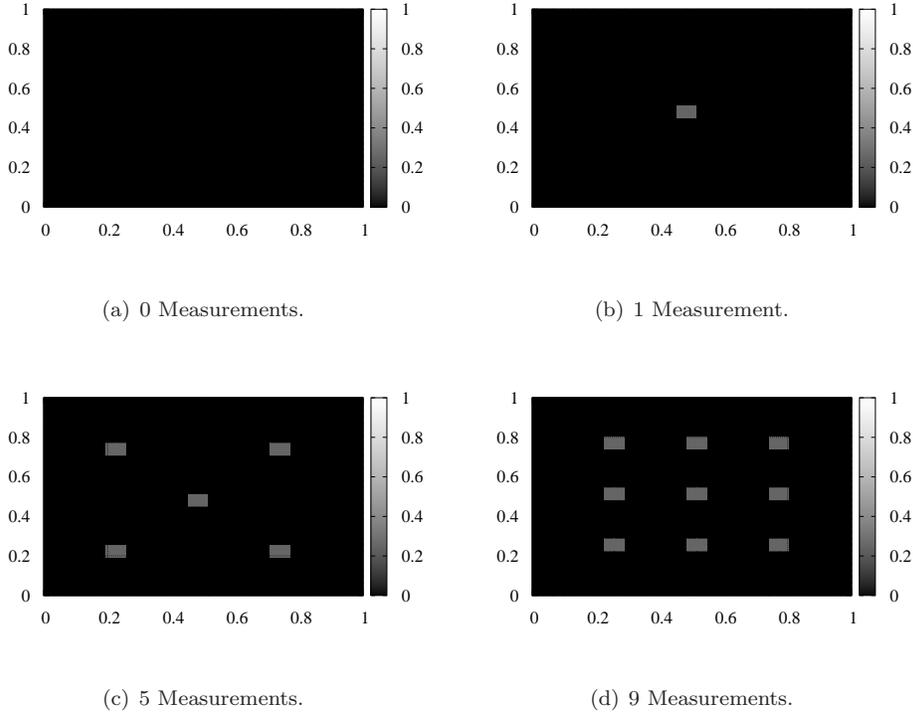
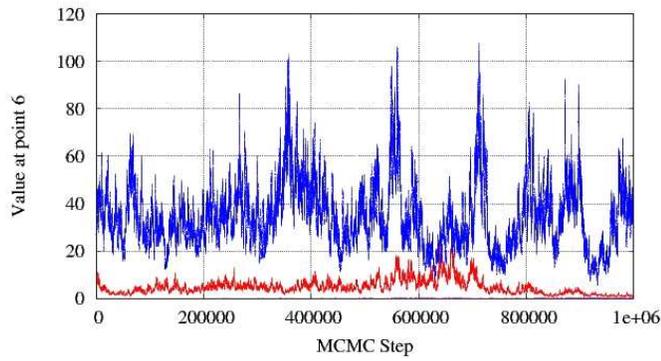


FIG. 6.7. *Measurement locations for subsurface experiments.*

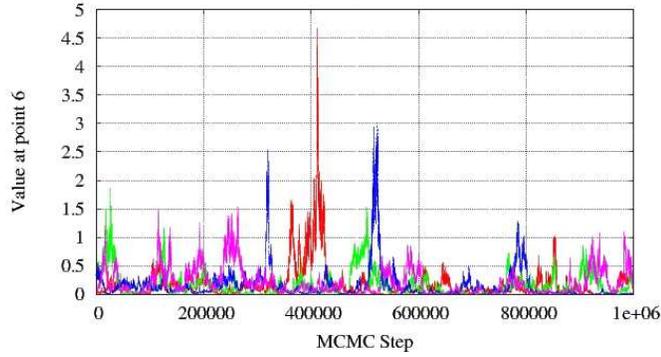
6.3.4. Numerical Results: Case 1. Recall that this corresponds to Gaussian covariance which is only just trace-class. Figure 6.8 shows trace plots obtained through application of the MwG and pCN methods to the Gaussian prior, and a pCN-based Gibbs sampler for the sieve prior, denote Sieve. Four different seeds are used. It is clear from these plots that only the MCMC chain generated by the sieve prior/algorithm combination converges in the available computational time. The other algorithms fail to converge under these test conditions. This demonstrates the importance of prior modeling assumptions for these under-determined inverse problems with multiple solutions.

6.3.5. Numerical Results: Case 2. Recall that this corresponds to Gaussian covariance which decays more rapidly in Fourier space. The experiments here are designed to resolve the question of whether the proposal pCN or MwG is better for this problem. Figure 6.9 show the corresponding trace plots for this experiment. The trace plots show that the MwG proposal is inferior. This fact is confirmed by the autocorrelation functions shown in Figure 6.10 and the integrated autocorrelation time (IACT) shown in Table 6.6, (measured with a maximum lag of 100). This also shows that the Sieve algorithm outperforms the other methods under these experimental conditions.

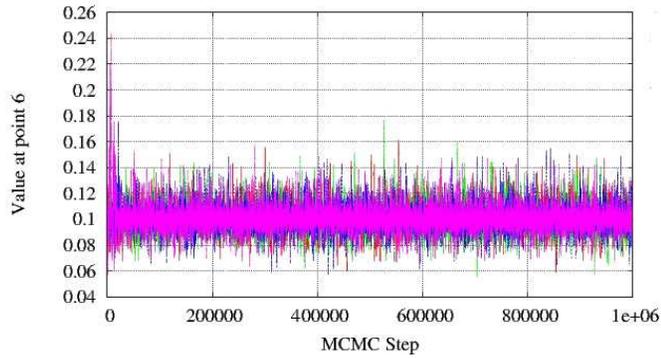
6.4. Image Registration. In this subsection we consider the image registration problem from subsection 2.4. Our primary purpose is to illustrate the idea that, in



(a) MwG



(b) pCN



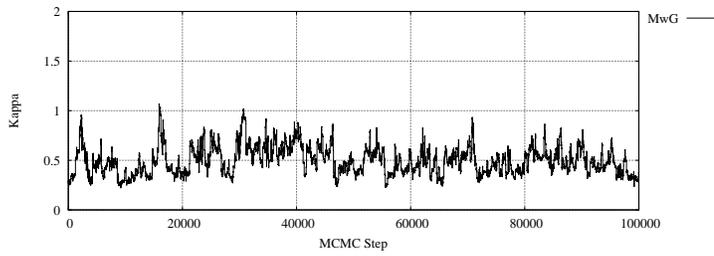
(c) Sieve

FIG. 6.8. Trace Plots for Case 1 of the subsurface geophysics application, using 5 measurements. The MwG, pCN and Sieve algorithms are compared. Different colours correspond to identical MCMC simulations with different random number generator seeds.

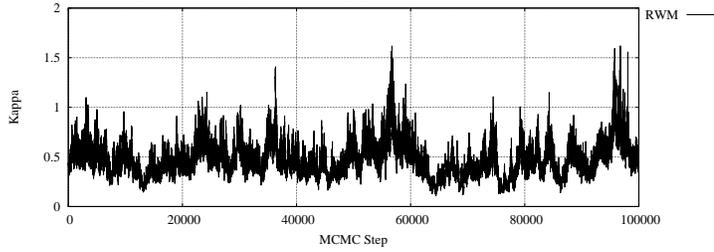
Algorithm	IACT
MwG	94.1
pCN	80.0
Sieve	73.2

TABLE 6.6

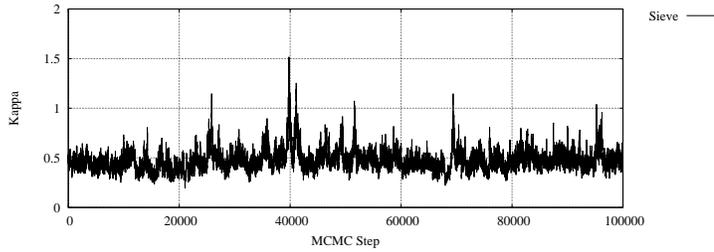
Approximate Integrated Autocorrelation Times for Target κ_2



(a) MwG



(b) pCN



(c) Sieve

FIG. 6.9. Trace Plots for Case 2 of the subsurface geophysics application, using 5 measurements. The MwG, pCN and Sieve algorithms are compared.

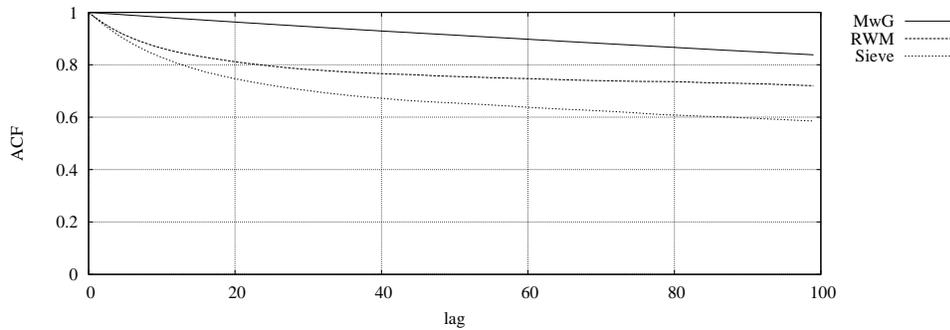


FIG. 6.10. Autocorrelation functions for Case 1 of the subsurface geophysics application. The MwG, pCN and Sieve algorithms are compared.

the function space setting, it is possible to extend the prior modeling to include an unknown observational precision, and to use conjugate Gamma priors for this parameter.

6.4.1. Target Distribution. We study the setup from subsection 2.4, with data generated from a noisily observed truth $u = (p, \eta)$ which corresponds to a smooth closed curve. We make J noisy observations of the curve where, as will be seen below, we consider the cases $J = 10, 20, 50, 100, 200, 500$ and 1000 . The noise is uncorrelated mean zero Gaussian at each location with variance $\sigma^2 = 0.01$. We will study the case where σ^2 is itself unknown, and introduce a prior on $\tau = \sigma^{-2}$.

6.4.2. Prior. The priors on the initial momentum and reparameterization are taken as

$$\begin{aligned}\mu_p(p) &= \mathcal{N}(0, \delta_1 \mathcal{H}^{-\alpha_1}), \\ \mu_\nu(\nu) &= \mathcal{N}(0, \delta_2 \mathcal{H}^{-\alpha_2}),\end{aligned}\tag{6.5}$$

where $\alpha_1 = 0.55$, $\alpha_2 = 1.55$, $\delta_1 = 30$ and $\delta_2 = 5 \times 10^{-2}$. Here $\mathcal{H} = (I - \Delta)$ denotes the Helmholtz operator in one dimension and hence the chosen values of α_i ensure trace-class prior covariance operators. As a consequence draws from the prior are continuous and square-integrable. The prior for τ is defined as

$$\mu_\tau = \text{Gamma}(\alpha_\sigma, \beta_\sigma),\tag{6.6}$$

noting that this leads to a conjugate posterior on this variable, since the observational noise is Gaussian. In the numerics that follow, we set $\alpha_\sigma = \beta_\sigma = 0.0001$.

6.4.3. Numerical Implementation. In each experiment, the data is produced using the same template shape Γ_{db} , with parameterization given by:

$$q_{\text{db}}(s) = (\cos(s) + \pi, \sin(s) + \pi), \quad s \in [0, 2\pi).\tag{6.7}$$

In the following numerics, the observed shape is chosen by first sampling an instance of p and ν from their respective prior distributions, and using the numerical approximation of the forward model to give us the parameterization of the target shape. The N observational points $\{s_i\}_{i=1}^N$ are then picked by evenly spacing them out over the interval $[0, 1)$, so that $s_i = (i - 1)/N$.

6.4.4. Finding the observational noise. We implement an MCMC method to sample from the joint distribution of (u, τ) where (recall) $\tau = \sigma^{-2}$ is the inverse observational precision. When sampling u we employ the pCN method. In this context it is possible to either: (i) implement a Metropolis-within-Gibbs sampler, alternating between use of pCN to sample $u|\tau$ and using explicit sampling from the Gamma distribution for $\tau|u$; or (ii) marginalize out τ and sample directly from the marginal distribution for u , generating samples from τ separately; we adopt the second approach.

We show that, by taking datasets with an increasing number of observations N , the true values of the functions u and the precision parameter τ can both be recovered: a form of posterior consistency.

This is demonstrated in Figure 6.11, for the posterior distribution on a low wave number Fourier coefficient in the expansion of the initial momentum p and the reparameterization η . Figure 6.12 shows the posterior distribution on the value of the observational variance σ^2 ; recall that the true value is 0.01. The posterior distribution becomes increasingly peaked close to this value as N increases.

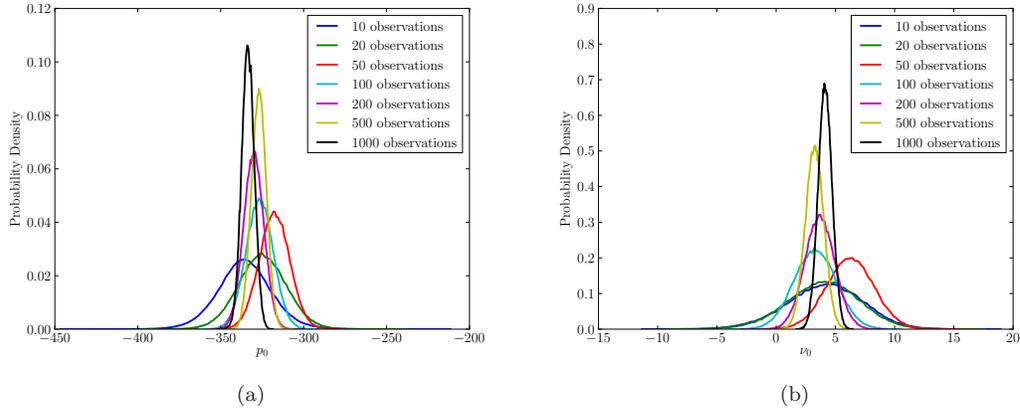


FIG. 6.11. Convergence of the lowest wave number Fourier modes in a) the initial momentum p_0 , and b) the reparameterization function ν , as the number of observations is increased, using the pCN .

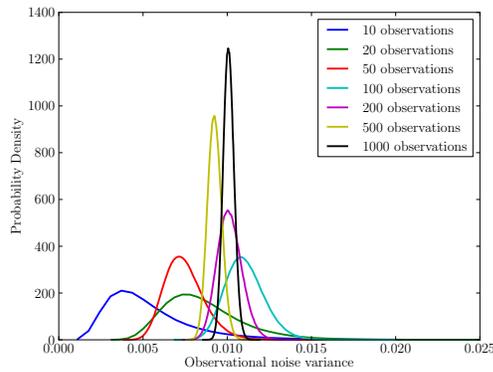


FIG. 6.12. Convergence of the posterior distribution on the value of the noise variance $\sigma^2 I$, as the number of observations is increased, sampled using the pCN .

6.5. Conditioned Diffusions. Numerical experiments which employ function space samplers to study problems arising in conditioned diffusions have been published in a number of articles. The paper [8] introduced the idea of function space samplers in this context, and demonstrated the advantage of the CNL method (5.12) over the standard Langevin algorithm for bridge diffusions; in the notation of that paper, the IA method with $\theta = \frac{1}{2}$ is our CNL method. The article [24] also demonstrated effectiveness of the CNL method, for smoothing problems arising in signal processing. The paper [6] contains numerical experiments showing comparison of the function-space HMC method from subsection 5.8 with the CNL variant of the MALA method from subsection 5.3, for a bridge diffusion problem; the function-space HMC method is superior in that context, demonstrating the power of methods which break random-walk type behaviour of local proposals.

7. Theoretical Analysis. The numerical experiments in this paper demonstrate that the function-space algorithms of Crank-Nicolson type behave well on a range of non-trivial examples. In this section we describe some basic analysis which adds theoretical weight to the choice of Crank-Nicolson discretizations which underlie these algorithms. We also show that the acceptance probability resulting from these proposals behaves as in finite dimensions: in particular that it is continuous as the scale factor δ for the proposal variance tends to zero. And finally we summarize briefly the theory available in the literature which relates to the function-space viewpoint that we highlight in this paper. We assume throughout that Φ satisfies the following assumptions:

ASSUMPTIONS 7.1. *The function $\Phi : X \rightarrow \mathbb{R}$ satisfies the following:*

1. *there exists $p > 0, K > 0$ such that, for all $u \in X$*

$$0 \leq \Phi(u; y) \leq K(1 + \|u\|^p);$$

2. *for every $r > 0$ there is $K(r) > 0$ such that, for all $u, v \in X$ with $\max\{\|u\|, \|v\|\} < r$,*

$$|\Phi(u) - \Phi(v)| \leq K(r)\|u - v\|.$$

These assumptions arise naturally in many Bayesian inverse problems where the data is finite dimensional [59]. Both the data assimilation inverse problems from subsection 2.2 are shown to satisfy Assumptions 7.1, for appropriate choice of X in [15] (Navier-Stokes) and [59] (Stokes). The groundwater flow inverse problem from subsection 2.3 is shown to satisfy these assumptions in [17], again for appropriate choice of X . It is shown in [12] that Assumptions 7.1 are satisfied for the image registration problem of subsection 2.4, again for appropriate choice of X . A wide range of conditioned diffusions satisfy Assumptions 7.1: see [24]. The density estimation problem from subsection 2.1 satisfies the second item from Assumptions 7.1, but not the first.

7.1. Why The Crank-Nicolson Choice?. In order to explain this choice we consider a one-parameter (θ) family of discretizations of equation (5.2), which reduces to the discretization (5.4) when $\theta = \frac{1}{2}$. This family is

$$v = u - \delta\mathcal{K}((1 - \theta)\mathcal{L}u + \theta\mathcal{L}v) - \delta\gamma\mathcal{K}D\Phi(u) + \sqrt{2\delta\mathcal{K}}\xi_0, \quad (7.1)$$

where $\xi_0 \sim \mathcal{N}(0, I)$ is a white noise on X . Note that $w := \sqrt{\mathcal{C}}\xi_0$ has covariance operator \mathcal{C} and is hence a draw from μ_0 . Recall that if u is the current state of the Markov chain, then v is the proposal. For simplicity we consider only Crank-Nicolson proposals, and not the MALA variants, so that $\gamma = 0$. However the analysis generalizes to the Langevin proposals in a straightforward fashion.

Rearranging (7.1) we see that the proposal v satisfies

$$v = (I - \delta\theta\mathcal{K}\mathcal{L})^{-1} \left((I + \delta(1 - \theta)\mathcal{K}\mathcal{L})u + \sqrt{2\delta\mathcal{K}}\xi_0 \right). \quad (7.2)$$

If $\mathcal{K} = I$ then the operator applied to u is bounded on X for any $\theta \in (0, 1]$. If $\mathcal{K} = \mathcal{C}$ it is bounded for $\theta \in [0, 1]$. The white noise term is almost surely in X for $\mathcal{K} = I, \theta \in (0, 1]$ and $\mathcal{K} = \mathcal{C}, \theta \in [0, 1]$. The Crank-Nicolson proposal (5.5) is found by letting $\mathcal{K} = I$ and $\theta = \frac{1}{2}$. The preconditioned Crank-Nicolson proposal (5.7) is found by setting $\mathcal{K} = \mathcal{C}$ and $\theta = \frac{1}{2}$. The following theorem explains the choice $\theta = \frac{1}{2}$.

THEOREM 7.2. Let $\mu_0(X) = 1$, let Φ satisfy Assumption 7.1(2) and assume that μ and μ_0 are equivalent as measures with Radon-Nikodym derivative (2.1). Consider the proposal $v|u \sim q(u, \cdot)$ defined by (7.2) and the resulting measure $\eta(du, dv) = q(u, dv)\mu(du)$ on $X \times X$. For both $\mathcal{K} = I$ and $\mathcal{K} = \mathcal{C}$ the measure $\eta^\perp = q(v, du)\mu(dv)$ is equivalent to η if and only if $\theta = \frac{1}{2}$. Furthermore, if $\theta = \frac{1}{2}$, then

$$\frac{d\eta^\perp}{d\eta}(u, v) = \exp(\Phi(u) - \Phi(v)).$$

Proof. We define $\eta_0(u, v)$ to be the measure $\eta(u, v)$ on $X \times X$ in the case where $\Phi \equiv 0$. The measure $\eta_0(u, v)$ is then Gaussian. First consider the case $\theta = \frac{1}{2}$. It then follows that $\eta_0(u, v)$ is symmetric in u, v . To see this it suffices to show that $\mathbb{E}^{\eta_0} v \otimes v = \mathbb{E}^{\eta_0} u \otimes u = \mathcal{C}$. This identity holds because

$$\begin{aligned} \mathbb{E} v \otimes v &= (2 - \delta\mathcal{K}\mathcal{L})^{-2}(2 + \delta\mathcal{K}\mathcal{L})^2\mathcal{C} + (2 - \delta\mathcal{K}\mathcal{L})^{-2}8\delta\mathcal{K} \\ &= (2 - \delta\mathcal{K}\mathcal{L})^{-2}\left((2 + \delta\mathcal{K}\mathcal{L})^2 - 8\delta\mathcal{K}\mathcal{L}\right)\mathcal{C} \\ &= \mathcal{C}. \end{aligned}$$

Now

$$\begin{aligned} \eta(du, dv) &= q(u, dv)\mu(du), \\ \eta_0(du, dv) &= q(u, dv)\mu_0(du), \end{aligned}$$

and the measures μ and μ_0 are equivalent. It follows that η and η_0 are equivalent and that

$$\frac{d\eta}{d\eta_0}(u, v) = \frac{d\mu}{d\mu_0}(u) = Z \exp(-\Phi(u)).$$

Since η_0 is symmetric in u, v we also have that η^\perp and η_0 are equivalent and that

$$\frac{d\eta^\perp}{d\eta_0}(u, v) = Z \exp(-\Phi(v)).$$

Since equivalence of measures is transitive it follows that η and η^\perp are equivalent with the desired Radon-Nikodym derivative.

To show that the measures are not absolutely continuous for $\theta \neq \frac{1}{2}$ it suffices to consider the purely Gaussian case where $\Phi \equiv 0$. Two Gaussian measures are either equivalent or mutually singular [16]. To show that the measures are mutually singular it suffices to show that the marginal measures $\int_X \eta_0(\cdot, dv)$ and $\int_X \eta_0^\perp(\cdot, dv)$ are mutually singular. But the first of these measures is distributed as $\mathcal{N}(0, \mathcal{C})$ and the distribution of the second is simply given by the distribution of v under the proposal (7.2). A straightforward calculation shows that v is mean zero with covariance operator

$$\mathbb{E} v \otimes v = \left(I + \delta^2(1 - 2\theta)(I - \delta\theta\mathcal{K}\mathcal{L})^{-2}\mathcal{K}^2\mathcal{L}^2\right)\mathcal{C}.$$

By the Feldman-Hajek theorem [16] mutual singularity is proved, for $\theta \neq \frac{1}{2}$, provided that the operator

$$(I - \delta\theta\mathcal{K}\mathcal{L})^{-2}\mathcal{K}^2\mathcal{L}^2$$

is not Hilbert-Schmidt. If $\mathcal{K} = \mathcal{C}$ then

$$(I - \delta\theta\mathcal{K}\mathcal{L})^{-2}\mathcal{K}^2\mathcal{L}^2$$

has all eigenvalues equal to $(I + \delta\theta)^{-2}$ and these are clearly not square summable. If $\mathcal{K} = I$ then $(I - \delta\theta\mathcal{K}\mathcal{L})^{-2}\mathcal{K}^2\mathcal{L}^2 = (\mathcal{C} + \delta\theta)^{-2}$. Since \mathcal{C} is trace class the eigenvalues of this operator are asymptotic to $(\delta\theta)^{-2}$ and are again not square summable. The result follows. \square

Note that the choice $\theta = \frac{1}{2}$ thus has the desirable property that $u \sim \mathcal{N}(0, \mathcal{C})$ implies that $v \sim \mathcal{N}(0, \mathcal{C})$: thus the prior measure is preserved under the proposal. This mimics the behaviour of the SDE (5.2) for which the prior is an invariant measure. We have thus justified the proposals (5.5) and (5.7) on function space. To complete our analysis it remains to rule out the standard random walk proposal (5.3).

THEOREM 7.3. *Consider the proposal $v|u \sim q(u, \cdot)$ defined by (5.3) and the resulting measure $\eta(du, dv) = q(u, dv)\mu(du)$ on $X \times X$. For both $\mathcal{K} = I$ and $\mathcal{K} = \mathcal{C}$ the measure $\eta^\perp = q(v, du)\mu(dv)$ is not absolutely continuous with respect to η . Thus the MCMC method is not defined on function space.*

Proof. The proof follows along similar lines to the proof of the second half of the preceding theorem. Note that, under the proposal (5.3),

$$\mathbb{E}v \otimes v = \mathcal{C} + 2\delta\mathcal{K}.$$

By the Feldman-Hajek theorem it suffices to show that $\mathcal{K}\mathcal{C}^{-1}$ is not Hilbert-Schmidt. This is immediate for both $\mathcal{K} = I$ (since \mathcal{C} is trace class) and $\mathcal{K} = \mathcal{C}$. \square

7.2. The Acceptance Probability. We now study the properties of the two Crank-Nicolson methods with proposals (5.5) and (5.7) in the limit $\delta \rightarrow 0$, showing that finite dimensional intuition carries over to this function space setting. We define

$$R(u, v) = \Phi(u) - \Phi(v)$$

and note from (5.11) that, for both of the Crank-Nicolson proposals,

$$a(u, v) = \min\{1, \exp(R(u, v))\}.$$

THEOREM 7.4. *Let μ_0 be a Gaussian measure on a Hilbert space $(X, \|\cdot\|)$ with $\mu_0(X) = 1$ and let μ be an equivalent measure on X given by the Radon-Nikodym derivative (2.1), satisfying Assumptions 7.1(1) and 7.1(2). Then both the pCN and CN algorithms with fixed δ are defined on X and, furthermore, the acceptance probability satisfies*

$$\lim_{\delta \rightarrow 0} \mathbb{E}^\eta a(u, v) = 1.$$

Proof. In [10] it is proved that, if $\mathbb{E}^\eta |R(u, v)| < \infty$, then the average acceptance probability satisfies

$$1 \geq \mathbb{E}^\eta a(u, v) \geq \exp(-\lambda) \left(1 - \frac{\mathbb{E}^\eta |R(u, v)|}{\lambda}\right)$$

for all $\lambda > \mathbb{E}^\eta |R(u, v)|$. Thus to demonstrate that

$$\lim_{\delta \rightarrow 0} \mathbb{E}^\eta a(u, v) = 1$$

it suffices to show that $\mathbb{E}^\eta |R(u, v)| \rightarrow 0$ as $\delta \rightarrow 0$.

In order to do this we first prove bounds on the moments of u, v and $u - v$ under η . Since $\rho(u, v) = \Phi(u)$ we deduce from (5.9) that

$$\frac{d\eta}{d\eta_0}(u, v) = \exp(-\Phi(u)).$$

Since Φ is bounded from below it suffices to prove the desired moment bounds under η_0 because the change of measure from η to η_0 is bounded.

Consider first the preconditioned random walk. The bound on all moments of u under \mathbb{E}^{η_0} follows because this is a Gaussian measure (Fernique theorem [16], used several times without comment in what follows), identical to \mathbb{E}^{μ_0} on u , and because $\mathbb{E}^{\mu_0} \|u\|^2$ is finite. From equation (5.7) we have that

$$v = \left(\frac{2 - \delta}{2 + \delta}\right)u + \sqrt{\frac{8\delta}{(2 + \delta)^2}}w.$$

Straightforward calculation, using the fact that u and w are independent and that $w \sim \mathcal{N}(0, \mathcal{C})$, shows that

$$\mathbb{E}^{\eta_0} \|u - v\|^2 \leq \frac{4\delta^2}{(2 + \delta)^2} \mathbb{E}^{\mu_0} \|u\|^2 + \frac{8\delta}{(2 + \delta)^2} \mathbb{E}^{\eta_0} \|w\|^2.$$

All moments of w are bounded because w is identically distributed to u under η_0 . Thus, since $u - v$ is Gaussian under η_0 , all moments $\mathbb{E}^{\eta_0} \|u - v\|^p$ are bounded and tend to zero as $\delta \rightarrow 0$ as required. We also have that

$$\mathbb{E}^{\eta_0} \|v\|^2 \leq 2 \left(\mathbb{E}^{\mu_0} \|u\|^2 + \mathbb{E}^{\eta_0} \|u - v\|^2 \right)$$

and hence the second, and all other, moments of v are bounded under η_0 too.

Now consider the standard random walk. Straightforward calculation with (5.5) shows that

$$\begin{aligned} v - u &= -2\delta(2\mathcal{C} + \delta)^{-1}u + \sqrt{8\delta\mathcal{C}}(2\mathcal{C} + \delta)^{-1}w \\ &:= a + b. \end{aligned}$$

Let the covariance operator \mathcal{C} have eigenvalues λ_k indexed by $k \in \mathbb{Z}^+$ and ordered so as to be decreasing. Note that $\sum_k \lambda_k < \infty$ as any covariance operator is necessarily trace class. Let

$$k^* = \sup_k \{\lambda_k \geq \delta^{1-\beta}\}.$$

$$\begin{aligned} \mathbb{E}^{\eta_0} \|a\|^2 &= 4\delta^2 \sum_k \frac{\lambda_k}{(2\lambda_k + \delta)^2} \\ &\leq \delta^{2\beta} \sum_{k \leq k^*} \lambda_k + 4 \sum_{k > k^*} \lambda_k. \end{aligned}$$

Both terms tend to zero as $\delta \rightarrow 0$, because \mathcal{C} is trace class. For the second term b note that

$$\begin{aligned}\mathbb{E}^{\eta_0} \|b\|^2 &\leq 8\delta \sum_k \frac{\lambda_k^2}{(2\lambda_k + \delta)^2} \\ &\leq 4\delta \sum_k \frac{\lambda_k}{2\lambda_k + \delta} \\ &\leq 2\delta^\beta \sum_{k \leq k^*} \lambda_k + 4 \sum_{k > k^*} \lambda_k.\end{aligned}$$

Again, both terms tend to zero as $\delta \rightarrow 0$, because \mathcal{C} is trace class. Having obtained these bounds on the moments of $v - u$, the remainder of the proof proceeds as in the preconditioned case.

Finally we demonstrate the desired property of $R(u, v)$. The symbol κ denotes a constant independent of u and δ . Recall that, for both proposals,

$$R(u, v) = \Phi(u) - \Phi(v). \quad (7.3)$$

By Assumption 7.1 it follows that

$$\begin{aligned}\mathbb{E}^\eta |R| &= \mathbb{E}^\eta |\Phi(u) - \Phi(v)| \\ &= \mathbb{E}^\eta \left(|\Phi(u) - \Phi(v)| \mathbb{I}(\|u\|, \|v\| \leq r) \right) + \mathbb{E}^\eta \left(|\Phi(u) - \Phi(v)| \mathbb{I}(\|u\| \text{ or } \|v\| > r) \right) \\ &\leq K(r) \mathbb{E}^\eta \|u - v\| \\ &\quad + \left(\mathbb{E}^\eta |\Phi(u) - \Phi(v)|^{1+\epsilon} \right)^{\frac{1}{1+\epsilon}} \left(\mathbb{P}^\eta(\|u\| \text{ or } \|v\| > r) \right)^{\frac{\epsilon}{1+\epsilon}} \\ &\leq K(r) \mathbb{E}^\eta \|u - v\| \\ &\quad + \kappa \left(\mathbb{E}^\eta (1 + \|u\|^{(1+\epsilon)p} + \|v\|^{(1+\epsilon)p}) \right)^{\frac{1}{1+\epsilon}} \left(\frac{\mathbb{E}^\eta \|u\|^2 + \mathbb{E}^\eta \|v\|^2}{r^2} \right)^{\frac{\epsilon}{1+\epsilon}}.\end{aligned}$$

We have shown that all moments of u and v are bounded under η . and that $\mathbb{E}^\eta \|u - v\|^{2p} \rightarrow 0$ as $\delta \rightarrow 0$. Fix arbitrary $\iota > 0$. The proof is completed by taking r sufficiently large to make the second term in the bound for $\mathbb{E}^\eta |R|$ less than ι , and then choosing δ sufficiently small to make the first term less than ι . \square

7.3. Scaling Limits and Spectral Gaps. There are two basic theories which have been developed to explain the advantage of using the algorithms introduced here which are based on the function-space viewpoints. The first is to prove scaling limits of the algorithms, and the second is to establish spectral gaps. The use of scaling limits was pioneered for local-proposal Metropolis algorithms in the papers [50, 51, 52], and recently extended to the hybrid Monte Carlo method [5]. All of this work concerned i.i.d. target distributions, but recently it has been shown that the basic conclusions of the theory, relating to optimal scaling of proposal variance with dimension, and optimal acceptance probability, can be extended to the target measures of the form (2.1) which are central to this paper; see [37, 44]. These results show that standard MCMC method must be scaled with proposal variance (or time-step in the case of HMC) which is inversely proportional to a power of d_u , the discretization dimension, and that the number of steps required grows under mesh refinement. The papers [45, 6] demonstrate that judicious modifications of these standard algorithms, as described in this paper, lead to scaling limits *without* the need for scalings of proposal variance or

time-step which depend on dimension. These results indicate that the number of steps required is stable under mesh refinement, for these new methods, as demonstrated numerically in this paper. The second approach, namely the use of spectral gaps, offers the opportunity to further substantiate these ideas: in [27] it is shown that the pCN method has a dimension independent spectral gap, whilst a standard random walk which closely resembles it has spectral gap which shrinks with dimension. This method of analysis, via spectral gaps, will be useful for the analysis of many other MCMC algorithms arising in high dimensions.

8. Conclusions. We have demonstrated the following:

- that a wide range of applications lead naturally to problems defined via density with respect to a Gaussian random field prior, or variants on this structure;
- that designing MCMC methods on function space, and then discretizing, produces better insight into algorithm design than discretizing and then applying MCMC methods;
- in particular we have highlighted new random walk, Langevin and Hybrid Monte Carlo Metropolis-type methods, appropriate for problems where the posterior distribution has density with respect to a Gaussian prior, all of which can be implemented by means of small modifications of existing codes;
- we have applied these MCMC methods to a range of problems, demonstrating their efficacy in comparison with standard methods, and shown their flexibility with respect to the incorporation of standard ideas from MCMC technology such as Gibbs sampling and estimation of noise precision through conjugate Gamma priors;
- we have pointed to the emerging body of theoretical literature which substantiates the desirable properties of the algorithms we have highlighted here.

The ubiquity of Gaussian priors means that the technology that is described in this article is of immediate applicability to a wide range of applications. The generality of the philosophy that underlies our approach also suggests the possibility of numerous further developments. In particular, many existing algorithms can be modified to the function space setting that is shown to be so desirable here, when Gaussian priors underlie the desired target; and many similar ideas can be expected to emerge for the study of problems with non-Gaussian priors, such as arise in wavelet based non-parametric estimation.

REFERENCES

- [1] F. Abramovich, T. Sapatinas, and B.W. Silverman. Wavelet thresholding via a Bayesian approach. *J. R. Stat. Soc. B*, 60:725–749, 1998.
- [2] R.P. Adams, I. Murray, and D.J.C. Mackay. The Gaussian process density sampler. Technical report, University of Cambridge, 1993.
- [3] R. J. Adler. *The Geometry of Random Fields*. Classics in Applied Mathematics. SIAM, Philadelphia, PA, 2010.
- [4] A. Bennett. *Inverse Modeling of the ocean and Atmosphere*. Cambridge, 2002.
- [5] A. Beskos, N.S. Pillai, G.O. Roberts, J.M. Sanz-Serna and A.M. Stuart, *Optimal Tuning of Hybrid Monte-Carlo*, To appear Bernoulli 2012.
- [6] A. Beskos, F.J. Pinski, J.M. Sanz-Serna, and A.M. Stuart. Hybrid Monte-Carlo on Hilbert space. *Stoch. Proc. and Applications*. **121**(2011), 2201–2230.
- [7] A. Beskos, G. O. Roberts, and A. M. Stuart. Optimal scalings for local Metropolis-Hastings chains on non-product targets in high dimensions. *Ann. Appl. Prob.*, 19:863–898, 2009.
- [8] A. Beskos, G. O. Roberts, A. M. Stuart, and J. Voss. MCMC methods for diffusion bridges. *Stochastic Dynamics*, 8(3):319–350, Sep 2008.

- [9] A. Beskos and A. M. Stuart. Computational complexity of Metropolis-Hastings methods in high dimensions. In Pierre L'Ecuyer and Art B. Owen, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2008*. Springer-Verlag, 2010.
- [10] A. Beskos and A.M. Stuart. MCMC methods for sampling function space. In *Invited Lectures, Sixth International Congress on Industrial and Applied Mathematics, ICIAM07, Editors Rolf Jeltsch and Gerhard Wanner*, pages 337–364. European Mathematical Society, 2009.
- [11] C. J. Cotter. The variational particle-mesh method for matching curves. *J. Phys. A: Math. Theor.*, 41:344003, 2008.
- [12] C.J. Cotter, S.L. Cotter and F.-X. Vialard. *Bayesian data assimilation in shape registration*. Submitted, 2011.
- [13] S.L. Cotter. Applications of MCMC methods on function spaces. *PhD Thesis, University of Warwick*, 2010.
- [14] S. L. Cotter, M. Dashti, and A. M. Stuart. Variational Data Assimilation Using Targetted Random Walks. *International Journal for Numerical Methods in Fluids*, to appear 2011.
- [15] S.L. Cotter, M. Dashti, J.C. Robinson, and A.M. Stuart. Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse Problems*, 25:doi:10.1088/0266-5611/25/11/115008, 2009.
- [16] G. Da Prato and J. Zabczyk. *Stochastic Equations in Infinite Dimensions*, volume 44 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1992.
- [17] M. Dashti, S. Harris, and A.M. Stuart. 2010.
- [18] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc. B*, 39:1–38, 1977.
- [19] P. Diaconis. Bayesian numerical analysis. pages 163–175. springer-verlag, 1988.
- [20] S. Duane, A.D. Kennedy, B. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216–222, 1987.
- [21] J. Glaunes, A. Trounev, and L. Younes. Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 712–718, 2004.
- [22] M. Hairer, A. M. Stuart, and J. Voss. Analysis of SPDEs arising in path sampling, part II: The nonlinear case. *Annals of Applied Probability*, 17:1657–1706, 2007.
- [23] M. Hairer, A. M. Stuart, and J. Voss. Sampling conditioned diffusions. In *Trends in Stochastic Analysis*, volume 353 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 2009.
- [24] M. Hairer, A.M. Stuart, and J. Voss. Signal processing problems on function space: Bayesian formulation, stochastic PDEs and effective MCMC methods. In *Oxford Handbook of Nonlinear Filtering*, 2010.
- [25] M. Hairer, A.M. Stuart, J. Voss, and P. Wiberg. Analysis of SPDEs arising in path sampling. Part I: The Gaussian case. *Comm. Math. Sci.*, 3:587–603, 2005.
- [26] M. Hairer, A.M. Stuart and J. Voss, Signal processing problems on function space: Bayesian formulation, stochastic PDEs and effective MCMC methods. Appears in "The Oxford Handbook of Nonlinear Filtering", Editors D. Crisan and B. Rozovsky, Oxford University Press 2011, pages 833-873.
- [27] M. Hairer, A.M. Stuart and S. Vollmer, Spectral gaps for Metropolis-Hastings algorithms in infinite dimensions. <http://arxiv.org/abs/1112.1392>
- [28] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [29] S.E. Hills and A.F.M. Smith. Parameterization issues in Bayesian inference. In *Bayesian Statistics IV*. Oxford University Press, 1992.
- [30] N. Hjort, C. Holmes, P. Muller, and S.G. Walker. *Bayesian Nonparametrics*. Cambridge University Press, 2010.
- [31] A. Iserles *A First Course in the Numerical Analysis of Differential Equations*. Cambridge Texts in Applied Mathematics, 2004.
- [32] E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge, 2003.
- [33] L'Ecuyer, Pierre and Simard, Richard. TestU01: A C Library for Empirical Testing of Random Number Generators. *ACM Transactions on Mathematical Software*, 2007, Vol 33, No 4.
- [34] J.C. Lemm. *Bayesian Field Theory*. Johns Hopkins, 2003.
- [35] Jun S. Liu. *Monte Carlo strategies in scientific computing*. Springer Series in Statistics. Springer, 2001.
- [36] D. McLaughlin and L.R. Townley. A reassessment of the groundwater inverse problem. *Water Res. Research*, 32:1131–1161, 1996.
- [37] J.C. Mattingly, N. Pillai and A.M. Stuart, Diffusion limits of random walk Metropolis algorithms in high dimensions. To appear, *Ann. Appl. Prob.*

- [38] X-L. Meng and D. van Dyk. The EM algorithm – an old folk-song sung to a new fast tune. *J. Roy. Stat. Soc. B*, 59:511–567, 1997.
- [39] N. Metropolis, R.W. Rosenbluth, M.N. Teller, and E. Teller. Equations of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.
- [40] M. T. Miller and L. Younes. Group actions, homeomorphisms, and matching: A general framework. *International Journal of Computer Vision*, 41:61–84, 2001.
- [41] R.M. Neal. MCMC using Hamiltonian dynamics. Appears in “Handbook of Markov Chain Monte Carlo: Methods and Applications”, page 113, 2010. Chapman and Hall.
- [42] R.M. Neal. *Bayesian learning for Neural networks*. Springer-Verlag, 1996.
- [43] A. O’Hagan. Uncertainty analysis and other inference tools for complex computer codes (with discussion). In *Bayesian Statistics VI*, pages 503–524. Oxford University Press, 1999.
- [44] N.S.Pillai, A.M. Stuart and A.H. Thiery, Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions. To appear, *Ann. Appl. Prob.* <http://arxiv.org/abs/1103.0542>
- [45] N.S.Pillai, A.M. Stuart and A.H. Thiery, On the random walk Metropolis algorithm for Gaussian random field priors and gradient flow. <http://arxiv.org/abs/1108.1494>
- [46] Press, William H. and Teukolsky, Saul A. and Vetterling, William T. and Flannery, Brian P. *Numerical Recipes in C* Cambridge University Press, 2002.
- [47] J.O. Ramsay and B.W. Silverman. *Functional Data Analysis*. Springer, 2005.
- [48] D. Richtmyer and K.W. Morton. *Difference Methods for Initial Value Problems*. Wiley, 1967.
- [49] C.P. Robert and G.C. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag, 1999.
- [50] G.O. Roberts, A. Gelman, and W.R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Prob.*, 7:110–120, 1997.
- [51] G.O. Roberts and J. Rosenthal. Optimal scaling discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B*, 60:255–268, 1998.
- [52] G.O. Roberts and J. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16:351–367, 2001.
- [53] G.O. Roberts and O. Stramer. On inference for partially observed nonlinear diffusion models using the MetropolisHastings algorithm. *Biometrika*, 88:603–621, 2001.
- [54] G.O. Roberts and R.L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- [55] H. Rue and L. Held. *Gaussian Markov Random Fields: Theory and Applications*. Chapman and Hall, 2005.
- [56] A.F.M. Smith and G.O. Roberts. Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods. *J. Roy. Stat. Soc. B*, 55:3–23, 1993.
- [57] A.D. Sokal. Monte Carlo methods in statistical mechanics: foundations and new algorithms, Université Lausanne, Bâtiment des sciences de physique Troisième Cycle de la physique en Suisse romande, 1989.
- [58] M.L. Stein. *Interpolation of Spatial Data*. Springer-Verlag, 1999.
- [59] A.M. Stuart. Inverse problems: a Bayesian approach. *Acta Numerica*, 19, 2010.
- [60] A.M. Stuart, J. Voss, and P. Wiberg. Conditional path sampling of SDEs and the Langevin MCMC method. *Comm. Math. Sci.*, 2:685–697, 2004.
- [61] M.A. Tanner and W.H. Wong. The calculation of posterior densities by data augmentation. *J. Amer. Stat. Assoc.*, 82:528–540, 1987.
- [62] L. Tierney. A note on Metropolis-Hastings kernels for general state spaces. *Ann. Appl. Probab.*, 8(1):1–9, 1998.
- [63] M. Vaillant and J. Glaunes. Surface matching via currents. *IPMI*, 381–392, 2005.