Published in final edited form as:

J Stat Comput Simul. 2017 ; 87(11): 2227–2252. doi:10.1080/00949655.2017.1326117.

Bayesian Computation for Log-Gaussian Cox Processes: A Comparative Analysis of Methods

Ming Teng,

Department of Biostatistics, University of Michigan

Farouk S. Nathoo, and

Department of Mathematics and Statistics, University of Victoria

Timothy D. Johnson*

Department of Biostatistics, University of Michigan

Abstract

The Log-Gaussian Cox Process is a commonly used model for the analysis of spatial point pattern data. Fitting this model is difficult because of its doubly-stochastic property, i.e., it is an hierarchical combination of a Poisson process at the first level and a Gaussian Process at the second level. Various methods have been proposed to estimate such a process, including traditional likelihood-based approaches as well as Bayesian methods. We focus here on Bayesian methods and several approaches that have been considered for model fitting within this framework, including Hamiltonian Monte Carlo, the Integrated nested Laplace approximation, and Variational Bayes. We consider these approaches and make comparisons with respect to statistical and computational efficiency. These comparisons are made through several simulation studies as well as through two applications, the first examining ecological data and the second involving neuroimaging data.

Keywords

Hamiltonian Monte Carlo; Integrated Nested Laplace Approximation; Log-Gaussian Cox Process; Variational Bayes

1 Introduction

Spatial point process models for point pattern data have many applications including those involving ecology, geology, seismology, and neuroimaging. The theory of point processes has its roots with the work of Poisson in the 19th century, while more modern treatments with a statistical focus include [1], [2] and [3]. Among models for a spatial point process, the homogeneous Poisson process is the most fundamental but its use is limited in many applications due to its simplistic nature. A related but more flexible process is the Log-Gaussian Cox Process (LGCP), a process that is obtained by assuming a hierarchical structure, where at the first level the process is assumed Poisson conditional on the intensity

*Corresponding author tdjtdj@umich.edu.

function, and at the second level the log of the intensity function is assumed to be drawn from a Gaussian process. The flexibility of the model arises from the Gaussian process prior specified over the log-intensity function. Given this hierarchical structure with a Gaussian process at the second level, fitting this model to observed spatial point pattern data is a computational challenge [4].

A number of approaches have been developed to estimate the LGCP model in both the classical and Bayesian frameworks. In [5] the authors propose an adaption of Rosenblatt's kernel method [6] for the purpose of non-parametric estimation and then derive an expression for the mean squared error based on a stationarity assumption. A Bayesian framework is considered in [2] where the authors propose the use of the Metropolis-Adjusted Langevin Algorithm (MALA) [7] for Monte Carlo sampling of the posterior distribution. These authors also introduce a discretization of the spatial domain in order to attain computational tractability. [8] proposes an exact estimation method to deal with a modification of such a point process which they term the sigmoidal Gaussian Cox process. To avoid discretization of the spatial domain, [9] proposed a Markov chain Monte Carlo algorithm that samples from the joint posterior distribution of the LGCP model. A particular choice of the dominating measure is used to obtain the likelihood function without discretization and this is shown to be crucial to devise a valid MCMC algorithm. Although avoiding the error from discretization, the authors describe potential identifiability problems that require careful choice of prior distributions. In addition the computational complexity of the proposed MCMC algorithm could limit the application of this approach when considering very large data sets.

With regards to MCMC algorithms for Bayesian approaches, MALA is a special case of the potentially more efficient Hamiltonian Monte Carlo (HMC) [10] algorithm that uses the notion of Hamiltonian dynamics to construct proposals for the Metropolis-Hastings algorithm and has been adopted recently for an increasing number of applications [11]. Extending further, the Riemann Manifold Hamiltonian Monte Carlo algorithm proposed by [12] is a generalization of both MALA and HMC that can lead to more efficient posterior sampling in some cases. One drawback of Riemann Manifold HMC is that it requires the inversion of a potentially large matrix (the expected Fisher information matrix plus the negative Hessian of the log prior) at each iteration and this is not computationally feasible for very high-dimensional problems such as those typically involving the LGCP model.

An advantage associated with the use of MCMC algorithms for Bayesian computation is the underlying theory which guarantees simulation consistent estimation of various important characteristics of the posterior distribution. Thus the practitioner is assured of an accurate Monte Carlo representation of the posterior distribution given a sufficient amount of sampling effort. A drawback is that MCMC can be computationally intense, and this has motivated several alternative deterministic approaches for approximate Bayesian inference.

One such approach is Variational Bayes (VB) [13], where the approximation to the posterior distribution is assumed to lie within some convenient family of distributions and then an optimization is carried out to minimize the Kullback-Leibler divergence measuring the discrepancy between the true posterior and the approximation. Often the family of

distributions within which the approximation is assumed to lie is based on the notion of a mean field approximation, which corresponds to assuming posterior independence between certain model parameters. The idea in this case is to replace stochastic posterior dependence between parameters with deterministic dependence between the posterior moments in a manner that minimizes Kullback-Leibler divergence. Variational Bayes approximations have been applied successfully to the analysis of hidden Markov models in [13] and to other mixture models [14]. [15] used Variational Bayes for models of spatiotemporal systems represented by linear stochastic differential equations and demonstrated quick and efficient approximate inference both for continuous observations and point process data.

Mean field Variational Bayes is well suited for dealing with models within the conjugate exponential family where closed form solutions for the iterative steps of the optimization algorithm are available. In general such closed form solutions may not be available and additional approximations are then required. This is the case with the LGCP model. Mean field Variational Bayes approximations for non-conjugate models can be obtained by incorporating further approximations based on the delta method or the Laplace method [16]. These approximations are successfully applied to a correlated topic model and a Bayesian logistic regression model in [16], and to robust Bayesian models in [17]. For the LGCP model, tractable variational approximations can be obtained following this approach, where a mean field approximation with further approximations based on the Laplace method are used to handle the non-conjugate structure of the model. The variational approach considered in this paper can be thought of as an application of the techniques in [16] to the case of the LGCP model.

In [18] the authors approximate the posterior distribution of models incorporating a latent Gaussian process (which includes the LGCP model) using a mixture of Gaussian distributions, and derive a fixed-form variational approach for implementing this approximation. However, the proposed approach appears limited to one-dimensional point processes and is not applied to spatial point processes which is the focus of our work. [19] proposed a variational Bayes approach for fitting a Cox process assuming that the intensity function is the square of a Gaussian process. Unfortunately, this variational approach cannot be applied for estimation of the LGCP model, where the intensity is assumed to be the exponential of a Gaussian process, as the update steps for computing the proposed variational approximation to the posterior then become intractable.

Variational Bayes approximations can work well in some settings and the corresponding approximations can be computed relatively fast. A drawback is that there is no underlying theory guaranteeing the accuracy of the approximation or characterizing its error, thus these approximations need to be evaluated on a case-by-case basis, and they may or may not achieve reasonable accuracy depending on the utility of the practitioner. One contribution of this paper is the derivation of a mean field VB approximation which incorporates the Laplace method for the LGCP model. As far as we are aware such approximations have not been considered previously for this model. Another contribution of this paper is to compare this VB approximation with HMC in terms of both statistical and computational efficiency.

An alternative approach for approximate Bayesian inference that has gained tremendous popularity in the statistical literature is the integrated nested Laplace approximation (INLA) [20]. INLA is less generally applicable than MCMC or VB as it assumes the model has a latent Gaussian structure with only a moderate number of hyper-parameters. For spatial modeling the approach makes use of the Gaussian Markov Random field [21] and corresponding approximations which are known to be computationally efficient. The basis of INLA is the use of the Laplace approximation and numerical integration with latent Gaussian models to derive approximate posterior marginal distributions. INLA does not produce an approximation to the joint posterior which is a drawback of the approach in settings where the joint posterior (as opposed to the marginals) is of interest.

For spatial models incorporating a Gaussian Random field (GRF) with a Matérn correlation structure, [22] develop an approximate approach based on stochastic partial differential equations (SPDE) and these approximations have been combined for use with INLA. The essence of the approach is to specify a SPDE that has as its solution the GRF and then the SPDE representation is used in conjunction with basis representations to approximate the process over the vertices of a 2-dimensional mesh covering the spatial domain. The value of the process at any location is then obtained based on interpolation of the values at the mesh vertices. In recent work, [23] evaluated this approximation applied to the LGCP model with spatially varying covariates and demonstrated adequate performance for the settings and data considered there.

A comparison between INLA and MALA for models incorporating GMRF approximations is considered in [24]. In our work, we compare for the LGCP model Bayesian computation based on HMC, VB with a Laplace approximation, INLA, and INLA with the SPDE approximation. The comparisons we make are with respect to computational time, properties of estimators, posterior variability, and goodness fitness checking based on posterior predictive methods. Our objective is to provide practical guidance for users of the LGCP model. In addition to these comparisons, there are two novel aspects to the work presented here. First, we develop a mean field variational Bayes approximation that incorporates the Laplace method to deal with the non-conjugacy of the LGCP model. To the best of our knowledge this is the first time such an approximation has been developed for approximate Bayesian inference with the LGCP model. Second, we apply HMC for fully Bayesian inference and a novel aspect of our implementation is that HMC is used to update the decay (correlation parameter) associated with the latent Gaussian process. A result is that the sampling algorithm mixes very well and to our knowledge the development of the HMC algorithm in this context is the first of its kind.

The remainder of the paper proceeds as follows. Section 2 discusses various approaches for conducting Bayesian inference for the LGCP model. Section 3 presents a comparison of approaches through simulation studies, while Section 4 makes comparisons using two real point pattern data sets, the first arising from an ecological application and the second arising from a neuroimaging study. The paper concludes with a discussion in Section 5.

2 Bayesian inference for log-Gaussian Cox processes

2.1 Model specification

Consider an inhomogeneous Poisson process $\mathbf{Z}(s)$ with intensity function $\lambda(s)$, $s \in \mathcal{S} \subseteq \mathbb{R}^2$. Without loss of generality we shall assume that $\mathcal{S}$ is the unit square. The density of a Poisson process does not exist with respect to Lebesgue measure, but the Radon-Nikodym derivative does exist with respect to a unit-rate Poisson process [2]. We will call this derivative the density of the Poisson process. Given a set of K points $\{s_k\} = \{s_1, \dots, s_K\} \subset \mathcal{S}$, where both the number K and the locations s_k are random, the density is given by

$$\pi[\{s_k\}|\lambda(s)] = \exp \left\{ \int_{\mathcal{S}} [1 - \lambda(s)] \, ds \right\} \prod_{k=1}^K \lambda(s_k). \quad (1)$$

where $\lambda(s)$ is the intensity function. If we further assume that the log of the intensity function arises from a Gaussian random field (GRF) $\mathcal{Y}(s)$ so that $\lambda(s) = \exp(\mathcal{Y}(s))$, then this hierarchical process is called a log-Gaussian Cox process (LGCP) [2]. The LGCP, assumed to be stationary and isotropic, is uniquely determined by the mean function $\mu(s)$ and the covariance function $\text{Cov}(s, s') = \sigma^2 r(\|s - s'\|)$ of the Gaussian process $\mathcal{Y}(s)$, where σ^2 is the marginal variance and $r(\|s - s'\|)$ denotes correlation as a function of the Euclidean distance $\|s - s'\|$. Two commonly used correlation functions are the power exponential function [25]

$$r_p(\|s - s'\|) = \exp(-\rho \|s - s'\|^\delta)$$

where $\rho > 0$ is the decay parameter, $\delta \in (0, 2]$ is the power exponential term (which we will take as a known constant throughout this manuscript); and the Matérn correlation function [26]

$$r_m(\|s - s'\|) = (\Gamma(\nu) 2^{\nu-1})^{-1} (\|s - s'\|/\phi)^\nu K_\nu(\|s - s'\|/\phi)$$

where $\phi > 0$ is the range parameter, $\nu > 0$ is the shape parameter, and K_ν is the modified Bessel function of the second kind.

To fit the model in a tractable way a common approach is to divide the spatial domain into an $n \times n$ uniform grid of equally spaced cells [2] and to make the simplifying assumption that the log-intensity is constant over each grid cell so that the log-intensity $\mathcal{Y}(s)$ within a given cell, say the i^{th} cell, is constant and characterized by its value at the corresponding centroid, c_i , of cell i , $i \in \{1 \dots n^2\}$. The unique log-intensity values then comprise a vector $\mathbf{Y} = (\mathcal{Y}(c_1), \mathcal{Y}(c_2), \dots, \mathcal{Y}(c_{n^2}))^T$. To simplify notation we let $Y_i = \mathcal{Y}(c_i)$ and y_i is a realized value of Y_i . From the defining property of a GRF, $\mathbf{Y}$ follows a multivariate normal distribution $\mathbf{Y} \sim N(\mu \mathbf{1}_{n^2}, \sigma^2 \mathbf{C})$, where $\mathbf{C}$ is the $n^2 \times n^2$ correlation matrix with elements $r(\|c_i - c_j\|)$. Let θ be the set of parameters determining the mean and covariance of the GRF (e.g. $\theta = (\mu, \sigma^2, \rho)$ for the power correlation and $\theta = (\mu, \sigma^2, \phi)$ for the Matérn), and let A denote the area of each cell in the uniform grid. Under this discretization, the log density (see Equation (1)) is

$$\log \pi(\{s_k\}|\theta, \mathbf{y}) = \text{constant} + \sum_i [y_i n_i - A \exp(y_i)]$$

where n_i is the number of points in $\{s_k\}$ occurring in the i^{th} grid cell. The log posterior can then be expressed as

$$\begin{aligned} \log \pi(\theta, \mathbf{y}|\{s_k\}) = & \text{constant} + \sum_i [y_i n_i - A \exp(y_i)] \\ & - 0.5(\mathbf{y} - \mu \mathbf{1}_{n^2})^T \sigma^{-2} \mathbf{C}^{-1} (\mathbf{y} - \mu \mathbf{1}_{n^2}) \\ & - 0.5 n^2 \log(\sigma^2) - 0.5 \log(|\mathbf{C}|) + \log \pi(\theta) \end{aligned} \quad (2)$$

where $\pi(\theta)$ is the prior density of the parameter vector θ . The computational problem for Bayesian inference is the calculation of $\pi(\theta, \mathbf{y} | \{s_k\})$ and its associated marginals or properties of these distributions. This computation is nontrivial because the calculation of the normalizing constant is nontrivial, particularly when the dimension of the parameter space is high. We now discuss three approaches for approximating the posterior distribution and/or its marginals.

2.2 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) has its origins with the work of [27] and [28] and was first introduced into the statistical literature by [10]. It is a Metropolis-Hastings algorithm that can be used for sampling a high-dimensional target distribution more efficiently than an algorithm based on random-walk proposals, especially when the parameters are highly correlated. The algorithm uses the notion of the (separable) Hamiltonian $H(\mathbf{q}, \mathbf{p})$ from physics that is defined as the sum of potential energy $U(\mathbf{q})$ and kinetic energy $K(\mathbf{p})$, where $\mathbf{q}$ and $\mathbf{p}$ are random vectors that refer to position and momentum. The connection to Bayesian computation lies with relating $U(\mathbf{q})$ to the posterior distribution and hence $\mathbf{q}$ to the model parameters, and with introducing auxiliary Gaussian random variables to represent momentum $\mathbf{p}$, a vector having the same length as $\mathbf{q}$. The evolution of this system is then described by the Hamilton equations from statistical mechanics:

$$\frac{dq_i}{dt} = \frac{\partial H}{\partial p_i}$$

$$\frac{dp_i}{dt} = - \frac{\partial H}{\partial q_i} \quad (3)$$

which, if an analytic solution exists, produces a draw from the posterior distribution. In practice this system is solved using numerical integration techniques [11] and the resulting approximate solution is accepted or rejected using a Metropolis-Hastings step.

To carry out the computations required for the LGCP model we use a combination of re-parametrization of the random field and numerical techniques based on the 2D Fast Fourier transform (FFT) as in [2]. Note that although [12] has applied RM-HMC to LGCP model, their computation is slow due to inverting the fisher information matrix and a Cholesky factorization of the correlation matrix. While here we can greatly speedup matrix multiplications involving the correlation matrix $\mathbf{C}$ (we note here that we use the power exponential correlation function) by using FFT. A second reason to use the re-parametrization is that we avoid inversion on the correlation matrix at each iteration, which in our simulations and data analyses below is a 4096×4096 matrix. Although this size of a matrix can be inverted on a computer, it is computationally expensive. Furthermore, this FFT trick can handle much larger matrices that would be too large to invert on most computers. To use this trick we require the matrix to be block-circulant as there is a direct relationship between the eigenvalue-eigenvector decomposition of a block-circulant matrix and the 2D discrete Fourier transform. However, the correlation matrix has a block-Toeplitz structure. A block-Toeplitz matrix can always be extended to a block-circulant matrix [29]. To do so, we extend the original $n \times n$ grid to an $m \times m$ grid and wrap it on a torus, where $m = 2^g$ and g is an integer such that $m \geq 2(n - 1)$. The metric of this toroidal space is then defined by the minimum distance between two points. It is easy to show that the new correlation matrix $\mathbf{E}$ (of which $\mathbf{C}$ is a submatrix), whose elements are based on the metric defined on the torus, is a block-circulant matrix [2]. In extending the space we must also expand the vector of latent variables $\mathbf{Y}$ in a corresponding manner, and we refer to this new vector as $\mathbf{Y}^{ext}$ (of which $\mathbf{Y}$ is a subvector). Also, we set the number of points in cell i , m_i , still to be n_i if i is on the original grid and equal to 0 otherwise.

The block-circulant extended correlation matrix can be decomposed as $\mathbf{E} = \mathbf{F}\mathbf{\Lambda}\mathbf{F}^H$, where $\mathbf{F}$ is the matrix of eigenvectors, $\mathbf{\Lambda}$ is the diagonal matrix containing the corresponding eigenvalues of $\mathbf{E}$, and H denotes the complex conjugate transpose. Given a random vector $\mathbf{v}$ of length m^2 the product $\mathbf{E}\mathbf{v}$ can be obtained by calculating $\mathbf{F}^H\mathbf{v}$, $\mathbf{\Lambda}\mathbf{F}^H\mathbf{v}$ and then $\mathbf{F}\mathbf{\Lambda}\mathbf{F}^H\mathbf{v}$ in order. Note that the first and last calculations amount to a discrete inverse Fourier transform and a discrete Fourier transform (DFT), respectively. The middle calculation is simply element-wise multiplication of $\mathbf{\Lambda}$ and the vector $\mathbf{F}^H\mathbf{v}$ [21]. As a result, the complexity of the required matrix operations can be reduced to $O(m^2 \log(m^2))$ using the FFT.

After extension of the grid we re-parametrize the latent variables $\mathbf{Y}^{ext}$ as

$\mathbf{Y}^{ext} = \mu \mathbf{1}_{m^2} + \sigma \mathbf{E}^{\frac{1}{2}} \boldsymbol{\gamma}$ where $\mathbf{1}_{m^2}$ denotes the m^2 -dimensional vector of ones. $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_{m^2})^T$ with $\gamma_i \stackrel{\text{iid}}{\sim} N(0, 1)$, $i = 1, \dots, m^2$. The gradients, used in the HMC algorithm, for all the parameters are straightforward to derive except for ρ , we give the expression for ρ here and refer the readers to Appendix A for more details.

$$\frac{\partial \log \pi(\rho|\cdot)}{\partial \rho} = -\frac{\sigma}{2} \left[\mathbf{m} - A \exp(\mu \mathbf{1}_{m^2} + \sigma \mathbf{E}^{\frac{1}{2}} \boldsymbol{\gamma}) \right]^T \mathbf{E}^{-\frac{1}{2}} \mathbf{E}^* \boldsymbol{\gamma}, \quad (4)$$

where $\pi(\cdot | \cdot)$ denotes the full conditional given the data and other parameters, $\mathbf{m}$, $\mathbf{E}^*$ are defined in the appendix. And as $\mathbf{E}^{\frac{1}{2}}$, $\mathbf{E}^{-\frac{1}{2}}$, $\mathbf{E}^*$ are all block-circulant matrices, FFT can be used.

With the stochastic representation as in Equation 3 the HMC algorithm is based on setting $\mathbf{U}(\mathbf{q}) = -\log [\pi(\{s_k\} | \theta, \boldsymbol{\gamma}) \pi(\theta, \boldsymbol{\gamma})]$ where $\mathbf{q} = (\boldsymbol{\gamma}^T, \theta^T)^T$, and the kinetic energy term is $K(\mathbf{p}) = \mathbf{p}^T \mathbf{M}^{-1} \mathbf{p} / 2$, where $\mathbf{M}$ is a symmetric, positive-definite 'mass matrix' and auxiliary momentum variables $\mathbf{p}$ (a vector of length $m^2 + 3$). In our work we set $\mathbf{M}$ to be a diagonal matrix with distinct diagonal components $m_{\boldsymbol{\gamma}}$, m_{μ} , m_{σ^2} and m_{ρ} corresponding to $\boldsymbol{\gamma}$ and θ .

Each iteration of the HMC algorithm involves a block update of $\boldsymbol{\gamma}$ and θ based on the Hamiltonian Monte Carlo scheme. Each such update requires $L + 1$ evaluations of the gradient vector $\nabla_{\theta, \boldsymbol{\gamma}} \log \pi(\theta, \boldsymbol{\gamma} | \{s_k\})$ for some $L \geq 1$. If $L = 1$, the HMC algorithm reduces to MALA, which typically mixes faster than the random walk Metropolis-Hastings algorithm [30] but not as fast as the more general HMC algorithm. Letting $\boldsymbol{\gamma}^*$, θ^* be the current value in the Markov chain for $\boldsymbol{\gamma}$, θ , the HMC update, based on a step size $\epsilon > 0$, proceeds as follows: At each iteration, the number of steps in the numerical integration, L , is drawn from a Poisson distribution with mean 100, while the step size, ϵ , is initially chosen to be 0.005 and adjusted adaptively during burning so that the acceptance rate is approximately 0.65 [31]. Trace plots of the parameters are examined and based on these we adjust the values of $\mathbf{M}$ to improve the mixing.

2.3 Mean field Variational Bayes with Laplace Approximation

As an alternative to HMC (or MCMC) sampling of the posterior distribution, a deterministic approximation can be employed. Mean field variational Bayes (MFVB) is one such approximation that has been applied successfully to a number of problems, including spatial models for high-dimensional problems requiring fast computations [32]. For the MFVB algorithm, we return to the parametrization of the model in Equation (2). Let $q(\mathbf{y}, \theta)$ be an arbitrary density function having the same support as the posterior density $\pi(\theta, \mathbf{y} | \{s_k\})$. Letting $\log \pi(\{s_k\})$ denote the marginal likelihood of the model, we can express its logarithm as

$$\begin{aligned} \log \pi(\{s_k\}) &= \int q(\mathbf{y}, \theta) \log \left\{ \frac{\pi(\{s_k\}, \mathbf{y}, \theta)}{q(\mathbf{y}, \theta)} \right\} \\ &\quad + \int q(\mathbf{y}, \theta) \log \left\{ \frac{q(\mathbf{y}, \theta)}{\pi(\theta, \mathbf{y} | \{s_k\})} \right\} \\ &\geq \int q(\mathbf{y}, \theta) \log \left\{ \frac{\pi(\{s_k\}, \theta, \mathbf{y})}{q(\mathbf{y}, \theta)} \right\} \equiv F(q) \end{aligned}$$

such that the functional $F(q)$ is a lower bound for $\log \pi(\{s_k\})$ for any q . The approximation is obtained by restricting q to a manageable class of density functions, and maximizing F over that class. We develop the approximation under the assumption that the GRF has a power exponential correlation function, and for now, we will assume that ρ is known, so that C is assumed known in what follows.

Algorithm 1

HMC algorithm

1. Simulate latent vector $\mathbf{p}^* \sim N_{m^2+3}(\mathbf{0}, \mathbf{M})$. Set

$$(\gamma^{(0)}, \theta^{(0)}) = (\gamma^*, \theta^*)$$

$$\mathbf{p}^{(0)} = \mathbf{p}^* + \frac{\varepsilon}{2} \nabla_{\theta^*, \gamma^*} [\log \{\pi(\gamma^*, \theta^* | \{s_k\}) \pi(\theta^*, \gamma^*)\}].$$

2. For $l=1, \dots, L$,

$$(\gamma^{(l)}, \theta^{(l)})^T = (\gamma^{(l-1)}, \theta^{(l-1)})^T + \varepsilon \mathbf{M}^{-1} \mathbf{p}^{(l-1)}$$

$$\mathbf{p}^{(l)} = \mathbf{p}^{(l-1)} + \varepsilon_l \nabla_{\theta^{(l)}, \gamma^{(l)}} [\log \{\pi(\gamma^{(l)}, \theta^{(l)} | \{s_k\}) \pi(\theta^{(l)}, \gamma^{(l)})\}]$$

where $\varepsilon_l = \varepsilon$ for $l < L$ and $\varepsilon_L = \varepsilon/2$.

3. Accept $(\gamma^{(L)}, \theta^{(L)})$ as the new state for (γ, θ) with probability

$$\alpha = \min(1, \exp \{-H(\mathbf{q}^{(L)}, \mathbf{p}^{(L)}) + H(\mathbf{q}^*, \mathbf{p}^*)\})$$

else remain in the current state γ^*, θ^* with probability $1 - \alpha$.

Repeat steps 1–3 for a sufficiently long time.

We assume that the approximating density q can be factorized

$$q(\mathbf{y}, \theta) = \left[\prod_{i=1}^{n^2} q(y_i) \right] q(\mu) q(\sigma^2). \quad (5)$$

Under this assumption, a coordinate ascent algorithm is applied to maximize F which leads to a sequence of coordinate-wise updates taking the form

$$q(y_i) \propto \exp\{E_{-q(y_i)}[\log \pi(\{s_k\} | \mathbf{y}) \pi(\mathbf{y} | \theta)]\} \quad i=1, \dots, n^2$$

$$q(\mu) \propto \exp\{E_{-q(\mu)}[\log \pi(\mathbf{y} | \theta) \pi(\mu)]\}$$

$$q(\sigma^2) \propto \exp\{E_{-q(\sigma^2)}[\log \pi(\mathbf{y} | \theta) \pi(\sigma^2)]\}$$

where $E_{-q(x)}[\cdot]$ denotes the expectation taken with respect to the set of random variables $\{\mathbf{y}, \theta\} \setminus x$ under the variational approximation q_{-x} , and the updates steps are iterated to the convergence of F . We describe the derivation of the update steps below. In what follows, $E[\cdot]$ will denote the expectation of its argument under the variational approximation q .

If conditionally conjugate Gaussian and inverse-gamma priors are chosen for $\pi(\mu)$ and $\pi(\sigma^2)$ respectively, $\mu \sim N(\mu_\mu, \sigma_\mu^2)$, $\sigma^2 \sim G^{-1}(\alpha, \beta)$ the distributions $q(\mu)$ and $q(\sigma^2)$ comprising the update steps will also be Gaussian and inverse-gamma. To derive the update step for μ we have

$$\begin{aligned} E_{-q(\mu)} \log \pi(\mathbf{y}|\theta)\pi(\mu) &= c - \frac{1}{2}[(\mu_{\mathbf{y}} - \mu\mathbf{1})^T E(\sigma^{-2})(C^{-1})(\mu_{\mathbf{y}} - \mu\mathbf{1})] - \frac{1}{2}(\mu - \mu_\mu)^2 / \sigma_\mu^2 \\ &= c - \frac{1}{2}[(E(\sigma^{-2})\mathbf{1}^T(C^{-1})\mathbf{1} + 1/\sigma_\mu^2)\mu^2 - 2E(\sigma^{-2})\mu_{\mathbf{y}}^T(C^{-1})\mathbf{1}\mu + \mu_\mu / \sigma_\mu^2] \end{aligned} \quad (6)$$

where $\mu_{\mathbf{y}} = E[\mathbf{y}]$ and c denotes a constant not depending on μ . As (6) is a quadratic function of μ it follows that $q(\mu)$ is the density of a normal distribution $N(\mu_\mu^*, \sigma_\mu^{*2})$ where after some algebra we have

$$\mu_\mu^* = (E(\sigma^{-2})\mu_{\mathbf{y}}^T(C^{-1})\mathbf{1} + \mu_\mu / \sigma_\mu^2)(\sigma_\mu^{*2})^{-1}, \quad \sigma_\mu^{*2} = (E(\sigma^{-2})\mathbf{1}^T(C^{-1})\mathbf{1} + 1/\sigma_\mu^2)^{-1}.$$

To derive the update step for σ^2 we have

$$E_{-q(\sigma^2)} \log \pi(\mathbf{y}|\theta)\pi(\sigma^2) = c - \frac{n^2}{2} \log \sigma^2 - \frac{1}{2} \sigma^{-2} E_\mu[(\mu_{\mathbf{y}} - \mu\mathbf{1})^T C^{-1}(\mu_{\mathbf{y}} - \mu\mathbf{1}) + \text{Tr}(C^{-1}\Sigma_{\mathbf{y}})] - (\alpha+1) \log \sigma^2 - \beta / \sigma^2$$

where $\Sigma_{\mathbf{y}}$ is the covariance matrix of $\mathbf{y}$ under $q(\mathbf{y})$ and c denotes a constant not depending on σ^2 . Simplifying this expression yields

$$\begin{aligned} E_{-q(\sigma^2)} \log \pi(\mathbf{y}|\theta)\pi(\sigma^2) &= c - (\alpha+1 + \frac{n^2}{2}) \log \sigma^2 - (\beta + \frac{1}{2}[(\mu_{\mathbf{y}} - \mu_\mu^*\mathbf{1})^T(C^{-1})(\mu_{\mathbf{y}} - \mu_\mu^*\mathbf{1}) \\ &\quad + \sigma_\mu^{*2}\mathbf{1}^T(C^{-1})\mathbf{1} + \text{Tr}(C^{-1}\Sigma_{\mathbf{y}})]) / \sigma^2 \end{aligned}$$

and thus $q(\sigma^2)$ is the density of an Inverse-Gamma distribution $G^{-1}(\alpha + \frac{n^2}{2}, \beta^*)$ where

$$\beta^* = \beta + \frac{1}{2}[(\mu_{\mathbf{y}} - \mu_\mu^*\mathbf{1})^T(C^{-1})(\mu_{\mathbf{y}} - \mu_\mu^*\mathbf{1}) + \sigma_\mu^{*2}\mathbf{1}^T(C^{-1})\mathbf{1} + \text{Tr}(C^{-1}\Sigma_{\mathbf{y}})].$$

As the Gaussian prior for y_i is not conditionally conjugate for the LGCP model, the variational Bayes update for $q(y_i)$ is not a standard distribution and is therefore not easy to compute without some further approximation. We derive an update step for $q(y_i)$ by applying the Laplace method [16] within the variational Bayes update. We have

$$\begin{aligned} E_{-q(y_i)} \log \pi(\{s_k\}|\mathbf{y})\pi(\mathbf{y}|\theta) &= c + (y_i n_i - e^{y_i} A) \\ &\quad - \frac{1}{2}[(\tilde{\mathbf{y}} - E(\mu)\mathbf{1})^T E(\sigma^{-2})(C^{-1})(\tilde{\mathbf{y}} - E(\mu)\mathbf{1}) + \text{Var}(\mu)E(\sigma^{-2})\text{Tr}((C^{-1})\mathbf{1}\mathbf{1}^T)] \\ &= c + (y_i n_i - e^{y_i} A) - \frac{1}{2}[(\tilde{\mathbf{y}} - E(\mu)\mathbf{1})^T E(\sigma^{-2})(C^{-1})(\tilde{\mathbf{y}} - E(\mu)\mathbf{1})] \end{aligned} \quad (7)$$

where $\tilde{\mathbf{y}}$ denotes $(E(y_1), \dots, E(y_{i-1}), y_i, E(y_{i+1}), \dots, E(y_n))^T$. Taking the derivative with respect to y_i yields,

$$\begin{aligned} f(y_i) &= \partial_{E(y_i)} \log \pi(\{s_k\} | \mathbf{y}) \pi(\mathbf{y} | \theta) / \partial y_i = n_i - A e^{y_i} - E(\sigma^{-2}) [(C^{-1})(\tilde{\mathbf{y}} - E(\mu)\mathbf{1})]_i \\ &= n_i - A e^{y_i} - E(\sigma^{-2}) \sum_j (C^{-1})_{ij} (\tilde{\mathbf{y}} - E(\mu)\mathbf{1})_j \\ &= -A e^{y_i} - E(\sigma^{-2}) (C^{-1})_{ii} y_i + n_i + E(\sigma^{-2}) (C^{-1})_{ii} E(\mu) - E(\sigma^{-2}) \sum_{j \neq i} (C^{-1})_{ij} (\tilde{\mathbf{y}} - E(\mu)\mathbf{1})_j \end{aligned}$$

(8)

Given (8) we find $\hat{y}_i$ such that $f(\hat{y}_i) = 0$. We use Newton's method to obtain a numerical solution where the starting value for Newton's method is obtained by omitting the linear term in equation (8) and solving the resulting simplified equation exactly. We then take the second derivative with respect to y_i , $H(y_i) = -f(y_i)'' = -A e^{y_i} - E(\sigma^{-2}) (C^{-1})_{ii}$ and given this and the solution $\hat{y}_i$, the Laplace method yields a normal distribution $\mathcal{N}(\hat{y}_i, -H(\hat{y}_i)^{-1})$ for $q(y_i)$ which approximates the VB update. The variational-Laplace approximation to the posterior of the latent field $\mathbf{y}$ is then $\mathbf{y} \sim \mathcal{N}(\mu_{\mathbf{y}}, \Sigma_{\mathbf{y}})$ where $\mu_{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_n)^T$ and $\Sigma_{\mathbf{y}}$ is a diagonal matrix with the i^{th} element being $-H(\hat{y}_i)^{-1}$.

Given the update steps derived above, the approximate posterior distribution under the variational-Laplace approximation takes the form (5) where the component densities are standard distributions

$$\mathbf{y} \sim \mathcal{N}(\mu_{\mathbf{y}}^{(q)}, \Sigma_{\mathbf{y}}^{(q)})$$

$$\mu \sim \mathcal{N}(\mu_{\mu}^{(q)}, \sigma_{\mu}^{2(q)})$$

$$\sigma^2 \sim G^{-1}(\alpha^{(q)}, \beta^{(q)}).$$

The parameters determining these distributions are called the 'variational parameters'. These parameters are obtained through the sequence of update steps derived above which are used to determine equations expressing each variational parameter in terms of the remaining variational parameters; beginning with initial values for these parameters the equations are iterated to convergence of F .

With respect to the parameter ρ of the power exponential correlation function, we have found that its inclusion as an unknown parameter into the variational approach leads to convergence problems in a number of trial examples. To deal with this problem, we estimate this parameter prior to running the VB algorithm using the method of minimum contrast

[33], i.e., to use non-linear least squares estimation to fit a non-parametric estimated covariance function. The mean field VB algorithm incorporating the Laplace method for the LGCP model is presented in detail in Algorithm 2.

2.4 INLA

The integrated nested Laplace approximation (INLA) is another approach for constructing a deterministic approximation to the posterior distribution that can be applied to the fairly broad class of latent Gaussian models. The details underlying the approach have been described in a number of recent papers including the seminal work of [20]. We provide here only a brief overview of aspects that are relevant for use with the LGCP model.

Algorithm 2

Mean Field VB Algorithm with Laplace Method

1. Initialize the priors $\mu_\mu, \sigma_\mu^2, \alpha, \beta$.

2. Initialize $\mu_{\mathbf{y}}^{(q)}, \sum_{\mathbf{y}}^{(q)}, \mu_\mu^{(q)}, \sigma_\mu^{2(q)}, \beta^{(q)}$ and $E(\sigma^{-2}) = (\alpha + n^2/2)/(\beta^{(q)})$

3. Obtain ρ using the minimum contrast method, compute $\mathbf{C}^{-1}$ where $c_{kl} = \exp(-\rho \|c_k - c_l\|^\delta)$.

4. For $i = 1, \dots, n^2$, compute $\mu_{y_i}^{(q)}$ such that

$$-A \exp(\mu_{y_i}^{(q)}) - E(\sigma^{-2})(\mathbf{C}^{-1})_{ii} \mu_{y_i}^{(q)} + n_i + \mu_\mu^{(q)} E(\sigma^{-2})(\mathbf{C}^{-1})_{ii} - E(\sigma^{-2}) \sum_{j \neq i} (\mathbf{C}^{-1})_{ij} [\mu_{y_j}^{(q)} - \mu_\mu^{(q)} \mathbf{1}]_j = 0$$

where $[\cdot]_j$ denotes the j^{th} element of a vector.

Compute $H(\mu_{y_i}^{(q)}) = -A \exp(\mu_{y_i}^{(q)}) - E(\sigma^{-2})(\mathbf{C}^{-1})_{ii}$.

Obtain $\mu_{\mathbf{y}}^{(q)}$ and $\sum_{\mathbf{y}}^{(q)}$ where $\mu_{\mathbf{y}}^{(q)} = (\mu_{y_1}^{(q)}, \dots, \mu_{y_{n^2}}^{(q)})^T$ and $\sum_{\mathbf{y}}^{(q)}$ is a diagonal matrix with diagonal elements $-H(\mu_{y_i}^{(q)})^{-1}$.

5. Compute $\mu_\mu^{(q)} = (E(\sigma^{-2}) \mu_{\mathbf{y}}^{(q)T} \mathbf{C}^{-1} \mathbf{1} + \mu_\mu / \sigma_\mu^2) (E(\sigma^{-2}) \mathbf{1}^T \mathbf{C}^{-1} \mathbf{1} + 1 / \sigma_\mu^2)^{-1}$. Compute $\sigma_\mu^{2(q)} = (E(\sigma^{-2}) \mathbf{1}^T \mathbf{C}^{-1} \mathbf{1} + 1 / \sigma_\mu^2)^{-1}$.

6. Compute $\beta^{(q)} = \beta + 0.5 [(\mu_{\mathbf{y}}^{(q)} - \mu_\mu^{(q)} \mathbf{1})^T \mathbf{C}^{-1} (\mu_{\mathbf{y}}^{(q)} - \mu_\mu^{(q)} \mathbf{1}) + \sigma_\mu^{2(q)} \mathbf{1}^T \mathbf{C}^{-1} \mathbf{1} + \text{Tr}(\mathbf{C}^{-1} \sum_{\mathbf{y}}^{(q)})]$.

Obtain $E(\sigma^{-2}) = (\alpha + n^2/2)/(\beta^{(q)})$

7. Compute the lower bound

$$F(q) = \sum_i (\mu_{y_i}^{(q)} n_i - A \exp(\mu_{y_i}^{(q)} - 0.5 H(\mu_{y_i}^{(q)})^{-1})) - 0.5 E \log(|\mathbf{C}|) - 0.5 [(\mu_\mu^{(q)} - \mu_\mu)^2 + \sigma_\mu^{2(q)}] / \sigma_\mu^2 + \sum_i 0.5 \log(-H(\mu_{y_i}^{(q)})^{-1}) + 0.5 \log \sigma_\mu^{2(q)}$$

Repeat 4–7 until the increase in $F(q)$ is negligible.

For spatial models, INLA makes extensive use of the Gaussian Markov random field (GMRF) which is a Gaussian distribution having a sparse precision matrix. Algorithms for

fitting models incorporating a GMRF can be made efficient through the use of numerical methods for sparse matrices. In the case of the LGCP model the latent GRF which is a spatially continuous process is approximated by a GMRF on a discrete lattice so that these numerical methods can be applied. We consider this approximation for the case where the GRF is a Matérn field with ν known, so that $\theta = (\mu, \sigma^2, \phi)$ and the log intensity is characterized by $\mathbf{Y}$ as before. An approximation $\tilde{\pi}(\theta | \{s_k\})$ to the marginal posterior distribution $\pi(\theta | \{s_k\})$ is first obtained using the Laplace method. An approximation $\tilde{\pi}(y_i | \theta, \{s_k\})$ to the density of the full conditional distribution of each component of the latent field is then obtained using one of the three methods: a Gaussian approximation, the Laplace approximation, or a simplified Laplace approximation. The latter option is based on a series expansion of the Laplace approximation which has a lower computational cost. The marginal posterior distributions for the log intensity values of the LGCP model are then approximated through numerical integration over a discrete grid for θ

$$\tilde{\pi}(y_i | \{s_k\}) = \sum_k \tilde{\pi}(y_i | \theta_k, \{s_k\}) \tilde{\pi}(\theta_k, \{s_k\}) \Delta_k$$

where $\{\Delta_k\}$ are a set of area weights associated with the grid.

Recently, [22] develop an approximation to certain GRFs with Matérn correlation functions by specifying stochastic partial differential equations (SPDEs) that have certain Matérn processes as their solution. This SPDE representation provides an explicit link to GMRFs through a basis function representation of the solution where the corresponding weights comprise a GMRF with dependencies determined by a triangular mesh covering the spatial domain. This approximation can also be embedded within INLA and is implemented within the R-INLA package (obtained at www.r-inla.org) which allows for different mesh sizes. Increasing the size of the mesh will increase the accuracy of this approximation but will also increase the required time for computation.

We note, and express, here that the INLA package only allows for a Matérn correlation structure, whereas for the HMC and VB algorithms we use the power exponential family for computational purposes (e.g. it allows for easy calculation of the gradients).

3 Simulation Studies

Here the methods described in the previous section are compared using simulation studies. The discretized spatial domain is taken to be a 64×64 grid on the unit square, so that the simulated data are based on 4096 spatial locations. Each study is based on 1000 data sets simulated under the discretized LGCP model where we compare a total of six approaches: HMC incorporating FFT methods on the extended grid, VB incorporating the Laplace method, INLA with a simplified Laplace approximation (INLA I), INLA with a full Laplace approximation (INLA II), INLA with the SPDE model based on a mesh size of 436 (INLA III), and INLA with the SPDE model based on a mesh size of 4075 (INLA IV). The first mesh size was chosen purposely small and the second mesh size was chosen to be approximately the same as the number of cells in the discretized grid. In what follows, INLA

I and INLA II are also referred to as INLA with the lattice method, INLA III and INLA IV are also referred to as INLA with SPDE.

Our HMC and VB algorithms were derived under the assumption of a power exponential correlation for the GRF; whereas, we use the implementation of INLA in the standard R-INLA package that assumes a Matérn correlation. When simulating data we assume the GRF has a Matérn correlation and we apply all six approaches to the resulting data. Thus, INLA is based on a correctly specified covariance function and therefore has an advantage over HMC and VB which have the correlation function mis-specified. We will also assume that the parameter ν (in the Matérn model) is fixed and known, so that using INLA we only estimate the decay parameter ϕ in the Matérn model. As we cannot directly compare ρ and ϕ we are not able to compare directly the properties of the corresponding estimators. As an alternative we make comparisons with respect to the distance at which the correlation function drops to 0.5, denoted as $d_{0.5}$ and defined by the equations

$$r_p(d_{0.5}) = r_m(d_{0.5}) = 0.5 \quad (9)$$

where $r_p(\cdot)$ and $r_m(\cdot)$ denote the power exponential and Matérn correlation functions respectively.

3.1 Simulation One

We first simulate data sets from the LGCP model where the Matérn field has $\mu = 5$, $\sigma^2 = 3.5$, $\phi = 0.02$, and $\nu = 1$. Based on these parameters we simulate the GRF once and, based on this realization of the latent field, we simulate 1000 independent replicates of the data. Along with estimation of the log intensity values we also estimate μ , σ^2 , $d_{0.5}$ and $E(N)$, where $E(N)$ is the expected total number of points of the process within the spatial domain. The estimators are evaluated with respect to bias, variance and MSE. For the HMC and VB algorithms we match the shape of the power correlation function with that of the Matérn correlation function based on nonlinear least squares to estimate, and fix, the value for δ in the power exponential model, and we obtain a value of $\delta = 1.312$.

In terms of priors, HMC assumes a flat prior $\pi(\mu) \propto 1$, $\sigma^2 \propto I(0, \infty)$ and $\rho \propto I(0, \infty)$. With VB we are constrained to use conditionally conjugate priors and set $\mu \sim N(0, 625)$, $\sigma^2 \sim G^{-1}(1, 1)$, while ρ is estimated by the method of minimum contrast and assumed known in the VB algorithm. For INLA with the lattice method, we assign diffuse priors $\sigma^{-1} \sim \mathcal{G}(0.001, 0.001)$ and $d_I \sim \mathcal{G}(0.001, 0.001)$ where $d_I = \sqrt{8\nu}\phi$, and for INLA with SPDE we use the default joint-normal prior. Additional discussion of the priors and related numerical issues associated with INLA SPDE are mentioned in Section 5. INLA III has a mesh size of 436 (based on a length 0.1 for the inner mesh and 0.5 for outer mesh) and INLA IV has a mesh size of 4075 (based on a length 0.03 for the inner mesh and 0.5 for outer mesh). We acknowledge that the use of different priors for the HMC, VB, and INLA methods is not ideal as the differences we observe in the simulation results may, to some extent, be driven by differences in the priors. However, as the priors are taken to be fairly diffuse in all cases we do not expect the differences in the priors to play a significant role. As a practical matter

the form of the prior may sometimes be driven by the choice of computational algorithm used for Bayesian computation. Indeed, convergence issues may also impact the prior used in some circumstances. While not ideal from a theoretical perspective these issues are unavoidable from a practical standpoint. For example, the use of VB typically calls for conditionally conjugate priors while for INLA we are constrained to use certain forms for the priors that are built into the R-INLA package. As the different computational algorithms are often implemented with different priors we feel that these comparisons offer useful practical guidance for users despite these differences.

Figure 1 shows the true value of the discretized latent field $\mathbf{Y}$ in comparison to the average marginal posterior mean obtained from each of the methods under consideration. The average marginal posterior mean is the average of the estimates obtained from each of the 1000 simulation replicates. In this case we see that HMC, VB, INLA I and INLA II all produce similar average reconstructions of $\mathbf{y}$; whereas, the results from INLA III and INLA IV appear more over-smoothed, with the degree of over-smoothing decreasing as the mesh size increases.

Taking HMC as the baseline, a plot of the log-relative mean squared error (MSE) associated with the posterior mean estimator of $\mathbf{Y}$ for VB and INLA is shown in Figure 2. Points above (below) the black line indicate larger (smaller) MSE relative to that obtained from HMC. Both VB and INLA I–IV tend to produce estimators that have a higher MSE than the corresponding estimator obtained from HMC when the value of the log-intensity approaches either tail of the distribution of log-intensity values. Conversely, the methods appear to outperform HMC for values of the log-intensity around the median of this distribution. These differences appear to be more variable for INLA III and INLA IV compared with VB, INLA I and INLA II.

Table 1 displays the bias, variance, and MSE for the posterior mean estimators of μ , σ^{-2} , $d_{0.5}$ and $E(N)$. In this and all subsequent tables we display, for HMC, the actual values of bias, variance and MSE, whereas for VB and INLA the values are relative to those obtained from HMC. For estimation of these parameters we see that VB and INLA generally have larger bias and MSE compared with HMC. Both INLA I and INLA II outperform VB slightly, while INLA III and INLA IV lag behind the alternatives rather significantly. By construction, INLA I and INLA II will provide identical estimates for σ^{-2} and $d_{0.5}$ and this is reflected in the table. For the estimation of $E(N)$, HMC outperforms VB with respect to bias, variance, and MSE, while these measures are not reported for INLA as we only obtain marginal posterior distributions of the log-intensity values in this case and thus cannot estimate $E(N)$.

While Table 1 displays the variance of the posterior mean estimators, we display in Table 2 the average (over simulation replicates) marginal posterior variance. The associated large sample theory guarantees that the posterior variance as obtained from HMC is simulation consistent. As such the posterior variance for VB and INLA are again listed relative to that obtained from HMC in order to determine the extent to which these approaches under-estimate or over-estimate posterior variability. For VB we see that the marginal posterior variance is under-estimated which is in line with expectations from the literature [34, 32].

INLA I and II provide measures of variability that are closer to that of HMC, while INLA III and IV tend to over-estimate the marginal posterior variance, with this over-estimation being substantial when the smaller mesh size is used. With respect to average computational time, HMC requires 679 seconds based on 1500 total iterations with the first 500 thrown away as burn-in; VB requires 453 seconds to run to convergence (about 300 iterations); INLA I runs for 46 seconds, INLA II requires 196 seconds, INLA III requires 10 seconds, and INLA IV requires 154 seconds. The algorithms are run on an iMac with a 3.2 GHz Intel Core i5 processor and 16GB memory.

3.2 Simulation Two

In order to make comparisons in a setting where there is a slower decay for the spatial correlation and smoother realizations of the random field our second study is based on setting $\phi = 0.05$ and $\nu = 3$ with all other settings remaining unchanged. Figure 3 presents the average posterior mean log-intensity over 1000 simulation replicates. Comparing the images the overall best reconstruction seems to arise from both VB and HMC, followed by INLA I, INLA II, and INLA IV all of which capture the general features of the true latent field, while INLA III seems subject to over-smoothing as before. Figure 4 displays the log-relative MSE of the five methods in comparison with HMC as in Figure 2. Interestingly, INLA IV seems to have the best performance in terms of MSE here, which taken together with Figure 3 suggests that the estimators from INLA IV may have lower variance while still achieving adequate bias.

With respect to the hyper-parameters, Table 3 displays the bias, variance, and MSE of the posterior mean estimators for μ , σ^{-2} , $d_{0.5}$ and $E(N)$. INLA III and INLA IV have the smallest MSE for μ , but have large MSE for $d_{0.5}$. As before HMC attains the lowest MSE for estimation of σ^{-2} and $d_{0.5}$. In terms of $E(N)$, HMC outperforms VB with respect to bias, variance and MSE. Table 4 shows the average marginal posterior variance for the hyper-parameters. INLA I and II generally under-estimate the marginal posterior variance. VB shows significant under-estimation of the marginal variance for σ^{-2} and $d_{0.5}$, and over-estimation of the posterior variance for μ . As with the previous study, INLA III and INLA IV over-estimate the marginal posterior variance for all three hyper-parameters. In terms of average timing, HMC requires 521 seconds for a total of 2000 iterations with the first 1000 iterations discarded as burn-in; VB requires 467 seconds and typically required a greater number of iterations to converge (about 1000) compared to simulation one; INLA I takes 134 seconds, INLA II takes 357 seconds, INLA III takes 11 seconds, INLA IV takes 227 seconds.

4 Application

We next compare the computational algorithms through an application to two data sets where the LGCP model is applied in both cases. In addition to comparing the methods with respect to posterior summaries of the parameters of interest, we also make comparisons with respect to goodness-of-fit checking using the posterior predictive distribution [35] which has been applied for checking hierarchical spatial models in a number of applications including disease ecology and neuroimaging [36, 37, e. g.]). The posterior predictive checks are based

on the L function [38] where we simulate, based on the model, posterior predictive replicates of the discrepancy measure $\delta(r) = L(r, \{s_k\}^{obs}, \mathbf{y}, \theta) - L(r, \{s_k\}^{rep}, \mathbf{y}, \theta)$ where $\{s_k\}^{obs}$ denotes the observed data, $\mathbf{y}$ and θ are drawn from the posterior distribution, and $\{s_k\}^{rep}$ denotes replicate data that is drawn from the posterior predictive distribution. For a given distance range r , if the value $\delta(r) = 0$ is observed to lie as an extreme value in either tail of the posterior predictive distribution we may question the fit of the model as characterized by the L function at that distance. As INLA does not provide the joint posterior distribution we are unable to simulate predictive realizations and thus we make the predictive comparisons only between HMC and VB.

All priors are the same as in the simulation studies unless otherwise indicated. We also compare VB and the INLA methods with HMC as HMC is simulation consistent.

4.1 Bramble Canes data

The data record the (x, y) locations of 823 bramble canes in a field of 9 m^2 , rescaled to a unit square. The data are depicted in Figure 5(a) and were recorded and analyzed by Hutchings (1979) and further analyzed by Diggle [39].

To determine the value of δ in the power-exponential correlation function and the value of ν in the Matérn correlation function, we use the method of minimum contrast which estimates the values as $\hat{\delta} = 0.51$, $\hat{\nu} = 0.02$. As the R-INLA package only offers three possible values $\nu = 1, 2, 3$, we select $\nu = 1$ as it is the closest of the three choices to the estimate obtained from the minimum contrast method. Figure 6 depicts the posterior mean of the log-intensity as obtained by each of the six methods. HMC, INLA I and INLA II result in estimated images that appear fairly consistent; results obtained from VB are also consistent with HMC but less so than INLA I and II, while both INLA III and INLA IV appear to over-smooth the estimated latent field relative to the other methods.

Posterior summaries of the hyper-parameters μ , σ^{-2} and $d_{0.5}$ are presented in Table 5. Here we see that VB is under-estimating the posterior variance for all hyper-parameters while the point estimates for μ , σ^{-2} are larger than those obtained from HMC, but within somewhat reasonable bounds. The point estimates obtained from INLA with the lattice method are closer to those obtained from HMC than those obtained from INLA with SPDE. In Figure 7 we compare the marginal posterior variance of each element of the latent field as obtained from all of the methods to the posterior variance obtained from HMC. In this case all of the methods under-estimate the posterior variance and this under-estimation is most severe for INLA with SPDE.

Figure 8 compares the 95% posterior predictive intervals for $\delta(r)$ as obtained from both HMC and VB. Comparing the figures indicates that the posterior predictive variability is under-estimated for VB, primarily at the lower distance ranges r . The implication of this for data analysis is that posterior predictive checks for VB under similar settings would be conservative. In terms of the data, the posterior predictive check as obtained from HMC does not reveal a lack-of-fit with respect to the chosen discrepancy measure. With respect to timing, HMC requires 597s for 1500 iterations and 500 burn-in iterations; VB actually

requires more time than HMC in this example and takes 1012s requiring 137 iterations to convergence. INLA I require 55s, INLA II 172s, INLA III 11s, and INLA IV takes 227s.

4.2 Multiple Sclerosis MRI Data

Our second application consists of a point pattern depicting the locations of Multiple Sclerosis (MS) lesions obtained from taking a slab of sagittal slices (10mm thick) obtained from magnetic resonance imaging from a cohort of MS patient and converting the spatial domain to the unit square. The point pattern consists of 1950 locations and is depicted in Figure 5(b). Aside from the application this data set differs from the first in that the observed level of aggregation is higher and the points are more unevenly distributed. The method of minimum contrast is used to select values of $\hat{\delta} = 1.165$ and $\hat{\nu} = 1$ for the covariance functions. The posterior mean of the log-intensity values are depicted in Figure 9. In this case all methods seem to capture the same general features of the image.

Table 6 presents posterior summaries of the hyper-parameters. The point estimates for μ obtained from all approaches are fairly consistent while, relative to HMC, posterior variance is not well estimated by either VB or INLA I–IV. Relative to HMC the precision σ^{-2} is not well estimated by any of the methods and similarly for the posterior variance of this parameter. As $d_{0.5}$ is estimated rather precisely by HMC, all of VB and INLA I–III either severely under-estimate or over-estimate the posterior variability though this is not the case for INLA IV. Figure 10 compares the marginal posterior variance for each cell of the discretized latent field obtained from all of the methods with the posterior variance obtained from HMC. Interestingly, in this case we find that VB over-estimates the posterior variability while INLA I–III all under-estimate the posterior variability to different degrees relative to HMC. The posterior variability arising from INLA IV gives extremely large values (up to 10 times larger than those obtained from HMC) which likely indicates a numerical problem, though we note again that INLA IV does give an adequate representation of the posterior mean for this data set.

Turning to posterior predictive checks which are depicted in Figure 11 we see that VB has much wider 95% posterior predictive intervals than HMC. Although neither algorithm shows a lack of fit, using HMC, the model appears to fit the data better as the mean and median are much closer to zero for all ranges of r and the 95% predictive intervals are much tighter at each value of r . The greater posterior predictive variability arising from VB may be in part a result of the posterior variance of μ being over-estimated by VB. In terms of timing, HMC takes 1456s with 2000 iterations and 1000 burn-in. Again, VB actually requires more computation time 1608s with 2763 iterations required for convergence. INLA I takes 47s, INLA II takes 166s, INLA III takes 57s, INLA IV takes 384s.

5 Discussion

We have compared HMC incorporating FFT matrix methods on an extended grid, VB incorporating the Laplace method, and four versions of INLA for Bayesian computation and estimation associated with the LGCP model. One would not expect to get back the true latent field following inference via any of these methods, rather, the gold standard is the posterior density of the field given the data. Of the methods considered, only HMC has

theoretical guarantees in terms of (simulation) consistency in attaining the gold standard so the empirical comparisons considered here are of practical value. A number of settings for both simulated and real data have been adopted for these comparisons. Overall, in terms of point estimation of the latent field we do observe some differences in some settings; however, generally, all of HMC, VB, INLA I, and INLA II perform reasonably well, while INLA with SPDE has a tendency to over-smooth the field, though this tendency is reduced as the size of the underlying mesh is increased. Thus if point estimation of the log-intensity is the only objective we recommend the use of INLA I based on the required computation time. If, in addition, inference on hyper-parameters is of importance then it seems clear that HMC and the additional required computation time is necessary. As expected from the literature VB has a tendency to under-estimate posterior variability although on occasion it also seems to over-estimate posterior variability (see e.g. [40] for discussion of the latter issue and how it may also arise with VB). We also find that posterior predictive checking based on VB may not be representative of the true posterior predictive distribution. While VB has been applied successfully in a wide range of applications and is often the method of choice in machine learning the required computation time for the LGCP model and for the settings considered here suggest that it is not as accurate as HMC and not as computationally efficient as INLA I. Of course, our implementation of VB did not incorporate the FFT methods for matrix multiplication as this is not straightforward to implement within the VB framework. We acknowledge that fixing ρ in the VB algorithm based on the minimum contrast estimate is not ideal; however, this was the only practical possibility given the convergence problems we see for VB when ρ is treated as unknown. We suspect that the convergence problem is somehow related to the mean-field approximation and its assumption of posterior independence between ρ and the other model parameters, which likely conflicts with the data. Further examination of this issue, perhaps more generally for mean field approximations is a potentially interesting avenue for further work. To improve the variational approximation one approach that may be worth considering is the use of fixed-form multivariate Gaussian variational approximations which may have improved performance over mean field approximations.

To compare within the four flavors of INLA, the INLA with simplified Laplace method and the full Laplace method are quite similar in terms of accuracy, for the settings considered here; whereas, we see genuine gains in computation with the use of the simplified Laplace version of INLA. In addition to a tendency to exhibit over-smoothing, we have also found that INLA with SPDE can be numerically unstable in some situations and an inappropriate choice of step size in the Newton-Raphson algorithm can lead to convergence problems. Thus the choice of mesh is an important consideration. We acknowledge that the different priors and different covariance functions used with INLA relative to the competing approaches does make the interpretation of the simulation results comparing to VB and MCMC more difficult than had the same priors and covariance functions been used in all cases; however, as a practical matter, practitioners will use the built-in priors and the Matérn covariance function available in INLA, so we feel that these comparisons are of direct interest and practical value.

Acknowledgments

Dr. Johnson would like to thank the National Institute of Neurological Disorders and Stroke for their financial support: NIH Grant number 5-R01-NS-075066. Dr. Nathoo is supported by funding from the Natural Sciences and Engineering Research Council of Canada and holds a Tier II Canada Research Chair in Biostatistics for Spatial and High-Dimensional Data. The work presented in this manuscript represents the views of the authors and not necessarily that of the NINDS, NIH, nor the NSEERC.

References

1. Daley, DJ., Vere-Jones, D. An introduction to the theory of point processes. Vol. 2. Springer; 1988.
2. Møller J, Syversveen AR, Waagepetersen RP. Log gaussian cox processes. *Scandinavian journal of statistics*. 1998; 25(3):451–482.
3. Illian, J., Penttinen, A., Stoyan, H., Stoyan, D. Statistical analysis and modelling of spatial point patterns. John Wiley & Sons; 2008.
4. Murray I, Ghahramani Z, MacKay D. MCMC for doubly-intractable distributions. 2012 arXiv preprint arXiv:12066848.
5. Diggle P. A kernel method for smoothing point process data. *Applied Statistics*. 1985; 34(2):138–147.
6. Rosenblatt M, et al. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*. 1956; 27(3):832–837.
7. Besag J. Discussion on the paper by Grenander and Miller. *Journal of the Royal Statistical Society, Series B*. 1994; 56:591–592.
8. Adams, RP., Murray, I., MacKay, DJ. Proceedings of the 26th Annual International Conference on Machine Learning. ACM; 2009. Tractable nonparametric Bayesian inference in Poisson processes with Gaussian process intensities; p. 9-16.
9. Gonçalves FB, Gamerman D. Exact Bayesian inference in spatio-temporal Cox processes driven by multivariate Gaussian processes. 2015 arXiv preprint arXiv:150406638.
10. Neal, RM. Bayesian learning for neural networks. Springer; 1996.
11. Neal, R. MCMC using Hamiltonian dynamics. In: Brooks, S, Gelman, A, Jones, GL., Meng, XL., editors. *Handbook of Markov Chain Monte Carlo*. CRC Press; 2011. p. 113-162.
12. Girolami M, Calderhead B. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2011; 73(2):123–214.
13. MacKay, DJ. Ensemble learning for hidden Markov models. Technical report, Cavendish Laboratory, University of Cambridge; 1997.
14. Humphreys, K., Titterton, D. COMPSTAT. Springer; 2000. Approximate Bayesian inference for simple mixtures; p. 331-336.
15. Zammit-Mangion A, Sanguinetti G, Kadirkamanathan V. Variational estimation in spatiotemporal systems from continuous and point-process observations. *Signal Processing, IEEE Transactions on*. 2012; 60(7):3449–3459.
16. Wang C, Blei DM. Variational inference in nonconjugate models. *The Journal of Machine Learning Research*. 2013; 14(1):1005–1031.
17. Wang C, Blei DM. A General Method for Robust Bayesian Modeling. 2015 arXiv preprint arXiv:151005078.
18. Nguyen TV, Bonilla EV. Automated variational inference for Gaussian process models. *Advances in Neural Information Processing Systems*. 2014:1404–1412.
19. Lloyd C, Gunter T, Osborne MA, Roberts SJ. Variational inference for Gaussian process modulated Poisson processes. 2014 arXiv preprint arXiv:14110254.
20. Rue H, Martino S, Chopin N. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the royal statistical society: Series b (statistical methodology)*. 2009; 71(2):319–392.
21. Rue, H., Held, L. Gaussian Markov random fields: theory and applications. CRC Press; 2005.

22. Lindgren F, Rue H, Lindström J. 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)*. 2011; 73(4):423–498.
23. Simpson D, Lindgren F, Rue H. Think continuous: Markovian Gaussian models in spatial statistics. *Spatial Statistics*. 2012; 1:16–29.
24. Taylor BM, Diggle PJ. INLA or MCMC? A tutorial and comparative evaluation for spatial prediction in log-Gaussian Cox processes. *Journal of Statistical Computation and Simulation*. 2013;1–9. ahead-of-print.
25. Møller, J., Waagepetersen, RP. Statistical inference and simulation for spatial point processes. CRC Press; 2003.
26. Matérn B. Spatial Variation. Meddelanden från Statens Skogsforskningsinstitut. 1960; 49(5)
27. Alder BJ, Wainwright T. Studies in molecular dynamics. I. General method. *The Journal of Chemical Physics*. 1959; 31(2):459–466.
28. Duane S, Kennedy AD, Pendleton BJ, Roweth D. Hybrid monte carlo. *Physics letters B*. 1987; 195(2):216–222.
29. Wood ATA, Chan G. Simulation of stationary Gaussian processes in $[0, 1]^d$. *Journal of Computational and Graphical Statistics*. 1994; 3(4):409–432.
30. Roberts GO, Rosenthal JS, et al. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical science*. 2001; 16(4):351–367.
31. Beskos A, Pillai N, Roberts G, Sanz-Serna JM, Stuart A, et al. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*. 2013; 19(5A):1501–1534.
32. Nathoo F, Babul A, Moiseev A, Virji-Babul N, Beg M. A variational Bayes spatiotemporal model for electromagnetic brain mapping. *Biometrics*. 2014; 70(1):132–143. [PubMed: 24354514]
33. Diggle PJ, Gratton RJ. Monte Carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society Series B (Methodological)*. 1984:193–227.
34. Nathoo F, Lesperance M, Lawson A, Dean C. Comparing variational bayes with markov chain monte carlo for bayesian computation in neuroimaging. *Statistical methods in medical research*. 2013; 22(4):398–423. [PubMed: 22642986]
35. Gelman A, Meng XL, Stern H. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*. 1996; 6(4):733–760.
36. Nathoo F. Joint spatial modeling of recurrent infection and growth with processes under intermittent observation. *Biometrics*. 2010; 66(2):336–346. [PubMed: 19673862]
37. Kang J, Johnson TD, Nichols TE, Wager TD. Meta analysis of functional neuroimaging data via Bayesian spatial point processes. *Journal of the American Statistical Association*. 2011; 106(493): 124–134. [PubMed: 21706069]
38. Illian JB, Møller J, Waagepetersen RP. Hierarchical spatial point process analysis for a plant community with high biodiversity. *Environmental and Ecological Statistics*. 2009; 16(3):389–405.
39. Diggle, PJ., et al. Statistical analysis of spatial point patterns. Academic Press; 1983.
40. Daunizeau J, Friston K, Kiebel S. Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Physica D: nonlinear phenomena*. 2009; 238(21):2089–2118. [PubMed: 19862351]

Appendices

Appendix A: Gradient derivation for ρ

For the circulant matrix $\mathbf{E}$ with base $\mathbf{e} = (e_0, \dots, e_{m^2-1})$, the j^{th} eigenvalue is given by:

$$\lambda_i = \sum_{j=0}^{m^2-1} \mathbf{e}_j \exp(i2\pi j i / m^2) \quad (10)$$

where $\iota = \sqrt{-1}$. For the power exponential family of correlations $e_j = \exp(-\rho d_j^\delta)$ where d_j is the distance from origin. So we have

$$\lambda_i = \sum_{j=0}^{m^2-1} \exp(-\rho d_j^\delta) \exp(\iota 2\pi j i / m^2) \quad (11)$$

Thus, we have $\mathbf{E} = \mathbf{F} \mathbf{\Lambda} \mathbf{F}^H$ where $\mathbf{F}$ is the matrix of eigenvectors, and $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues with i^{th} value to be λ_i .

To derive the partial derivative of $\log \pi(\rho | \cdot)$, we first derive the partial derivative of $\mathbf{E}^{\frac{1}{2}} \gamma$

$$\frac{\partial}{\partial \rho} (\mathbf{E}^{\frac{1}{2}} \gamma)_i = (\mathbf{F} \frac{\partial}{\partial \rho} \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{F}^H \gamma)_i \quad (12)$$

So we need the partial derivative of each diagonal element of $\mathbf{\Lambda}^{1/2}$ w.r.t ρ .

$$\frac{\partial \lambda_i^{\frac{1}{2}}}{\partial \rho} = \frac{\partial}{\partial \rho} \left(\sum_{j=0}^{m^2-1} \exp(-\rho d_j^\delta) \exp(\iota 2\pi j i / m^2) \right)^{\frac{1}{2}} \quad (13)$$

$$= -\frac{1}{2} \lambda_i^{-\frac{1}{2}} \sum_{j=0}^{m^2-1} d_j^\delta e_j \exp(\iota 2\pi j i / m^2) \quad (14)$$

The summand in the last line turns out to be the base of a particular matrix with base $\mathbf{e}^* = (d_0^\delta e_0, \dots, d_{m^2-1}^\delta e_{m^2-1})$. Consider the circulant matrix $\mathbf{D}$ with base $\mathbf{d} = (d_j^\delta, \dots, d_{m^2-1}^\delta)$. Then it is easy to show that $\mathbf{E}^* = \mathbf{D} \odot \mathbf{E}$ is a circulant matrix with base $\mathbf{e}^* = \mathbf{d} \odot \mathbf{e}$ where $\odot$ represents element wise multiplication. And $\sum_{j=0}^{m^2-1} d_j^\delta e_j \exp(\iota 2\pi j i / m^2)$ is the i^{th} eigenvalue of $\mathbf{E}^*$. Call it ψ_i . Thus,

$$\frac{\partial \lambda_i^{\frac{1}{2}}}{\partial \rho} = -\frac{1}{2} \lambda_i^{-\frac{1}{2}} \psi_i \quad (15)$$

Putting this all together we have

$$\frac{\partial}{\partial \rho} (\mathbf{E}^{\frac{1}{2}} \gamma)_i = \frac{\partial}{\partial \rho} - \frac{1}{2} (\mathbf{F} \mathbf{\Lambda}^{-\frac{1}{2}} \Psi \mathbf{F}^H \gamma)_i \quad (16)$$

$$= -\frac{1}{2}(\mathbf{F}\Lambda^{-\frac{1}{2}}\mathbf{F}^H\mathbf{F}\Psi\mathbf{F}^H\boldsymbol{\gamma})_i \quad (17)$$

$$= -\frac{1}{2}(\mathbf{E}^{-\frac{1}{2}}\mathbf{E}^*\boldsymbol{\gamma})_i \quad (18)$$

Now we can derive the gradient for ρ :

$$\frac{\partial}{\partial \rho} \log \pi(\rho|\cdot) = \frac{\partial}{\partial \rho} \left\{ \sum_i [y_i m_i - A \exp(y_i)] + \log \pi(\rho) \right\} \quad (19)$$

$$= \sum_i \frac{\partial}{\partial \rho} \left\{ \mu m_i + \sigma (\mathbf{E}^{\frac{1}{2}}\boldsymbol{\gamma})_i m_i - A \exp [\mu + \sigma (\mathbf{E}^{\frac{1}{2}}\boldsymbol{\gamma})_i] \right\} + \frac{\pi'(\rho)}{\pi(\rho)} \quad (20)$$

$$= \sum_i \left\{ \sigma \frac{\partial}{\partial \rho} (\mathbf{E}^{\frac{1}{2}}\boldsymbol{\gamma})_i m_i - A \frac{\partial}{\partial \rho} \exp [\mu + \sigma (\mathbf{E}^{\frac{1}{2}}\boldsymbol{\gamma})_i] \right\} \quad (21)$$

$$= -\frac{1}{2}\sigma \sum_i \left\{ m_i - \frac{A\sigma}{2} \exp [\mu + \sigma (\mathbf{E}^{\frac{1}{2}}\boldsymbol{\gamma})_i] \right\} (\mathbf{E}^{-\frac{1}{2}}\mathbf{E}^*\boldsymbol{\gamma})_i \quad (22)$$

$$= -\frac{\sigma}{2} \left[\mathbf{m} - A \exp (\mu \mathbf{1}_{m^2} + \sigma \mathbf{E}^{\frac{1}{2}}\boldsymbol{\gamma}) \right]^T \mathbf{E}^{-\frac{1}{2}}\mathbf{E}^*\boldsymbol{\gamma} \quad (23)$$

where $\mathbf{m} = (m_1, \dots, m_{m^2})$ and we use the fact that $\boldsymbol{\pi}'(\rho) = 0$ for flat prior of ρ .

Because $\mathbf{E}^{\frac{1}{2}}\mathbf{E}^*\boldsymbol{\gamma} = \mathbf{F}\Lambda^{-\frac{1}{2}}\mathbf{F}^H\mathbf{F}\Psi\mathbf{F}^H\boldsymbol{\gamma} = \mathbf{F}\Lambda^{-\frac{1}{2}}\Psi\mathbf{F}^H\boldsymbol{\gamma}$, we can use the DFT to compute all the matrix operations in the equation above.

Figure 1.
Average marginal posterior mean of the log-intensity over 1000 samples from the first simulation study. Upper left panel is the true GRF.

Figure 2.

Log-relative MSE of estimated latent GRF from VB and INLA I–IV to that from HMC for the first simulation study. Each point represents the log-relative MSE of the discretized GRF. Points located above, on and below the horizontal line denotes bigger, equal and smaller MSE than that from HMC.

Figure 3.

Average marginal posterior mean of the log-intensity over 1000 samples from the second simulation study. Upper left panel is the true GRF.

Figure 4.

Log-relative MSE of estimated latent GRF from VB and INLA I–IV to that from HMC for the second simulation study. Each point represents the log-relative MSE of the discretized GRF. Points located above, on and below the horizontal line denotes bigger, equal and smaller MSE than that from HMC.

bramble canes location

MS lesion location

Figure 5.

(a) shows the Bramble canes locations, (b) shows the MS lesion locations. Both are represented by black dots, rescaled to the unit square.

Figure 6. Posterior mean of the latent GRF for the bramble canes data set, estimated from HMC, VB and INLA I–IV.

Figure 7. Scatter plot of the marginal posterior variance of the latent GRF from VB and INLA I–IV compared with those from HMC. Bramble canes data set.

95% posterior predictive interval for HMC

95% posterior predictive interval for VB

Figure 8.

95% posterior predictive interval for HMC (a) and VB (b) for the bramble canes data set.

The bounds are denoted by solid lines while the mean and median are denoted by dashed lines. These are obtained at 20 distinct distances.

Figure 9.
Posterior mean of the latent GRF, estimated from HMC, VB and INLA I–IV. MS data set

Figure 10.
Scatter plot of the marginal posterior variance of the latent GRF from VB and INLA I–IV compared with those from HMC for the MS data set.

95% posterior predictive interval for HMC

95% posterior predictive interval for VB

Figure 11.

95% posterior predictive interval for HMC (a) and VB (b) for the MS data set. The bounds are denoted by solid lines while the mean and median are denoted by dashed lines. These are obtained at 20 distinct distances.

Summary of the statistical properties for the hyper-parameters from the first simulation study. The values shown in table from VB and INLA I–IV are relative to that from HMC.

Table 1

Parm	Measure	Relative Measure							
		HMC	VB	INLA I	INLA II	INLA III	INLA IV		
$\mu = 5$	Bias	0.046	12.681	5.469	8.471	25.805	15.435		
	Var	0.010	0.476	0.903	0.663	8.453	0.632		
	MSE	0.012	29.441	6.143	13.507	127.227	43.558		
$\sigma^2 = 0.286$	Bias	-0.010	6.984	-3.727	-3.727	-72.906	-14.545		
	Var	3.20-04	0.454	1.338	1.338	36.531	4.810		
	MSE	4.08-04	11.436	4.191	4.191	1236.295	51.801		
$d_{0.5} = 0.025$	Bias	3.30E-05	-122.467	122.245	122.245	1704.511	1115.598		
	Var	2.20E-06	0.465	1.244	1.244	1199.323	4.574		
	MSE	2.20E-06	7.895	8.588	8.588	2626.682	616.258		
$E(N) = 910.29$	Bias	-1.337	-252.070	-	-	-	-		
	Var	918.282	1.420	-	-	-	-		
	MSE	920.069	124.820	-	-	-	-		

Table 2

Marginal variance estimates of the parameters for the first simulation study. VB and INLA I–IV are relative to HMC

Average Marginal Var			Relative Ave. Marg. Var			
Param	HMC	VB	INLA I	INLA II	INLA III	INLA IV
μ	0.028	0.511	0.868	0.333	85.232	2.287
σ^2	0.001	0.029	1.576	1.576	6.780	17.061
$d_{0.5}$	4.80E-06	0.022	1.667	1.667	87549.788	7.161

Table 3

Summary of the statistical properties for the hyper-parameters from the second simulation study. The values shown in table from VB and INLA I–IV are relative to that from HMC.

Parm	Measure	Relative Measure							
		HMC	VB	INLA I	INLA II	INLA III	INLA IV		
$\mu = 5$	Bias	-0.458	2.752	-1.513	-1.488	0.047	0.458		
	Var	0.029	1.169	4.682	4.617	0.932	1.232		
	MSE	0.239	6.797	2.579	2.507	0.115	0.334		
$\sigma^2 = 0.286$	Bias	-0.017	15.316	3.355	3.355	2.038	4.477		
	Var	1.20E-03	0.003	3.75	3.75	1.131	0.849		
	MSE	0.001	48.492	5.301	5.301	1.755	4.817		
$d_{0.5} = 0.13$	Bias	-1.00E-02	4.364	6.893	6.893	-43.224	-31.924		
	Var	7.90E-05	0.191	1.236	1.236	24.439	17.627		
	MSE	1.80E-04	10.868	27.431	27.431	1068.207	584.571		
$E(N) = 494.12$	Bias	0.077	1201.419	-	-	-	-		
	Var	496.455	1.199	-	-	-	-		
	MSE	496.461	18.527	-	-	-	-		

Table 4

Marginal variance estimates of the parameters from the second simulation study. VB and INLA I–IV are relative to HMC

Parm	Average Marginal Var		Relative Ave. Marg. Var			
	HMC	VB	INLA I	INLA II	INLA III	INLA IV
μ	0.255	6.084	0.301	0.299	5.952	4.531
σ^2	0.006	3.36E-05	0.650	0.650	23.322	20.661
$d_{0.5}$	3.10E-04	0.001	0.164	0.164	12.023	5.935

Table 5
Summary of parameter estimation for the bramble canes data set, VB and INLA I–IV are relative to HMC.

Parm	Measure	Relative Measure					
		HMC	VB	INLA I	INLA II	INLA III	INLA IV
μ	Post. Mean	5.019	1.223	0.996	1.078	1.197	1.161
	Post. Var	0.016	0.52	1.503	0.535	32.469	4.744
σ^2	Post. Mean	0.272	1.743	1.135	1.135	2.461	1.998
	Post. Var	0.001	0.194	1.581	1.581	114.182	26.765
$d_{0.5}$	Post. Mean	0.025	0.165	0.768	0.768	9.215	2.894
	Post. Var	8.00E-05	9.98E-06	0.025	0.025	46.779	0.945

Table 6
Summary of parameter estimation for the MS data set. VB and INLA I–IV are relative to HMC.

Parm	Measure	Relative Measure							
		HMC	VB	INLA I	INLA II	INLA III	INLA IV		
μ	Post. Mean	5.098	0.959	1.074	1.072	1.193	0.814		
	Post. Var	0.301	2.638	0.081	0.084	0.197	32.731		
σ^2	Post. Mean	0.468	0.282	0.1	0.1	1.729	0.192		
	Post. Var	0.005	0.002	31.83	31.83	1.19E-36	0.404		
$d_{0.5}$	Post. Mean	0.18	0.886	14.379	14.379	0.503	2.789		
	Post. Var	1.30E-04	2.41E-10	31653.155	31653.155	2.91E-37	0.959		