

Chapter 6

Bayesian Computation with Intractable Likelihoods



Matthew T. Moores, Anthony N. Pettitt, and Kerrie L. Mengersen

Abstract This chapter surveys computational methods for posterior inference with intractable likelihoods, that is where the likelihood function is unavailable in closed form, or where evaluation of the likelihood is infeasible. We survey recent developments in pseudo-marginal methods, approximate Bayesian computation (ABC), the exchange algorithm, thermodynamic integration, and composite likelihood, paying particular attention to advancements in scalability for large datasets. We also mention R and MATLAB source code for implementations of these algorithms, where they are available.

Keywords Composite likelihood · Likelihood-free inference · Markov random fields · Pseudo-marginal methods

The likelihood function plays an important role in Bayesian inference, since it connects the observed data with the statistical model. Both simulation-based (e.g. MCMC) and optimisation-based (e.g. variational Bayes) algorithms require the likelihood to be evaluated pointwise, up to an unknown normalising constant. However, there are some situations where this evaluation is analytically and computationally intractable. For example, when the complexity of the likelihood grows at a combinatorial rate in terms of the number of observations, then likelihood-based inference quickly becomes infeasible for the scale of data that is regularly encountered in applications.

M. T. Moores (✉)

National Institute for Applied Statistics Research Australia, School of Mathematics & Applied Statistics, University of Wollongong, Keiraville, NSW, Australia
e-mail: mmoores@uow.edu.au

A. N. Pettitt · K. L. Mengersen

School of Mathematical Sciences, Queensland University of Technology, Brisbane, QLD, Australia
e-mail: a.pettitt@qut.edu.au; k.mengersen@qut.edu.au

© The Editor(s) (if applicable) and The Author(s), under exclusive licence to Springer Nature Switzerland AG 2020

K. L. Mengersen et al. (eds.), *Case Studies in Applied Bayesian Data Science*, Lecture Notes in Mathematics 2259, https://doi.org/10.1007/978-3-030-42553-1_6

Intractable likelihoods arise in a variety of contexts, including models for DNA mutation in population genetics [43, 64], models for the spread of disease in epidemiology [46, 60], models for the formation of galaxies in astronomy [12], and estimation of the model evidence in Bayesian model choice [29]. This chapter will mainly focus on Markov random field (MRF) models with discrete state spaces, such as the Ising, Potts, and exponential random graph models (ERGM). These models are used for image segmentation or analysis of social network data, two areas where millions of observations are commonplace. There is therefore a need for scalable inference algorithms that can handle these large volumes of data.

The Ising, Potts, or ERGM likelihood functions can be expressed in the form of an exponential family:

$$p(\mathbf{y} \mid \boldsymbol{\theta}) = \frac{\exp\{\boldsymbol{\theta}^T \mathbf{s}(\mathbf{y})\}}{C(\boldsymbol{\theta})}, \quad (6.1)$$

where the observed data $\mathbf{y} = y_1, \dots, y_n$ is in the form of an undirected graph, $\boldsymbol{\theta}$ is a vector of unknown parameters, $\mathbf{s}(\mathbf{y})$ is a corresponding vector of jointly-sufficient statistics for these parameters, and $C(\boldsymbol{\theta})$ is an intractable normalising constant, also known as a partition function:

$$C(\boldsymbol{\theta}) = \sum_{\mathbf{y} \in \mathcal{Y}} \exp\{\boldsymbol{\theta}^T \mathbf{s}(\mathbf{y})\}, \quad (6.2)$$

where the sum is over all possible configurations of states, $\mathbf{y} \in \mathcal{Y}$.

In the case of an Ising model, a single node can take one of two possible values, $y_i \in \{0, 1\}$. For example, in image analysis the value 1 might represent a foreground pixel, while 0 represents the background. The q -state Potts model generalises this construction to more than two states, so $y_i \in \{1, \dots, q\}$. The cardinality of the configuration space, $\#\mathcal{Y}$, is then q^n . Even with only 2 states and $n = 100$ pixels, computation of (6.2) requires more than 10^{30} floating point operations. It would take a supercomputer with 100 PetaFLOPS over 400,000 years to find an answer.

Both the Ising and Potts models possess a single parameter, β , known as the inverse temperature. The corresponding sufficient statistic is then:

$$s(\mathbf{y}) = \sum_{i \sim \ell \in \mathcal{E}} \delta(y_i, y_\ell), \quad (6.3)$$

where \mathcal{E} is the set of all unique pairs of neighbours $i \sim \ell$ in the observed graph, and $\delta(x, y)$ is the Kronecker delta function, which equals 1 when $x = y$ and 0 otherwise. We assume a first-order neighbourhood structure, so a given pixel y_i would have up to 4 neighbours in a regular 2D lattice, or 6 neighbours in 3D. Pixels on the boundary of the image domain have less than 4 (or 6) neighbours, so $\#\mathcal{E} = 2(n - \sqrt{n})$ for a square 2D lattice, or $3(n - n^{2/3})$ for a cube.

The observed data for an ERGM can be represented as a binary adjacency matrix Y , encoding the presence or absence of a neighbourhood relationship between nodes i and j : $[Y]_{i,j} = 1$ if $i \sim j$; $[Y]_{i,j} = 0$ otherwise. $\#\mathcal{Y}$ for an ERGM is equal to 2^M , where $M = n(n - 1)/2$ is the maximum number of ties in an undirected graph with n nodes. As with the Ising or Potts models, computing the normalising constant (6.2) is therefore intractable for non-trivial graphs. Various kinds of ERGM can be defined by the choice of sufficient statistics. The simplest example is the Bernoulli random graph [23], which has a single statistic $s_1(Y) = m$, the number of connected neighbours in the graph. In an undirected graph, this is half the number of nonzero entries in the adjacency matrix. An important class of graph statistics are the numbers of k -stars [26], which can be defined in terms of the degree distribution [59]:

$$n_k = \sum_{i=1}^n \binom{d_i}{k}, \quad (6.4)$$

where the degree d_i is the number of neighbours of node i :

$$d_i = \sum_{j=1}^n [Y]_{ij}. \quad (6.5)$$

Note that under this definition $n_1 = 2m$, since each tie is counted twice. An alternative definition, which avoids double-counting, is given by:

$$\begin{aligned} n_1 &= \sum_{i < j} [Y]_{ij} && \text{number of edges} \\ n_2 &= \sum_{i < j < k} [Y]_{ik} [Y]_{jk} && \text{number of 2-stars} \\ n_3 &= \sum_{i < j < l < k} [Y]_{ik} [Y]_{jk} [Y]_{lk} && \text{number of 3-stars.} \end{aligned}$$

The remainder of this chapter will describe various MCMC methods that target the posterior distribution $\pi(\boldsymbol{\theta} \mid \mathbf{y})$, or some approximation thereof. This will be in the context of a random walk Metropolis (RWM) algorithm that proposes a new value of $\boldsymbol{\theta}'$ at iteration t using a (multivariate) Gaussian proposal distribution, $q(\cdot \mid \boldsymbol{\theta}^{t-1}) \sim \mathcal{N}(\boldsymbol{\theta}^{t-1}, \Sigma_t)$. Methods for tuning the proposal bandwidth Σ_t have been described by Andrieu and Thoms [3] and Roberts and Rosenthal [67]. Normally, the proposed parameter value would be accepted with probability $\min\{1, \rho_t\}$, or else rejected, where ρ_t is the Radon–Nikodým derivative:

$$\rho_t = \frac{q(\boldsymbol{\theta}^{(t-1)} \mid \boldsymbol{\theta}') p(\mathbf{y} \mid \boldsymbol{\theta}') \pi_0(\boldsymbol{\theta}')}{q(\boldsymbol{\theta}' \mid \boldsymbol{\theta}^{(t-1)}) p(\mathbf{y} \mid \boldsymbol{\theta}^{(t-1)}) \pi_0(\boldsymbol{\theta}^{(t-1)})}, \quad (6.6)$$

$\pi_0(\boldsymbol{\theta})$ is the prior density for the parameter/s, and $p(\mathbf{y} \mid \boldsymbol{\theta})$ is the likelihood (6.1). If we use a symmetric proposal distribution q and a uniform prior π_0 , then these terms

will cancel, leaving:

$$\rho_t = \frac{\psi(\mathbf{y} | \boldsymbol{\theta}')}{\psi(\mathbf{y} | \boldsymbol{\theta}^{(t-1)})} \frac{C(\boldsymbol{\theta}^{(t-1)})}{C(\boldsymbol{\theta}')}, \quad (6.7)$$

which is the ratio of unnormalised likelihoods $\psi = \exp\{\boldsymbol{\theta}^T \mathbf{s}(\mathbf{y})\}$, multiplied by the ratio of intractable normalising constants (6.2). It is clearly infeasible to evaluate (6.7) directly, so alternative algorithms are required. One option is to estimate ρ_t by simulation, which we categorise as auxiliary variable methods: pseudo-marginal algorithms, the exchange algorithm, and approximate Bayesian computation (ABC). Other methods include path sampling, also known as thermodynamic integration (TI), pseudolikelihood, and composite likelihood.

6.1 Auxiliary Variable Methods

6.1.1 Pseudo-Marginal Algorithms

Pseudo-marginal algorithms [2, 6] are computational methods for fitting latent variable models, that is where the observed data \mathbf{y} can be considered as noisy observations of some unobserved or hidden states, \mathbf{x} . For example, hidden Markov models (HMMs) are commonly used in time series analysis and signal processing. Models of this form can also arise as the result of data augmentation approaches, such as for mixture models [17, 73]. The marginal likelihood is of the following form:

$$p(\mathbf{y} | \boldsymbol{\theta}) = \int_{\mathcal{X}} p(\mathbf{y} | \mathbf{x}) p(\mathbf{x} | \boldsymbol{\theta}) d\mathbf{x}, \quad (6.8)$$

which can be intractable if the state space is very high-dimensional and non-Gaussian. In this case, we can substitute an unbiased, non-negative estimate of the likelihood.

O'Neill et al. [60] introduced the Monte Carlo within Metropolis (MCWM) algorithm, which replaces both $p(\mathbf{y} | \boldsymbol{\theta}')$ and $p(\mathbf{y} | \boldsymbol{\theta}^{(t-1)})$ in the Metropolis–Hastings ratio ρ_t (6.6) with importance-sampling estimates:

$$\tilde{\rho}_{IS}(\mathbf{y} | \boldsymbol{\theta}) \approx \frac{1}{M} \sum_{m=1}^M p(\mathbf{y} | X_m) \frac{p(X_m | \boldsymbol{\theta})}{q(X_m | \boldsymbol{\theta})}, \quad (6.9)$$

where the samples X_1, \dots, X_M are drawn from a proposal distribution $q(X_m | \boldsymbol{\theta})$ for $\boldsymbol{\theta}'$ and $\boldsymbol{\theta}^{(t-1)}$. MCWM is generally considered as an approximate algorithm, since it does not target the exact posterior distribution for $\boldsymbol{\theta}$. However,

Medina-Aguayo et al. [47] have established some conditions under which MCWM converges to the correct target distribution as $M \rightarrow \infty$. See also [55] and [1] for further theoretical analysis of approximate pseudo-marginal methods.

Beaumont [6] introduced the grouped independence Metropolis–Hastings (GIMH) algorithm, which does target the exact posterior. The key difference is that $\tilde{p}_{IS}(\mathbf{y} \mid \boldsymbol{\theta}^{(t-1)})$ is reused from the previous iteration, rather than being recalculated every time. The theoretical properties of this algorithm have been an active area of research, with notable contributions by Andrieu and Roberts [2], Maire et al. [41], Andrieu and Vihola [4], and Sherlock et al. [69]. Andrieu et al. [5] introduced the particle MCMC algorithm, which is a pseudo-marginal method that uses sequential Monte Carlo (SMC) in place of importance sampling. This is particularly useful for HMMs, where SMC methods such as the bootstrap particle filter provide an unbiased estimate of the marginal likelihood [62]. Although importance sampling and SMC are both unbiased estimators, it is necessary to use a large enough value of M so that the variance is kept at a reasonable level. Otherwise, the pseudo-marginal algorithm can fail to be variance-bounding or geometrically ergodic [39]. Doucet et al. [18] recommend choosing M so that the standard deviation of the log-likelihood estimator is between 1 and 1.7.

Pseudo-marginal algorithms can be computationally intensive, particularly for large values of M . One strategy to reduce this computational burden, known as the Russian Roulette algorithm [40], is to replace $\tilde{p}_{IS}(\mathbf{y} \mid \boldsymbol{\theta})$ (6.9) with a truncated infinite series:

$$\tilde{p}_{RR}(\mathbf{y} \mid \boldsymbol{\theta}) = \sum_{j=0}^{\tau} V_{\boldsymbol{\theta}}^{(j)}, \quad (6.10)$$

where τ is a random stopping time and $V_{\boldsymbol{\theta}}^{(j)}$ are random variables such that (6.10) is almost surely finite and $\mathbb{E}[\tilde{p}_{RR}(\mathbf{y} \mid \boldsymbol{\theta})] = p(\mathbf{y} \mid \boldsymbol{\theta})$. There is a difficulty with this method, however, in that the likelihood estimates are not guaranteed to be non-negative. Jacob and Thiery [37] have established that there is no general solution to this sign problem, although successful strategies have been proposed for some specific models.

Another important class of algorithms for accelerating pseudo-marginal methods involve approximating the intractable likelihood function using a surrogate model. For example, the delayed-acceptance (DA) algorithm of [14] first evaluates the Metropolis–Hastings ratio (6.6) using a fast, approximate likelihood $\tilde{p}_{DA}(\mathbf{y} \mid \boldsymbol{\theta})$. The proposal $\boldsymbol{\theta}'$ is rejected at this screening stage with probability $1 - \min\{1, \rho_t\}$. Otherwise, a second ratio $\rho_{DA}^{(2)}$ is calculated using a full evaluation of the likelihood function (6.9). The acceptance probability $\min\{1, \rho_{DA}^{(2)}\}$ is modified at the second stage according to:

$$\rho_{DA}^{(2)} = \frac{\tilde{p}_{IS}(\mathbf{y} \mid \boldsymbol{\theta}') \pi_0(\boldsymbol{\theta}')}{\tilde{p}_{IS}(\mathbf{y} \mid \boldsymbol{\theta}^{(t-1)}) \pi_0(\boldsymbol{\theta}^{(t-1)})} \frac{\tilde{p}_{DA}(\mathbf{y} \mid \boldsymbol{\theta}^{(t-1)}) \pi_0(\boldsymbol{\theta}^{(t-1)})}{\tilde{p}_{DA}(\mathbf{y} \mid \boldsymbol{\theta}') \pi_0(\boldsymbol{\theta}')}, \quad (6.11)$$

which corrects for the conditional dependence on acceptance at the first stage and therefore preserves the exact target distribution. DA has been used for PMCMC by Golightly et al. [33], where the linear noise approximation [25] was used for \tilde{p}_{DA} . Sherlock et al. [70] instead used k -nearest-neighbours for \tilde{p}_{DA} in a pseudo-marginal algorithm.

Drovandi et al. [22] proposed an approximate pseudo-marginal algorithm, using a Gaussian process (GP) as a surrogate log-likelihood. The GP is trained using a pilot run of MCWM, then at each iteration $\log \tilde{p}(\mathbf{y} \mid \boldsymbol{\theta}')$ is either approximated using the GP or else using SMC or importance sampling, depending on the level of uncertainty in the surrogate model for $\boldsymbol{\theta}'$. MATLAB source code is available from <http://www.runmycode.org/companion/view/2663>. Stuart and Teckentrup [71] have shown that, under certain assumptions, a GP provides a consistent estimator of the negative log-likelihood, and they provide error bounds on the approximation.

6.1.2 Exchange Algorithm

Møller et al. [50] introduced a MCMC algorithm for the Ising model that targets the exact posterior distribution for β . An auxiliary variable \mathbf{x} is defined on the same state space as \mathbf{y} , so that $\mathbf{x}, \mathbf{y} \in \mathcal{Y}$. This is a data augmentation approach, where we simulate from the joint posterior $\pi(\beta, \mathbf{x} \mid \mathbf{y})$, which admits the posterior for β as its marginal. Given a proposed parameter value β' , a proposal \mathbf{x}' is simulated from the model to obtain an unbiased sample from (6.1). This requires perfect simulation methods, such as coupling from the past [65], perfect slice sampling [49], or bounding chains [8, 35]. Refer to [36] for further explanation of perfect simulation. Instead of (6.7), the joint ratio for β' and \mathbf{x}' becomes:

$$\rho_t = \frac{\psi(\mathbf{y} \mid \beta')}{\psi(\mathbf{y} \mid \beta^{(t-1)})} \frac{\psi(\mathbf{x}' \mid \tilde{\beta})}{\psi(\mathbf{x}^{(t-1)} \mid \tilde{\beta})} \frac{\psi(\mathbf{x}^{(t-1)} \mid \beta^{(t-1)})}{\psi(\mathbf{x}' \mid \beta')}, \quad (6.12)$$

where the normalising constants $C(\beta')$ and $C(\beta^{(t-1)})$ cancel out with each other. This is analogous to an importance-sampling estimate of the normalising constant with $M = 1$ samples, since:

$$\mathbb{E}_{\mathbf{x}} \left[\frac{\psi(\mathbf{x} \mid \beta)}{q(\mathbf{x} \mid \beta)} \right] = C(\beta), \quad (6.13)$$

where the proposal distribution $q(\mathbf{x} \mid \beta)$ is (6.1). This algorithm is therefore closely-related with pseudo-marginal methods such as GIMH.

Murray et al. [54] found that (6.12) could be simplified even further, removing the need for a fixed value of $\hat{\beta}$. The exchange algorithm replaces (6.7) with the ratio:

$$\rho_t = \frac{\psi(\mathbf{y} \mid \beta')}{\psi(\mathbf{y} \mid \beta^{(t-1)})} \frac{\psi(\mathbf{x}' \mid \beta^{(t-1)})}{\psi(\mathbf{x}' \mid \beta')}. \quad (6.14)$$

However, perfect sampling is still required to simulate \mathbf{x}' at each iteration, which can be infeasible when the state space is very large. Cucala et al. [15] proposed an approximate exchange algorithm (AEA) by replacing the perfect sampling step with 500 iterations of Gibbs sampling. Caimo and Friel [9] were the first to employ AEA for fully-Bayesian inference on the parameters of an ERGM. AEA for the hidden Potts model is implemented in the R package ‘bayesImages’ [51] and AEA for ERGM is implemented in ‘Bergm’ [10].

6.1.3 Approximate Bayesian Computation

Like the exchange algorithm, ABC uses an auxiliary variable \mathbf{x} to decide whether to accept or reject the proposed value of θ' . In the terminology of ABC, \mathbf{x} is referred to as “pseudo-data.” Instead of a Metropolis–Hastings ratio such as (6.7), the summary statistics of the pseudo-data and the observed data are directly compared. The proposal is accepted if the distance between these summary statistics is within the ABC tolerance, ϵ . This produces the following approximation:

$$p(\theta \mid \mathbf{y}) \approx \pi_\epsilon(\theta \mid \|\mathbf{s}(\mathbf{x}) - \mathbf{s}(\mathbf{y})\| < \epsilon), \quad (6.15)$$

where $\|\cdot\|$ is a suitable norm, such as Euclidean distance. Since $\mathbf{s}(\mathbf{y})$ are jointly-sufficient statistics for Ising, Potts, or ERGM, the ABC approximation (6.15) approaches the true posterior as $n \rightarrow \infty$ and $\epsilon \rightarrow 0$. In practice there is a tradeoff between the number of parameter values that are accepted and the size of the ABC tolerance.

Grelaud et al. [34] were the first to use ABC to obtain an approximate posterior for β in the Ising/Potts model. Everitt [24] used ABC within sequential Monte Carlo (ABC-SMC) for Ising and ERGM. ABC-SMC uses a sequence of target distributions $\pi_{\epsilon_t}(\theta \mid \|\mathbf{s}(\mathbf{x}) - \mathbf{s}(\mathbf{y})\| < \epsilon_t)$ such that $\epsilon_1 > \epsilon_2 > \dots > \epsilon_T$, where the number of SMC iterations T can be determined dynamically using a stopping rule. The ABC-SMC algorithm of [19] uses multiple MCMC steps for each SMC iteration, while the algorithm of [16] uses multiple replicates of the summary statistics for each particle. Everitt [24] has provided a MATLAB implementation of ABC-SMC with the online supplementary material accompanying his paper.

The computational efficiency of ABC is dominated by the cost of drawing updates to the auxiliary variable, as reported by Everitt [24]. Thus, we would expect that the execution time for ABC would be similar to AEA or pseudo-

marginal methods. Various approaches to improving this runtime have recently been proposed. “Lazy ABC” [63] involves early termination of the simulation step at a random stopping time, hence it bears some similarities with Russian Roulette. Surrogate models have also been applied in ABC, using a method known as Bayesian indirect likelihood [BIL; 20, 21]. Gaussian processes (GPs) have been used as surrogate models by Wilkinson [75] and Meeds and Welling [48]. Järvenpää et al. [38] used a heteroskedastic GP model and demonstrated how the output of the precomputation step could be used for Bayesian model choice. Moores et al. [52] introduced a piecewise linear approximation for ABC-SMC with Ising/Potts models. Boland et al. [7] derived a theoretical upper bound on the bias introduced by this and similar piecewise approximations. They also developed a piecewise linear approximation for ERGM. Moores et al. [53] introduced a parametric functional approximate Bayesian (PFAB) algorithm for the Potts model, which is a form of BIL where $\tilde{p}_{BIL}(\mathbf{y} | \boldsymbol{\theta})$ is derived from an integral curve.

6.2 Other Methods

6.2.1 Thermodynamic Integration

Since the Ising, Potts, and ERGM are all exponential families of distributions, the expectation of their sufficient statistic/s can be expressed in terms of the normalising constant:

$$\mathbb{E}_{\mathbf{y}|\boldsymbol{\theta}}[\mathbf{s}(\mathbf{y})] = \frac{d}{d\boldsymbol{\theta}} \log\{C(\boldsymbol{\theta})\}. \quad (6.16)$$

Gelman and Meng [31] derived an approximation to the log-ratio of normalising constants for the Ising/Potts model, using the path sampling identity:

$$\log \left\{ \frac{C(\beta_{t-1})}{C(\beta')} \right\} = \int_{\beta'}^{\beta_{t-1}} \mathbb{E}_{\mathbf{y}|\beta}[s(\mathbf{y})] d\beta, \quad (6.17)$$

which follows from (6.16). The value of the expectation can be estimated by simulating from the Gibbs distribution (6.1) for fixed values of β . At each iteration, $\log\{\rho_t\}$ (6.7) can then be approximated by numerical integration methods, such as Gaussian quadrature or the trapezoidal rule. Figure 6.1 illustrates linear interpolation of $\mathbb{E}_{\mathbf{y}|\beta}[s(\mathbf{y})]$ on a 2D lattice for $q = 6$ labels and β ranging from 0 to 2 in increments of 0.05. This approximation was precomputed using the algorithm of [72].

TI is explained in further detail by Chen et al. [13, chap. 5]. A reference implementation in \mathbb{R} is available from the website accompanying [42]. Friel and Pettitt [29] introduced the method of power posteriors to estimate the marginal likelihood or model evidence using TI. Calderhead and Girolami [11] provide bounds on the discretisation error and derive an optimal temperature schedule by




Fig. 6.1 Approximation of $\mathbb{E}_{\mathbf{y}|\beta}[s(\mathbf{y})]$ by simulation for fixed values of β , with linear interpolation

minimising the variance of the Monte Carlo estimate. Oates et al. [56] introduced control variates for further reducing the variance of TI.

The TI algorithm has an advantage over auxiliary variable methods because the additional simulations are performed prior to fitting the model, rather than at each iteration. This is particularly the case when analysing multiple images that all have approximately the same dimensions. Since these simulations are independent, they can make use of massively parallel hardware. However, the computational cost is still slightly higher than pseudolikelihood, which does not require a pre-computation step.

6.2.2 Composite Likelihood

Pseudolikelihood is the simplest of the methods that we have considered and also the fastest. Rydén and Titterton [68] showed that the intractable distribution (6.1) could be approximated using the product of the conditional densities:

$$\tilde{p}_{PL}(\mathbf{y} \mid \boldsymbol{\theta}) \approx \prod_{i=1}^n p(y_i \mid y_{\setminus i}, \boldsymbol{\theta}). \tag{6.18}$$

This enables the Metropolis–Hastings ratio ρ_t (6.6) to be evaluated using (6.18) to approximate both $p(\mathbf{y} \mid \boldsymbol{\theta}')$ and $p(\mathbf{y} \mid \boldsymbol{\theta}^{(t-1)})$ at each iteration. The conditional

density function for the Ising/Potts model is given by:

$$p(y_i | y_{\setminus i}, \beta) = \frac{\exp \left\{ \beta \sum_{\ell \in \partial(i)} \delta(z_i, z_\ell) \right\}}{\sum_{j=1}^q \exp \left\{ \beta \sum_{\ell \in \partial(i)} \delta(j, z_\ell) \right\}}, \quad (6.19)$$

where $\ell \in \partial(i)$ are the first-order (nearest) neighbours of pixel i . The conditional density for an ERGM is given by the logistic function:

$$p([Y]_{ij} = 1 | [Y]_{\setminus ij}, \theta) = \text{logit}^{-1} \left\{ \theta^T \mathbf{s}(Y) \right\}. \quad (6.20)$$

Pseudolikelihood is exact when $\theta = 0$ and provides a reasonable approximation for small values of the parameters. However, the approximation error increases rapidly for the Potts/Ising model as β approaches the critical temperature, β_{crit} , as illustrated by Fig. 6.2. This is due to long-range dependence between the labels, which is inadequately modelled by the local approximation. Similar issues can arise for ERGM, which can also exhibit a phase transition.

Rydén and Titterington [68] referred to Eq. (6.18) as point pseudolikelihood, since the conditional distributions are computed for each pixel individually. They suggested that the accuracy could be improved using block pseudolikelihood. This is where the likelihood is calculated exactly for small blocks of pixels, then (6.18) is modified to be the product of the blocks:

$$\tilde{p}_{BL}(\mathbf{y} | \theta) \approx \prod_{i=1}^{N_B} p(\mathbf{y}_{B_i} | \mathbf{y}_{\setminus B_i}, \theta) \quad (6.21)$$

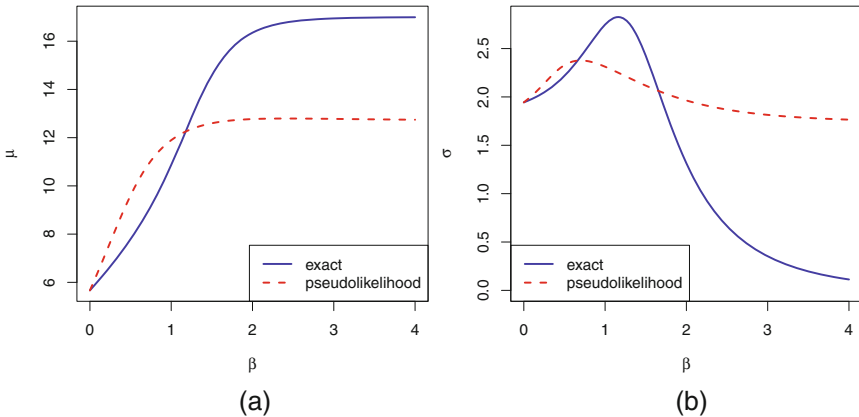


Fig. 6.2 Approximation error of pseudolikelihood for $n = 12$, $q = 3$ in comparison to the exact likelihood calculated using a brute force method: (a) $\sum_{\mathbf{y} \in \mathcal{Y}} s(\mathbf{y}) p(\mathbf{y} | \beta)$ using either Eq. (6.1) or (6.18); (b) $\sqrt{\sum_{\mathbf{y} \in \mathcal{Y}} (s(\mathbf{y}) - \mathbb{E}_{\mathbf{y} | \beta}[s(\mathbf{y})])^2 p(\mathbf{y} | \beta)}$. (a) Expectation. (b) Standard deviation

where N_B is the number of blocks, \mathbf{y}_{B_i} are the labels of the pixels in block B_i , and $\mathbf{y}_{\setminus B_i}$ are all of the labels except for \mathbf{y}_{B_i} . This is a form of composite likelihood, where the likelihood function is approximated as a product of simplified factors [74]. Friel [27] compared point pseudolikelihood to composite likelihood with blocks of 3×3 , 4×4 , 5×5 , and 6×6 pixels. Friel showed that (6.21) outperformed (6.18) for the Ising ($q = 2$) model with $\beta < \beta_{crit}$. Okabayashi et al. [58] discuss composite likelihood for the Potts model with $q > 2$ and have provided an open source implementation in the R package ‘potts’ [32].

Evaluating the conditional likelihood in (6.21) involves the normalising constant for \mathbf{y}_{B_i} , which is a sum over all of the possible configurations \mathcal{Y}_{B_i} . This is a limiting factor on the size of blocks that can be used. The brute force method that was used to compute Fig. 6.2 is too computationally intensive for this purpose. Pettitt et al. [61] showed that the normalising constant can be calculated exactly for a cylindrical lattice by computing eigenvalues of a $k^r \times k^r$ matrix, where r is the smaller of the number of rows or columns. The value of (6.2) for a free-boundary lattice can then be approximated using path sampling. Friel and Pettitt [28] extended this method to larger lattices using a composite likelihood approach.

The reduced dependence approximation (RDA) is another form of composite likelihood. Reeves and Pettitt [66] introduced a recursive algorithm to calculate the normalising constant using a lag- r representation. Friel et al. [30] divided the image lattice into sub-lattices of size $r_1 < r$, then approximated the normalising constant of the full lattice using RDA:

$$C(\beta) \approx \frac{C_{r_1 \times n}(\beta)^{r-r_1+1}}{C_{r_1-1 \times n}(\beta)^{r-r_1}} \quad (6.22)$$

McGrory et al. [44] compared RDA to pseudolikelihood and the exact method of [50], reporting similar computational cost to pseudolikelihood but with improved accuracy in estimating β . Ogden [57] showed that if r is chosen proportional to n , then RDA gives asymptotically valid inference when $\beta < \beta_{crit}$. However, the error increases exponentially as β approaches the phase transition. This is similar to the behaviour of pseudolikelihood in Fig. 6.2. Source code for RDA is available in the online supplementary material for McGrory et al. [45].

6.3 Conclusion

This chapter has surveyed a variety of computational methods for Bayesian inference with intractable likelihoods. Auxiliary variable methods, such as the exchange algorithm and pseudo-marginal algorithms, target the exact posterior distribution. However, their computational cost can be prohibitive for large datasets. Algorithms such as delayed acceptance, Russian Roulette, and “lazy ABC” can accelerate inference by reducing the number of auxiliary variables that need to be simulated, without modifying the target distribution. Bayesian indirect likelihood (BIL) algo-

rithms approximate the intractable likelihood using a surrogate model, such as a Gaussian process or piecewise function. As with thermodynamic integration, BIL can take advantage of a precomputation step to train the surrogate model in parallel. This enables these methods to be applied to much larger datasets by managing the tradeoff between approximation error and computational cost.

Acknowledgments This research was conducted by the Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (project number CE140100049) and funded by the Australian Government.

References

1. P. Alquier, N. Friel, R. Everitt, A. Boland, Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Stat. Comput.* **26**(1–2), 29–47 (2016). <https://doi.org/10.1007/s11222-014-9521-x>
2. C. Andrieu, G.O. Roberts, The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.* **37**(2), 697–725 (2009). <https://doi.org/10.1214/07-AOS574>
3. C. Andrieu, J. Thoms, A tutorial on adaptive MCMC. *Stat. Comput.* **18**(4), 343–373 (2008). <https://doi.org/10.1007/s11222-008-9110-y>
4. C. Andrieu, M. Vihola, Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *Ann. Appl. Prob.* **25**(2), 1030–1077, 04 (2015). <https://doi.org/10.1214/14-AAP1022>
5. C. Andrieu, A. Doucet, R. Holenstein, Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B* **72**(3), 269–342 (2010). <https://doi.org/10.1111/j.1467-9868.2009.00736.x>
6. M.A. Beaumont, Estimation of population growth or decline in genetically monitored populations. *Genetics* **164**(3), 1139–1160 (2003)
7. A. Boland, N. Friel, F. Maire, Efficient MCMC for Gibbs random fields using pre-computation. *Electron. J. Statist.* **12**(2), 4138–4179 (2018). <https://doi.org/10.1214/18-EJS1504>
8. C.T. Butts, A perfect sampling method for exponential family random graph models. *J. Math. Soc.* **42**(1), 17–36 (2018). <https://doi.org/10.1080/0022250X.2017.1396985>
9. A. Caimo, N. Friel, Bayesian inference for exponential random graph models. *Soc. Networks* **33**(1), 41–55 (2011). <https://doi.org/10.1016/j.socnet.2010.09.004>
10. A. Caimo, N. Friel, Bergm: Bayesian exponential random graphs in R. *J. Stat. Soft.* **61**(2), 1–25 (2014). <https://doi.org/10.18637/jss.v061.i02>
11. B. Calderhead, M. Girolami, Estimating Bayes factors via thermodynamic integration and population MCMC. *Comput. Stat. Data Anal.* **53**(12), 4028–4045 (2009). <https://doi.org/10.1016/j.csda.2009.07.025>
12. E. Cameron, A.N. Pettitt, Approximate Bayesian computation for astronomical model analysis: a case study in galaxy demographics and morphological transformation at high redshift. *Mon. Not. R. Astron. Soc.* **425**(1), 44–65 (2012). <https://doi.org/10.1111/j.1365-2966.2012.21371.x>
13. M.-H. Chen, Q.-M. Shao, J.G. Ibrahim, *Monte Carlo Methods in Bayesian Computation*. Springer Series in Statistics (Springer, New York, 2000)
14. J.A. Christen, C. Fox, Markov chain Monte Carlo using an approximation. *J. Comput. Graph. Stat.* **14**(4), 795–810 (2005). <https://doi.org/10.1198/106186005X76983>
15. L. Cucala, J.-M. Marin, C.P. Robert, D.M. Titterton, A Bayesian reassessment of nearest-neighbor classification. *J. Am. Stat. Assoc.* **104**(485), 263–273 (2009). <https://doi.org/10.1198/jasa.2009.0125>

16. P. Del Moral, A. Doucet, A. Jasra, An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Stat. Comput.* **22**(5), 1009–20 (2012). <https://doi.org/10.1007/s11222-011-9271-y>
17. A.P. Dempster, N.M. Laird, D.B. Rubin, Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B* **39**(1), 1–38 (1977).
18. A. Doucet, M. Pitt, G. Deligiannidis, R. Kohn, Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika* **102**(2), 295–313 (2015). <https://doi.org/10.1093/biomet/asu075>
19. C.C. Drovandi, A.N. Pettitt, Estimation of parameters for macroparasite population evolution using approximate Bayesian computation. *Biometrics* **67**(1), 225–233 (2011). <https://doi.org/10.1111/j.1541-0420.2010.01410.x>
20. C.C. Drovandi, A.N. Pettitt, M.J. Faddy, Approximate Bayesian computation using indirect inference. *J. R. Stat. Soc. Ser. C* **60**(3), 317–337 (2011). <https://doi.org/10.1111/j.1467-9876.2010.00747.x>
21. C.C. Drovandi, A.N. Pettitt, A. Lee, Bayesian indirect inference using a parametric auxiliary model. *Stat. Sci.* **30**(1), 72–95 (2015). <https://doi.org/10.1214/14-STS498>
22. C.C. Drovandi, M.T. Moores, R.J. Boys, Accelerating pseudo-marginal MCMC using Gaussian processes. *Comput. Stat. Data Anal.* **118**, 1–17 (2018). <https://doi.org/10.1016/j.csda.2017.09.002>
23. P. Erdős, A. Rényi, On random graphs. *Publ. Math. Debr.* **6**, 290–297 (1959)
24. R.G. Everitt, Bayesian parameter estimation for latent Markov random fields and social networks. *J. Comput. Graph. Stat.* **21**(4), 940–960 (2012). <https://doi.org/10.1080/10618600.2012.687493>
25. P. Fearnhead, V. Giagos, C. Sherlock, Inference for reaction networks using the linear noise approximation. *Biometrics* **70**(2), 457–466 (2014). <https://doi.org/10.1111/biom.12152>
26. O. Frank, D. Strauss, Markov graphs. *J. Amer. Stat. Assoc.* **81**(395), 832–842 (1986)
27. N. Friel, Bayesian inference for Gibbs random fields using composite likelihoods, in ed. by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, A.M. Uhrmacher, *Proceedings of the 2012 Winter Simulation Conference (WSC)* (2012), pp. 1–8. <https://doi.org/10.1109/WSC.2012.6465236>
28. N. Friel, A.N. Pettitt, Likelihood estimation and inference for the autologistic model. *J. Comp. Graph. Stat.* **13**(1), 232–246 (2004). <https://doi.org/10.1198/1061860043029>
29. N. Friel, A.N. Pettitt, Marginal likelihood estimation via power posteriors. *J. R. Stat. Soc. Ser. B* **70**(3), 589–607 (2008). <https://doi.org/10.1111/j.1467-9868.2007.00650.x>
30. N. Friel, A.N. Pettitt, R. Reeves, E. Wit, Bayesian inference in hidden Markov random fields for binary data defined on large lattices. *J. Comp. Graph. Stat.* **18**(2), 243–261 (2009). <https://doi.org/10.1198/jcgs.2009.06148>
31. A. Gelman, X.-L. Meng, Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statist. Sci.* **13**(2), 163–185 (1998). <https://doi.org/10.1214/ss/1028905934>
32. C.J. Geyer, L. Johnson, *potts: Markov Chain Monte Carlo for Potts Models*. R package version 0.5-2 (2014). <http://CRAN.R-project.org/package=potts>
33. A. Golightly, D.A. Henderson, C. Sherlock, Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Stat. Comput.* **25**(5), 1039–1055 (2015). <https://doi.org/10.1007/s11222-014-9469-x>
34. A. Grelaud, C.P. Robert, J.-M. Marin, F. Rodolphe, J.-F. Taly, ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Anal.* **4**(2), 317–336 (2009). <https://doi.org/10.1214/09-BA412>
35. M.L. Huber, A bounding chain for Swendsen-Wang. *Random Struct. Algor.* **22**(1), 43–59 (2003). <https://doi.org/10.1002/rsa.10071>
36. M.L. Huber, *Perfect Simulation* (Chapman & Hall/CRC Press, London/Boca Raton, 2016)
37. P.E. Jacob, A.H. Thiery, On nonnegative unbiased estimators. *Ann. Statist.* **43**(2), 769–784 (2015). <https://doi.org/10.1214/15-AOS1311>

38. M. Järvenpää, M. Gutmann, A. Vehtari, P. Marttinen, Gaussian process modeling in approximate Bayesian computation to estimate horizontal gene transfer in bacteria. *Ann. Appl. Stat.* **12**(4), 2228–2251 (2018). <https://doi.org/10.1214/18-AOAS1150>
39. A. Lee, K. Łatuszyński, Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for approximate Bayesian computation. *Biometrika* **101**(3), 655–671 (2014). <https://doi.org/10.1093/biomet/asu027>
40. A.-M. Lyne, M. Girolami, Y. Atchadé, H. Strathmann, D. Simpson, On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statist. Sci.* **30**(4), 443–467 (2015). <https://doi.org/10.1214/15-STS523>
41. F. Maire, R. Douc, J. Olsson, Comparison of asymptotic variances of inhomogeneous Markov chains with application to Markov chain Monte Carlo methods. *Ann. Statist.* **42**(4), 1483–1510, 08 (2014). <https://doi.org/10.1214/14-AOS1209>
42. J.-M. Marin, C.P. Robert, *Bayesian Core: A Practical Approach to Computational Bayesian Statistics*. Springer Texts in Statistics (Springer, New York, 2007)
43. P. Marjoram, J. Molitor, V. Plagnol, S. Tavaré, Markov chain Monte Carlo without likelihoods. *Proc. Natl Acad. Sci.* **100**(26), 15324–15328 (2003). <https://doi.org/10.1073/pnas.0306899100>
44. C.A. McGrory, D.M. Titterton, R. Reeves, A.N. Pettitt, Variational Bayes for estimating the parameters of a hidden Potts model. *Stat. Comput.* **19**(3), 329–340 (2009). <https://doi.org/10.1007/s11222-008-9095-6>
45. C.A. McGrory, A.N. Pettitt, R. Reeves, M. Griffin, M. Dwyer, Variational Bayes and the reduced dependence approximation for the autologistic model on an irregular grid with applications. *J. Comput. Graph. Stat.* **21**(3), 781–796 (2012). <https://doi.org/10.1080/10618600.2012.632232>
46. T.J. McKinley, I. Vernon, I. Andrianakis, N. McCreesh, J.E. Oakley, R.N. Nsubuga, M. Goldstein, R.G. White, et al., Approximate Bayesian computation and simulation-based inference for complex stochastic epidemic models. *Statist. Sci.* **33**(1), 4–18 (2018). <https://doi.org/10.1214/17-STS618>
47. F.J. Medina-Aguayo, A. Lee, G.O. Roberts, Stability of noisy Metropolis-Hastings. *Stat. Comput.* **26**(6), 1187–1211 (2016). <https://doi.org/10.1007/s11222-015-9604-3>
48. E. Meeds, M. Welling, GPS-ABC: Gaussian process surrogate approximate Bayesian computation, in *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, Quebec City, Canada (2014)
49. A. Mira, J. Møller, G.O. Roberts, Perfect slice samplers. *J. R. Stat. Soc. Ser. B* **63**(3), 593–606 (2001). <https://doi.org/10.1111/1467-9868.00301>
50. J. Møller, A.N. Pettitt, R. Reeves, K.K. Berthelsen, An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika* **93**(2), 451–458 (2006). <https://doi.org/10.1093/biomet/93.2.451>
51. M.T. Moores, D. Feng, K. Mengersen, *bayesImageS: Bayesian Methods for Image Segmentation Using a Potts Model*. R package version 0.5-3 (2014). URL <http://CRAN.R-project.org/package=bayesImageS>
52. M.T. Moores, C.C. Drovandi, K. Mengersen, C.P. Robert, Pre-processing for approximate Bayesian computation in image analysis. *Stat. Comput.* **25**(1), 23–33 (2015). <https://doi.org/10.1007/s11222-014-9525-6>
53. M.T. Moores, G.K. Nicholls, A.N. Pettitt, K. Mengersen, Scalable Bayesian inference for the inverse temperature of a hidden Potts model. *Bayesian Anal.* **15**, 1–27 (2020). <https://doi.org/10.1214/18-BA1130>
54. I. Murray, Z. Ghahramani, D.J.C. MacKay, MCMC for doubly-intractable distributions, in *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, Arlington (AUAI Press, Tel Aviv-Yafo, 2006), pp. 359–366
55. G.K. Nicholls, C. Fox, A. Muir Watt, Coupled MCMC with a randomized acceptance probability (2012). Preprint arXiv:1205.6857 [stat.CO]. <https://arxiv.org/abs/1205.6857>
56. C.J. Oates, T. Papamarkou, M. Girolami, The controlled thermodynamic integral for Bayesian model evidence evaluation. *J. Am. Stat. Assoc.* **111**(514), 634–645 (2016). <https://doi.org/10.1080/01621459.2015.1021006>

57. H.E. Ogden, On asymptotic validity of naive inference with an approximate likelihood. *Biometrika* **104**(1), 153–164 (2017). <https://doi.org/10.1093/biomet/asx002>
58. S. Okabayashi, L. Johnson, C.J. Geyer, Extending pseudo-likelihood for Potts models. *Statistica Sinica* **21**, 331–347 (2011)
59. E. Olbrich, T. Kahle, N. Bertschinger, N. Ay, J. Jost, Quantifying structure in networks. *Eur. Phys. J. B* **77**(2), 239–247 (2010). <https://doi.org/10.1140/epjb/e2010-00209-0>
60. P.D. O’Neill, D.J. Balding, N.G. Becker, M. Eerola, D. Mollison, Analyses of infectious disease data from household outbreaks by Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. C* **49**(4), 517–542 (2000). <https://doi.org/10.1111/1467-9876.00210>
61. A.N. Pettitt, N. Friel, R. Reeves, Efficient calculation of the normalizing constant of the autologistic and related models on the cylinder and lattice. *J. R. Stat. Soc. Ser. B* **65**(1), 235–246 (2003). <https://doi.org/10.1111/1467-9868.00383>
62. M.K. Pitt, R. dos Santos Silva, P. Giordani, R. Kohn, On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *J. Econometr.* **171**(2), 134–151 (2012). <https://doi.org/10.1016/j.jeconom.2012.06.004>
63. D. Prangle, Lazy ABC. *Stat. Comput.* **26**(1), 171–185 (2016). <https://doi.org/10.1007/s11222-014-9544-3>
64. J.K. Pritchard, M.T. Seielstad, A. Perez-Lezaun, M.W. Feldman, Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Mol. Biol. Evol.* **16**(12), 1791–1798 (1999). <https://doi.org/10.1093/oxfordjournals.molbev.a026091>
65. J.G. Propp, D.B. Wilson, Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algor.* **9**(1–2), 223–252 (1996). [https://doi.org/10.1002/\(SICI\)1098-2418\(199608/09\)9:1/2<223::AID-RSA14>3.0.CO;2-O](https://doi.org/10.1002/(SICI)1098-2418(199608/09)9:1/2<223::AID-RSA14>3.0.CO;2-O)
66. R. Reeves, A.N. Pettitt, Efficient recursions for general factorisable models. *Biometrika* **91**(3), 751–757 (2004). <https://doi.org/10.1093/biomet/91.3.751>
67. G.O. Roberts, J.S. Rosenthal, Examples of adaptive MCMC. *J. Comput. Graph. Stat.* **18**(2), 349–367 (2009). <https://doi.org/10.1198/jcgs.2009.06134>
68. T. Rydén, D.M. Titterton, Computational Bayesian analysis of hidden Markov models. *J. Comput. Graph. Stat.* **7**(2), 194–211 (1998). <https://doi.org/10.1080/10618600.1998.10474770>
69. C. Sherlock, A.H. Thiery, G.O. Roberts, J.S. Rosenthal, On the efficiency of pseudo-marginal random walk Metropolis algorithms. *Ann. Statist.* **43**(1), 238–275, 02 (2015). <https://doi.org/10.1214/14-AOS1278>
70. C. Sherlock, A. Golightly, D.A. Henderson, Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods. *J. Comput. Graph. Stat.* **26**(2), 434–444 (2017). <https://doi.org/10.1080/10618600.2016.1231064>
71. A.M. Stuart, A.L. Teckentrup, Posterior consistency for Gaussian process approximations of Bayesian posterior distributions. *Math. Comp.* **87**, 721–753 (2018). <https://doi.org/10.1090/mcom/3244>
72. R.H. Swendsen, J.-S. Wang, Nonuniversal critical dynamics in Monte Carlo simulations. *Phys. Rev. Lett.* **58**, 86–88 (1987). <https://doi.org/10.1103/PhysRevLett.58.86>
73. M.A. Tanner, W.H. Wong, The calculation of posterior distributions by data augmentation. *J. Am. Stat. Assoc.* **82**(398), 528–40 (1987)
74. C. Varin, N. Reid, D. Firth, An overview of composite likelihood methods. *Statistica Sinica* **21**, 5–42 (2011)
75. R.D. Wilkinson, Accelerating ABC methods using Gaussian processes, in ed. by S. Kaski, J. Corander, *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics AISTATS (JMLR: Workshop and Conference Proceedings)*, vol. 33 (2014), pp. 1015–1023