THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A determinant-free method to simulate the parameters of large Gaussian fields

Citation for published version:

Ellam, L, Strathmann, H, Girolami, M & Murray, I 2017, 'A determinant-free method to simulate the parameters of large Gaussian fields', *Stat*, vol. 6, no. 1, pp. 271-281. <https://doi.org/10.1002/sta4.153>

Digital Object Identifier (DOI):

[10.1002/sta4.153](https://doi.org/10.1002/sta4.153)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Stat

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A determinant-free method to simulate the parameters of large Gaussian fields

Louis Ellam^{a*}, Heiko Strathmann^b, Mark Girolami^{ac}, Iain Murray^d

Received 00 Month 2012; Accepted 00 Month 2012

We propose a determinant-free approach for simulation-based Bayesian inference in high-dimensional Gaussian models. We introduce auxiliary variables with covariance equal to the inverse covariance of the model. The joint probability of the auxiliary model can be computed without evaluating determinants, which are often hard to compute in high dimensions. We develop a Markov chain Monte Carlo sampling scheme for the auxiliary model that requires no more than the application of inverse-matrix-square-roots and the solution of linear systems. These operations can be performed at large scales with rational approximations. We provide an empirical study on both synthetic and real-world data for sparse Gaussian processes and for large-scale Gaussian Markov random fields. Copyright © 2012 John Wiley & Sons, Ltd.

Keywords: Bayesian inference; Markov chain Monte Carlo (MCMC); Gaussian processes (GPs); Gaussian Markov random fields (GMRFs); Data augmentation; Rational approximations.

1. Introduction

Linear Gaussian models are one of the most fundamental tools for statistical analysis, with diverse applications such as spatial modelling, uncertainty quantification, dimensionality reduction, and time-series modelling, e.g. speech recognition or even analysing high-dimensional recordings of neural activity (Roweis & Ghahramani, 1999; Rosti, 2004; Macke et al., 2011). A set of n observations, $\mathbf{y}$, is assumed to be generated from a set of latent Gaussian variables, $\mathbf{x}$, which are subject to a linear transformation $\mathbf{A}$ and corrupted by independent Gaussian noise ϵ . Popular examples include non-parametric regression with Gaussian processes (GPs, Rasmussen & Williams, 2006) or modelling spatial phenomena using Gaussian Markov random fields (GMRFs, Rue & Held, 2005; Lindgren et al., 2011). GPs are often based on full covariance matrices, which means none of the variables are marginally independent. In contrast GMRFs directly encode conditional independence relationships via zero elements in the precision matrix, which as a result is often sparse.

^aDepartment of Mathematics, Imperial College London, London, SW7 2AZ, UK

^bGatsby Unit of Computational Neuroscience, University College London, London, W1T 4JG, UK

^cThe Alan Turing Institute, The British Library, London, NW1 2DB

^dSchool of Informatics, University of Edinburgh, Edinburgh, EH8 9AB

*Email: l.ellam16@imperial.ac.uk

We focus on Bayesian posterior inference for linear Gaussian models. It is possible to obtain closed form expressions for key quantities such as the marginal likelihood and posterior predictions. However, computing these quantities in practice is difficult for large problems. For n modelled variables, an evaluation of the (normalised or unnormalised) probability density function costs $\mathcal{O}(n^3)$ time and $\mathcal{O}(n^2)$ memory. These costs arise from the solution of a linear system and the evaluation of a matrix determinant. There is a large literature on computationally motivated approximations to the model itself, for example via sub-sampling or low-rank approximations (Quiñero-Candela & Rasmussen, 2005), however, exact Bayesian inference for large models remains a challenging problem.

We can compute some properties of models with sparse covariance or precision matrices using two families of numerical linear algebra methods. These methods exploit the fact that large sparse matrices can be stored, and their products with vectors can be computed. Krylov methods, such as the linear solver conjugate gradient and multi-shift variants (Simpson, 2008; Jegerlehner, 1996), can solve linear systems with just matrix-vector products. Rational approximations reduce the computation of matrix functions, such as the determinant, inverse, or square root, to solving a family of linear equations – with error guarantees that are straight-forward to control and can even be set to floating point precision (Higham & Lin, 2013; Kennedy, 2004). A combination of these methods has been successfully applied to sample from high-dimensional Gaussian distributions (Aune et al., 2013), and to perform maximum likelihood inference by estimating the log-determinant term (Aune et al., 2014). A number of numerical approximations are available for estimating the log-determinant term, for example, those discussed by Bai et al. (1996), Han et al. (2015), and Saibaba et al. (2016). Recently, an inversion-free approach for inferring point-estimates of covariance parameters of large-scale Gaussian distributions was proposed by Anitescu et al. (2016). The approach is based on an alternative objective function than maximum likelihood, while asymptotically converging to the same solution.

Bayesian inference of the covariance parameters, however, remains largely an open problem. Markov chain Monte Carlo (MCMC) methods are ‘asymptotically exact’, but most previous work has limited scalability as the methods need to compute determinants (e.g., Murray & Adams, 2010). Lyne et al. (2015) realised that the log-determinant estimator by Aune et al. (2014) is unbiased, and can therefore be combined with the ‘exact-approximate’ framework within a pseudo-marginal MCMC algorithm (Andrieu & Roberts, 2009) – as also used by Filippone & Girolami (2014). A complication is the need to transform an unbiased estimator of the log-determinant into an unbiased estimator of the determinant itself, which Lyne et al. (2015) achieved via random truncation (Russian roulette) of the infinite series expression of the exponential function. Apart from theoretical issues around the existence of positive unbiased estimators for Russian roulette (Jacob et al., 2015), unfortunately, for a real-world model of global ozone distributions (Lindgren et al., 2011), the combination of Russian roulette and pseudo-marginal MCMC turned out to be too fragile and the resulting Markov chain in practice did not converge reliably.

In this paper we introduce an alternative MCMC approach to perform asymptotically exact inference on large-scale Gaussian models. We introduce auxiliary variables so that the joint distribution of the model and auxiliary variables contains no determinants. In return for removing determinants we need to update the auxiliary variables, which can be performed with the application of inverse-matrix-square-roots. In the case that the inverse-matrix-square-root of the covariance is unknown, we use rational approximations in the spirit of Aune et al. (2013) to perform tuning-free updates. Our scheme is considerably simpler than the approach taken by Lyne et al. (2015), and by avoiding the pseudo-marginal framework there is need to tune the internal unbiased estimators. Our approach scales well to high-dimensional models, provided that the application of an inverse-matrix-square-root can be carried out reliably. In the case of a poorly-conditioned model, in which the inverse-matrix-square-root is unknown, the underlying Krylov method may converge slowly, or not at all to within the desired tolerance.

In the remainder of the paper, we give details on the linear Gaussian model itself and the computational challenges, here for the case of using MCMC for posterior inference. We introduce our auxiliary model, which avoids the need to

Det-free method to simulate large Gaussian fields

Stat

compute matrix determinants, and describe a sampling scheme for the resulting joint distribution – including some necessary background on rational approximations. Finally, we empirically study the method on some toy and real-world examples. For models with well-behaved covariance or precision matrices, our determinant-free method outperforms MCMC using standard Cholesky factorisations.

2. Background

We consider linear models of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}. \quad (1)$$

Here, observations $\mathbf{y} \in \mathbb{R}^n$ are a linear transformation $\mathbf{A} \in \mathbb{R}^{n \times m}$ of independent Gaussian latent variables,

$$\mathbf{x} \in \mathbb{R}^m \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta), \quad (2)$$

with independent Gaussian additive noise[†] $\boldsymbol{\epsilon} \in \mathbb{R}^n$ and covariance matrix $\tau^{-1}\mathbf{I}$. Since the model is linear and Gaussian, the marginal likelihood, $p(\mathbf{y}|\boldsymbol{\theta})$, is also Gaussian, i.e.

$$\mathbf{y}|\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{A}\boldsymbol{\mu}_\theta, \mathbf{S}_\theta) \quad (3)$$

with covariance $\mathbf{S}_\theta = \tau^{-1}\mathbf{I} + \mathbf{A}\boldsymbol{\Sigma}_\theta\mathbf{A}^\top$.

Two common examples with this set-up are finite-dimensional realizations of a Gaussian process (GP) and Gaussian Markov random field (GMRF) models, often with dense covariance and sparse precision matrices, respectively.

For these models to be of any practical use, it is necessary to determine suitable parameter values $\boldsymbol{\theta}$. In a Bayesian setting, we define a prior density $p(\boldsymbol{\theta})$ and combine it with the likelihood in (3). Bayes' theorem induces the posterior

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\theta})p(\boldsymbol{\theta}), \quad (4)$$

which is intractable for most non-trivial applications. While there are many approximate inference schemes to explore the posterior (4), e.g. variational methods, we here focus on simulation via Markov chain Monte Carlo (MCMC), which has the advantage of being asymptotically exact. MCMC methods seek to generate samples $\{\boldsymbol{\theta}^{(i)}\}_{i=1,2,\dots} \sim p(\boldsymbol{\theta}|\mathbf{y})$ to represent the posterior. These samples can be used to estimate the posterior expectation of arbitrary functions, such as the posterior mean. In this paper we will use the standard Metropolis–Hastings (MH) algorithm.

2.1. Determinants in high dimensions

To evaluate the likelihood in (3), the standard approach is to compute the Cholesky factorisation $\mathbf{S}_\theta = \mathbf{R}_\theta^\top \mathbf{R}_\theta$, where $\mathbf{R}_\theta$ is upper triangular. This factorisation allows any linear system of the form $\mathbf{S}_\theta^{-1}\mathbf{x}$ to be solved via cheap back-substitution, and provides the log-determinant as

$$\log |\mathbf{S}_\theta| = 2 \sum_{i=1}^n \log(\text{diag}(\mathbf{R}_\theta)_i).$$

For sparse high-dimensional matrices $\mathbf{S}_\theta$, however, this approach is unsuitable, as even storing the Cholesky factorisation is infeasible. Cholesky factors suffer from a so called *fill-in* effect and are generally not sparse. Currently standard computers struggle to hold Cholesky factors in memory from around $n=10,000$. Examples of large sparse covariance

[†]For simplicity we include τ in the parameters $\boldsymbol{\theta}$, although τ doesn't effect the covariance of the latents $\mathbf{x}$.

matrices are those of Gaussian processes with compactly supported covariance functions. The Cholesky factor of sparse precision matrices from GMRF models can also be used.

Unbiased estimates of log-determinants can be obtained with rational approximations and Krylov subspace methods (Aune et al., 2014), although at greater expense than solving linear systems. Using these approximations to construct an MCMC method is complicated (Lyne et al., 2015).

3. Methodology

We augment the state space of the linear Gaussian model by introducing auxiliary variables

$$\mathbf{z}|\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}_{\boldsymbol{\theta}}^{-1}). \quad (5)$$

The resulting joint posterior is then given by the product of the posterior $p(\boldsymbol{\theta}|\mathbf{y})$ in (4) and the distribution over the new auxiliary variables,

$$p(\boldsymbol{\theta}, \mathbf{z}|\mathbf{y}) \propto p(\boldsymbol{\theta}) e^{-\frac{1}{2}(\mathbf{y}-\mathbf{A}\boldsymbol{\mu}_{\boldsymbol{\theta}})^{\top} \mathbf{S}_{\boldsymbol{\theta}}^{-1}(\mathbf{y}-\mathbf{A}\boldsymbol{\mu}_{\boldsymbol{\theta}}) - \frac{1}{2}\mathbf{z}^{\top} \mathbf{S}_{\boldsymbol{\theta}} \mathbf{z}}. \quad (6)$$

The normalising terms, i.e. determinants, of the original model and the augmented variables cancel,

$$\frac{1}{\int \int p(\boldsymbol{\theta}|\mathbf{y}) p(\mathbf{z}|\boldsymbol{\theta}) d\mathbf{z} d\boldsymbol{\theta}} \propto \frac{1}{|\mathbf{S}_{\boldsymbol{\theta}}|^{\frac{1}{2}} |\mathbf{S}_{\boldsymbol{\theta}}|^{-\frac{1}{2}}} \propto 1,$$

and (6) is left with only the quadratic forms for the latent and auxiliary Gaussian variables. The original posterior in (4) is restored by marginalizing over $\mathbf{z}$,

$$p(\boldsymbol{\theta}|\mathbf{y}) = \int p(\boldsymbol{\theta}, \mathbf{z}|\mathbf{y}) d\mathbf{z}.$$

In other words: in order to obtain posterior samples, we can sample from the augmented distribution and subsequently discard the auxiliary variables. Our MCMC scheme alternates between two updates:

1. Update the auxiliary variables $\mathbf{z}$ keeping $\boldsymbol{\theta}$ fixed using an MCMC method for target density (5).
2. Update the parameters $\boldsymbol{\theta}$ keeping $\mathbf{z}$ fixed using an MCMC method with target proportional to (6).

We next specify how we implement these updates.

3.1. Model parameter updates

The model parameters can be updated with standard Metropolis–Hastings (MH) updates. We use a preliminary run to estimate the covariance of the posterior $p(\boldsymbol{\theta}|\mathbf{y})$. In our experiments we use a random walk proposal that is tuned to achieve an acceptance rate in the range of 20–40% (Rosenthal et al., 2011).

3.2. Auxiliary variable updates

The auxiliary variables $\mathbf{z}$ could potentially be updated with MH or Hamiltonian Monte Carlo methods. However, the tuning of these updates would be challenging. Instead we perform Gibbs updates, directly resampling the auxiliary $\mathbf{z}$ from its conditional distribution $\mathcal{N}(\mathbf{0}, \mathbf{S}_{\boldsymbol{\theta}}^{-1})$. In the case that the inverse-matrix-square-root of covariance is unknown, we can't use a Cholesky factorization to perform this update for large scale problems. Instead we employ a technique from numerical linear algebra, a combination of rational approximations and Krylov sub-space methods (Aune et al.,

Det-free method to simulate large Gaussian fields

Stat

2014; Kennedy, 2004; Higham & Lin, 2013; Jegerlehner, 1996). These methods, outlined in the next section, are able to compute the product of the square-root-inverse of a matrix with arbitrary vectors – only requiring matrix-vector products involving the matrix itself. In particular, we can sample from high-dimensional Gaussian distributions as long as we can quickly apply the covariance or precision matrix to a vector, and that the matrix is reasonably well-conditioned. We now distinguish the cases where we have direct access either to the sparse latent covariance matrix Σ_θ from equation (2), or to the precision $\mathbf{Q}_\theta = \Sigma_\theta^{-1}$.

Sparse covariance For the case where we have direct access to a sparse latent covariance matrix Σ_θ , we directly compute $\mathbf{S}_\theta^{-\frac{1}{2}} \mathbf{w}$ for a standard normal $\mathbf{w}$, i.e. we generate a sample with desired auxiliary covariance $\mathbf{S}_\theta^{-1}$, via only using matrix-vector products of the form $\mathbf{S}_\theta \mathbf{w} = \tau^{-1} \mathbf{w} + \mathbf{A} \Sigma_\theta \mathbf{A}^\top \mathbf{w}$, i.e.

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \mathbf{z} \leftarrow \mathbf{S}_\theta^{-1/2} \mathbf{w}. \quad (7)$$

Sparse precision For the case where we know a sparse precision $\mathbf{Q}_\theta$, we create a set of ‘fantasy observations’, $\tilde{\mathbf{y}}$ from the model. For that, we first sample the latent variables $\tilde{\mathbf{x}}$ from (2) with covariance $\mathbf{Q}_\theta^{-1}$ by computing $\mathbf{Q}_\theta^{-\frac{1}{2}} \mathbf{w}'$, using matrix-vector products of the form $\mathbf{Q}_\theta \mathbf{w}'$, where $\mathbf{w}'$ is standard Gaussian. We then generate $\tilde{\mathbf{y}} = \mathbf{A} \tilde{\mathbf{x}} + \tau^{-1} \mathbf{w}'$. Finally, we pre-multiply $\tilde{\mathbf{y}}$, which has covariance $\mathbf{S}_\theta$, with $\mathbf{S}_\theta^{-1}$, to get

$$\mathbf{S}_\theta^{-1} \tilde{\mathbf{y}} = \mathbf{S}_\theta^{-1} \mathbf{S}_\theta^{\frac{1}{2}} \mathbf{w}' = \mathbf{S}_\theta^{-\frac{1}{2}} \mathbf{w}',$$

which has the desired auxiliary variable covariance $\mathbf{S}_\theta^{-1}$. In practice, we employ the matrix inversion lemma:

$$\mathbf{S}_\theta^{-1} \tilde{\mathbf{y}} = \tau \tilde{\mathbf{y}} - \tau^2 \mathbf{A} (\mathbf{Q}_\theta + \tau \mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \tilde{\mathbf{y}},$$

which only requires a single additional linear solve in terms of matrix-vector products involving the known $\mathbf{Q}_\theta$. In summary, given a sparse precision, we compute

$$\begin{aligned} \mathbf{w}' &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \tau^{-1} \mathbf{I}), \\ \tilde{\mathbf{x}} &\leftarrow \mathbf{Q}_\theta^{-\frac{1}{2}} \mathbf{w}', \quad \tilde{\mathbf{y}} \leftarrow \mathbf{A} \tilde{\mathbf{x}} + \boldsymbol{\epsilon}, \\ \tilde{\mathbf{z}} &\leftarrow \tau \tilde{\mathbf{y}} - \tau^2 \mathbf{A} (\mathbf{Q}_\theta + \tau \mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \tilde{\mathbf{y}}. \end{aligned} \quad (8)$$

3.3. Rational approximations and Krylov methods

We now briefly review the methodology required to solve the large sparse linear systems in (8) and matrix-inverse-square-roots in (8). Crucially, this is done without ever storing dense matrices, and only requires computing matrix-vector products $\mathbf{S}_\theta \mathbf{w}$ and $\Sigma_\theta^{-1} \mathbf{w}$ respectively (we assume either the ability to compute $\mathbf{S}_\theta \mathbf{w}$ or that a sparse Σ_θ^{-1} is given). We mainly follow the approach taken by Aune et al. (2013). Rational approximations are used to construct matrix-vector products with matrix-inverse-square-roots, e.g. $\mathbf{S}_\theta^{-\frac{1}{2}} \mathbf{w}$, by solving a family of sparse linear systems. Krylov space methods may be used to solve linear systems using only matrix-vector products.

The key identity, the Cauchy integral formula (Kennedy, 2004; Clark, 2006), relates a square matrix $\mathbf{A}$ and a function f that is analytic on and inside the closed contour Γ enclosing the Eigenvalues of $\mathbf{A}$ as

$$f(\mathbf{A}) = \frac{1}{2\pi i} \oint_\Gamma f(z) (z\mathbf{I} - \mathbf{A})^{-1} dz.$$

Right multiplying with vector $\mathbf{w}$ and applying a quadrature of the integral yields

$$f(\mathbf{A})\mathbf{w} \approx \sum_{i=1}^N \alpha_i (\mathbf{A} + \sigma_i \mathbf{I})^{-1} \mathbf{w}, \quad (9)$$

where we have to pick integration weights $\{\alpha_i\}_{i=1}^N$ and shifts $\{\sigma_i\}_{i=1}^N$ in order to ensure convergence. For positive definite $\mathbf{A}$ and $f(\mathbf{A}) = \mathbf{A}^{-\frac{1}{2}}$ and standard normal $\mathbf{w}$, Equation (9) can be used to sample from a multivariate Gaussian with covariance $\mathbf{A}^{-1}$.

Integration weights and shifts in (9) can be efficiently computed from the contour Γ , i.e. the smallest and largest Eigenvalue of $\mathbf{A}$, m and M respectively, in a standard way, e.g. Aune et al. (2014). To our knowledge, this is the best known rational approximation (Hale et al., 2008). Crucially, the Frobenius norm or the error in estimating $f(\mathbf{A})$ decays rapidly as

$$\mathcal{O} \left(\exp \left(\frac{-2\pi^2 N}{\ln(M/m) + 3} \right) \right),$$

and we can choose the number of quadrature points N to reach a desired accuracy. In practice, the contour Γ should be chosen to enclose the smallest region that contains the spectral range of $\mathbf{A}$ for optimal efficiency. We will see in the experiments, however, that we only need $N \approx 20$ quadrature points to reach floating point precision, e.g. 10^{-15} when using $m=10^{-6}$ and $M=10^6$. Due to space constraints, we refer to the literature for further details.

Sparse (shifted family) linear systems The conjugate gradient algorithm can solve the N linear systems required in (9), only requiring sparse matrix-vector products. Conjugate gradient is guaranteed to converge after n iterations for an n -dimensional system, where each matrix multiplication costs $\mathcal{O}(n^2)$. Depending on the condition number of the underlying matrix, however, convergence up to a reasonable tolerance can happen at a fraction of n , and a preconditioning method can improve convergence rates (Benzi, 2002; Chow & Saad, 2014). The N linear equation systems in (9) exhibit a special structure: only the diagonal term differs, arising from the various shifts $\{\sigma_i\}_{i=1}^N$. Shifted family Krylov methods can solve these systems simultaneously at the cost of a *single* solve of an unshifted system, i.e. $(\mathbf{A} + 0 \cdot \mathbf{I})^{-1} \mathbf{w}$ (Freund, 1993; Clark, 2006; Aune et al., 2013; Jegerlehner, 1996). Alternatively, depending on the conditioning of $\mathbf{A}$, it might be preferable to solve all N system separately, each with a different preconditioning approach (Simpson, 2008).

4. Experiments

In this section we apply our methodology to selected linear Gaussian models[‡]. We begin with a simple toy model with a random precision matrix. Our proposed methodology runs faster than a Cholesky based approach for large models, and scales to model sizes where Cholesky factors are not practical at all due to their memory costs.

4.1. Random pattern precision matrices

Consider the following model

$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, (\gamma^{-1} \mathbf{Q} + \gamma \mathbf{I})^{-1}),$$

where $\mathbf{Q}$ is a random, symmetric and positive definite matrix, generated by first generating a random Jacobi rotation of a positive definite diagonal matrix with elements in $[-1/2, 1/2]$. By adding elements to the diagonal, the precision

[‡]Code used for the results in this section is available at https://github.com/lellam/det_free_method.

Det-free method to simulate large Gaussian fields

$\gamma^{-1}\mathbf{Q} + \gamma\mathbf{I}$ is diagonally dominant and therefore its condition number is bounded. Each random precision matrix has $\sim 3n$ non-zero elements; an example is illustrated in Figure 1. We use $\ln \gamma = -3$ to generate the data for this example.

To perform inference, we use a log-uniform prior on γ . We run chains of length 10,000 for each example and tune the random walk proposal to obtain an acceptance rate of 20–40%. We initialize the chain near its true value for this example. We use 20 terms in the rational approximation and run the shifted family conjugate gradient solver (Simpson, 2008; Jegerlehner, 1996) to solve all linear systems in (9) in a single run. The covariance matrices in this example are well-conditioned, so the conjugate gradient method converges rapidly without the help of a preconditioner.

We compare our proposed determinant-free method to a standard random walk MH sampler where we evaluate the model likelihood using a Cholesky factorisation. As the table in Figure 1 shows, the Cholesky based approach produces more independent samples for a fixed number of MCMC iterations, as measured by effective sample size (ESS). It is expected that dependencies between the model parameters and auxiliary variables will cause the Markov chain to mix slower. Computing Cholesky factorisations requires more time per iteration and scales worse with problem size, which means that our scheme produces more effective samples per unit time. For $n \geq 10^5$ a standard laptop doesn't have enough memory to run the Cholesky based scheme at all, while our method scales to much larger problems.

Matrix size	10^3	10^4	10^5	10^6
Standard MCMC				
ESS	1768	1529	NA	NA
ESS/time	6.9009	0.0362	NA	NA
Our method				
ESS	883	543	613	1271
ESS/time	7.6241	0.9460	0.3159	0.0558

Figure 1 & Table 1. Left: Sparsity pattern of a well conditioned randomly generated $10^3 \times 10^3$ matrix. **Right:** Comparison of MCMC efficiency for $\ln \gamma = -3$ for a number of different sized matrices.

4.2. Sparse covariances for spatial modelling of anti-social crime data

We begin by applying our method to model the spatial distribution of anti-social criminal activity in London, using count data obtained from the UK government website. Log-crime-rates were calculated for each Lower Layer Super Output Area (LSOA), defined as number of crimes divided by LSOA population, and each crime-rate was assigned to the central location of the LSOA. After pre-processing, our dataset consists of 4,826 log-crime-rates, one for each region. LSOA data was obtained from the Office of National Statistics. Figure 2 (top) illustrates the dataset.

The classical geostatistical model (Gelfand et al., 2010), decomposes observations (here one-dimensional) of a spatial stochastic process at locations $\mathbf{s} \subseteq \mathbb{R}^2$ as $y(\mathbf{s}) = \mu(\mathbf{s}) + \eta(\mathbf{s}) + \epsilon(\mathbf{s})$, where $\mu(\mathbf{s})$ is a deterministic mean function, $\eta(\mathbf{s})$ is a continuous zero-mean stochastic process and $\epsilon(\mathbf{s})$ is white noise.

Our choice of mean function $\mu(\mathbf{s})$ is $\mu(\mathbf{s}) = \beta_0 + \sum_{h=1}^5 \beta_h \exp \left\{ -\frac{1}{2\sigma_h^2} \|\mathbf{s} - \mathbf{s}_h\|_2^2 \right\}$, where the constant β_0 represents the background crime-rate and the radial basis functions capture the general trend of an increase in crime in built-up areas. The five coefficients $\{\beta_h\}_{h=1}^5$ are held fixed at their maximum-likelihood values, after fitting a linear regression model. The radial basis functions are centred at $\mathbf{s}_h$ found using the k-means algorithm and the scaling σ_h is set equal

Figure 2. Top: Crime dataset showing log-crime-rate for reported anti-social crimes. **Left:** Comparison of MCMC efficiency of standard MCMC and the proposed method.

to the smallest pairwise distance between the s_h . Centering the GP around the linear combination of basis functions focusses the model's attention on deviations from the general trend. The stochastic process $\eta(s)$ is expected to capture localised crime 'hot-spots'. Those tend to only influence surrounding neighbourhoods – a phenomenon that can be appropriately modelled with a compactly supported Wendland kernel (Rasmussen & Williams, 2006):

$$k_{\theta}(x_i, x_j) = \sigma_f^2 \left(1 - \frac{\|x_i - x_j\|_2}{l}\right)_+^4 \left(\frac{4\|x_i - x_j\|_2}{l} + 1\right).$$

We use independent weakly informative log-normal priors on $p(\tau)$, $p(s)$ and $p(l)$. The posterior over all unknown parameters $\theta = (\tau, s, l)$ is then

$$p(\tau, s, l | y) \propto p(y | \theta) p(\tau) p(s) p(l).$$

We explore the posterior using a standard random walk MH both with determinant computations and our proposed augmented scheme. We run chains of length 10,000 with a burn-in of 3,000. The results and predictions are summarized in Figure 2. Due to the relatively small size of the dataset, the standard MCMC approach here produces slightly more independent samples per unit time – while our method runs significantly faster, it mixes slower due to the augmented sampling space.

We expect that our determinant-free method will provide more effective samples per unit time on larger systems, with a greater spatial extent. To test that idea we generated $n=10^4$ data locations from the uniform distribution over $[-0.5, 0.3] \times [51.25, 51.75]$ and sampled from the predictive distribution with the covariance parameters set to their posterior means. We ran the above procedure for a chain of length 1,000 and report the results for the model parameter τ . The ESS/time measured in $s^{-1} \times 10^{-4}$, for standard MCMC and our method was 9.3 and 78.4 respectively. Our method produced an order of magnitude more effective samples per unit time.

4.3. Gaussian Markov random field models specified by a whitening matrix

We now present a case where the model's precision matrix is known and is specified by a whitening matrix $\mathbf{L}_\theta$

$$\mathbf{L}_\theta \mathbf{x} = \mathbf{w}, \quad \mathbf{w} \sim N(\mathbf{0}, \mathbf{I}), \quad (10)$$

$$\mathbf{Q}_\theta = \mathbf{L}_\theta^\top \mathbf{L}_\theta, \quad (11)$$

This general setting is applicable for several models of interest, such as stochastic partial differential equation models discussed in Kaipio & Somersalo (2006); Lindgren et al. (2011). In this case, the GMRF has sparse precision $\mathbf{Q}_\theta = \Sigma_\theta^{-1}$, and the latent process in (2) is $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \mathbf{Q}_\theta^{-1})$. We may simulate from the centered latent process by drawing white noise and solving the sparse linear system in (10); the Krylov methods that raise concerns over conditioning are not required in this case. The covariance of the marginal likelihood in (3) is $\mathbf{S}_\theta = \tau^{-1} \mathbf{I} + \mathbf{A} \mathbf{Q}_\theta^{-1} \mathbf{A}^\top$, which is not necessarily sparse. To evaluate the marginal likelihood without having to work with non-sparse matrices we can use the matrix inversion lemma:

$$\ln p(\mathbf{y}|\theta) = \frac{1}{2} \left\{ \ln |\mathbf{Q}_\theta| + n \ln \tau - \ln |\mathbf{Q}_\theta + \tau \mathbf{A}^\top \mathbf{A}| - \tau \mathbf{y}^\top \mathbf{y} + \tau^2 \mathbf{y}^\top \mathbf{A} (\mathbf{Q}_\theta + \tau \mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{y} \right\} + \text{const}. \quad (12)$$

For moderate scale models, the log-determinants in (12) can be evaluated with 2 Cholesky decompositions.

In this example, specify the precision by a smooth latent process: $\mathbf{L}_\theta = \frac{1}{\gamma} \mathbf{L}_D$, where $\mathbf{L}_D$ is a discrete approximation of the Laplacian with Dirichlet boundary conditions generated from the mask

$$\begin{bmatrix} & -1 & \\ -1 & 4 & -1 \\ & -1 & \end{bmatrix}.$$

We define the latent process over a uniform quadrilateral mesh and consider the cases with 120×120 nodes and 200×200 nodes, so that $m = 14,400$ and $m = 40,000$ for each case, respectively. We generate synthetic datasets of size n at locations drawn uniformly over $[0, 1]^2$. The matrix $\mathbf{A}$ performs linear interpolation to estimate the latent process at the observation locations. We use the values $\ln \tau = -2$ and $\ln \gamma = -1$ to generate the true latent process in both cases and the latent process for the 120×120 case is presented in Figure 3. We perform inference on the marginal posterior for γ and τ using log-uniform priors for each.

Our preliminary runs revealed that $\mathbf{Q}_\theta$ is poorly conditioned, e.g. the 120×120 example had a condition number of 5.57×10^7 , which is expected to increase with m . Unlike for models specified by covariance, Krylov methods must be performed directly on the latent process and the observation noise does not help with conditioning as it does for standard GP models. As a result, Krylov methods are expected to require a large number of iterations. In this setting, we make use of the sparse linear system in (10), which can be solved using a banded linear system solver.

We again explore the posterior using a standard random walk MH: first with determinant computations, using 2 Cholesky decompositions per iteration; and then with our proposed augmented scheme. We run chains of length 10,000 and initialize the parameters close to their true values. The results in Figure 3 show that, for the both examples, our method out-performs the standard MCMC approach in cost per iteration. For large scale problems, our method can outperform standard MCMC in ESS/time, despite the slower mixing caused by adding a larger number of auxiliary variables to the model.

5. Discussion

We have shown how to introduce auxiliary random variables to replace the determinant computation arising in Gaussian models whose covariance is specified by unknown model parameters. The Markov chains for our method mix more

Grid	120 × 120		200 × 200	
Case	$\ln \tau$	$\ln \gamma$	$\ln \tau$	$\ln \gamma$
Std MCMC				
Mean	-1.99	-1.05	-1.98	-0.99
SD	0.008	0.050	0.011	0.037
ESS	1078	1373	796	1567
ESS/time	0.387	0.493	0.053	0.105
Det-free MCMC				
Mean	-1.99	-1.05	-1.98	-0.99
SD	0.012	0.068	0.015	0.053
ESS	512	636	414	786
ESS/time	0.238	0.295	0.067	0.128

Figure 3 & Table 2. Left: The true latent process for the 120×120 grid. **Right:** Comparison of the efficiency of standard MCMC and the proposed method for 120×120 grid (with 15,000 observations) and 200×200 grid (with 10,000 observations). Observations were drawn at random locations by interpolating the latent process and adding scaled white noise.

slowly per iteration than a Cholesky-based approach, due to the additional auxiliary variables, but can be much cheaper per iteration. Our method can be fast because it exploits fast matrix-vector operations, such as for models specified by sparse matrices, and it never creates dense matrices. These properties are particularly beneficial when the Cholesky decomposition is prohibitively expensive due to the fill-in, or where the Cholesky decomposition is expensive to obtain.

In practice, latent functions require a degree of smoothness, which forces the smallest eigenvalue in a large system close to zero and can result in a poorly conditioned system. The irony is that a system being poorly conditioned implies that it is almost low rank. Low rank systems are more constrained than an arbitrary process, and so should be less expensive to work with. In our crime example we exploited the low rank structure and worked with a reasonably well conditioned covariance matrix. However, Krylov methods do not work as well with extremely poorly-conditioned precision matrices as found in some GMRFs. In future work we will explore ways to exploit low rank structure more generally, so that our auxiliary variable scheme can focus on exploring the degrees of freedom of the process that have significant posterior uncertainty. We are also exploring applications of the methodology to inverse problems arising in partial differential equation models.

Acknowledgement

L.E. was supported by the EPSRC. H.S. was supported by the Gatsby Charitable Foundation. M.G. was supported by EPSRC [EP/J016934/1, EP/K034154/1, EP/P020720/1], an EPSRC Established Career Fellowship.

References

- Andrieu, C & Roberts, G (2009), 'The pseudo-marginal approach for efficient Monte Carlo computations,' *The Annals of Statistics*, **37**(2), pp. 697–725.
- Anitescu, M, Chen, J & Stein, ML (2016), 'An inversion-free estimating equations approach for Gaussian process models,' *Journal of Computational and Graphical Statistics*, **to appear**.

Det-free method to simulate large Gaussian fields

- Aune, E, Eidsvik, J & Pokern, Y (2013), 'Iterative numerical methods for sampling from high dimensional Gaussian distributions,' *Statistics and Computing*, **23**(4), pp. 501–521.
- Aune, E, Simpson, DP & Eidsvik, J (2014), 'Parameter estimation in high dimensional Gaussian distributions,' *Statistics and Computing*, **24**(2), pp. 247–263.
- Bai, Z, Fahey, G & Golub, G (1996), 'Some large-scale matrix computation problems,' *Journal of Computational and Applied Mathematics*, **74**(1), pp. 71–89.
- Benzi, M (2002), 'Preconditioning techniques for large linear systems: a survey,' *Journal of Computational Physics*, **182**(2), pp. 418–477.
- Chow, E & Saad, Y (2014), 'Preconditioned krylov subspace methods for sampling multivariate gaussian distributions,' *SIAM Journal on Scientific Computing*, **36**(2), pp. A588–A608.
- Clark, MA (2006), 'The rational hybrid Monte Carlo algorithm,' in *Proceedings from Lattice*, arXiv:hep-lat/0610048.
- Filippone, M & Girolami, M (2014), 'Pseudo-marginal Bayesian inference for Gaussian processes,' *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **36**(11), pp. 2214–2226.
- Freund, RW (1993), 'Solution of shifted linear systems by quasi-minimal residual iterations,' *Numerical Linear Algebra*, pp. 101–121.
- Gelfand, AE, Diggle, P, Guttorp, P & Fuentes, M (2010), *Handbook of Spatial Statistics*, CRC press.
- Hale, N, Higham, NJ & Trefethen, LN (2008), 'Computing A^α , $\log(A)$ and related matrix functions by contour integrals,' *SIAM Journal on Numerical Analysis*, **46**(5), pp. 2505–2523.
- Han, I, Malioutov, D & Shin, J (2015), 'Large-scale log-determinant computation through stochastic Chebyshev expansions,' *arXiv preprint arXiv:1503.06394*.
- Higham, NJ & Lin, L (2013), 'Matrix functions: A short course,' *Matrix Functions and Matrix Equations*, **19**, pp. 1–27.
- Jacob, PE, Thiery, AH et al. (2015), 'On nonnegative unbiased estimators,' *The Annals of Statistics*, **43**(2), pp. 769–784.
- Jegerlehner, B (1996), 'Krylov space solvers for shifted linear systems,' ArXiv:hep-lat/9612014.
- Kaipio, J & Somersalo, E (2006), *Statistical and computational inverse problems*, vol. 160, Springer Science & Business Media.
- Kennedy, A (2004), 'Approximation theory for matrices,' *Nuclear Physics B-Proceedings Supplements*, **128**, pp. 107–116.
- Lindgren, F, Rue, H & Lindström, J (2011), 'An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach,' *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **73**(4), pp. 423–498.
- Lyne, A, Girolami, M, Atchade, Y, Strathmann, H & Simpson, D (2015), 'On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods,' *Statistical Science*, **30**(4), pp. 443–467.
- Macke, JH, Buesing, L, Cunningham, JP, Byron, MY, Shenoy, KV & Sahani, M (2011), 'Empirical models of spiking in neural populations,' in *Advances in Neural Information Processing Systems*, vol. 24, pp. 1350–1358.
- Murray, I & Adams, RP (2010), 'Slice sampling covariance hyperparameters of latent Gaussian models,' in *Advances in Neural Information Processing Systems*, vol. 23, pp. 1723–1731.

- Quiñonero-Candela, J & Rasmussen, CE (2005), 'A unifying view of sparse approximate Gaussian process regression,' *Journal of Machine Learning Research*, **6**(Dec), pp. 1939–1959.
- Rasmussen, CE & Williams, C (2006), *Gaussian processes for machine learning*, MIT Press.
- Rosenthal, JS et al. (2011), 'Optimal proposal distributions and adaptive MCMC,' *Handbook of Markov Chain Monte Carlo*, pp. 93–112.
- Rosti, AVI (2004), *Linear Gaussian models for speech recognition*, Ph.D. thesis, University of Cambridge.
- Roweis, S & Ghahramani, Z (1999), 'A unifying review of linear Gaussian models,' *Neural Computation*, **11**(2), pp. 305–345.
- Rue, H & Held, L (2005), *Gaussian Markov random fields: theory and applications*, CRC Press.
- Saibaba, AK, Alexanderian, A & Ipsen, IC (2016), 'Randomized Matrix-free Trace and Log-Determinant Estimators,' *arXiv preprint arXiv:1605.04893*.
- Simpson, DP (2008), *Krylov subspace methods for approximating functions of symmetric positive definite matrices with applications to applied statistics and anomalous diffusion*, Ph.D. thesis, Queensland University of Technology.