

# Bayesian computation for statistical models with intractable normalizing constants

*Yves F. Atchadé<sup>\*</sup>, Nicolas Lartillot<sup>†</sup> and Christian P. Robert<sup>‡</sup>*

(April 2008)

**Abstract:** This paper deals with some computational aspects in the Bayesian analysis of statistical models with intractable normalizing constants. In the presence of intractable normalizing constants in the likelihood function, traditional MCMC methods cannot be applied. We propose an approach to sample from such posterior distributions. The method can be thought as a Bayesian version of the MCMC-MLE approach of [8]. To the best of our knowledge, this is the first general and asymptotically consistent Monte Carlo method for such problems. We illustrate the method with examples from image segmentation and social network modeling. We study as well the asymptotic behavior of the algorithm and obtain a strong law of large numbers for empirical averages.

**AMS 2000 subject classifications:** Primary 60C05, 60J27, 60J35, 65C40.

**Keywords and phrases:** Monte Carlo methods, Adaptive MCMC, Bayesian inference, Ising model, Image segmentation, Social network modeling.

## 1. Introduction

Statistical inference for models with intractable normalizing constants poses a major computational challenge. This problem occurs in the statistical modeling of many scientific problems. Examples include the analysis of spatial point processes ([13]), image analysis ([10]), protein design ([11]) and many others. The problem can be described as follows. Suppose we have a dataset  $x_0 \in (\mathcal{X}, \mathcal{B})$  generated from a statistical model  $e^{E(x, \theta)} \lambda(dx) / Z(\theta)$  with parameter  $\theta \in (\Theta, \Xi)$ , where the normalizing constant  $Z(\theta) = \int_{\mathcal{X}} e^{E(x, \theta)} \lambda(dx)$  depends on  $\theta$  and is not available in closed form. Let  $\mu$  be the prior density of the parameter  $\theta \in (\Theta, \Xi)$ . The posterior distribution of

---

<sup>\*</sup>Department of Statistics, University of Michigan, email: yvesa@umich.edu

<sup>†</sup>LIRM, Université de Montpellier 2, email: nicolas.lartillot@lirmm.fr

<sup>‡</sup>CEREMADE, Université Paris-Dauphine and CREST, INSEE, email: xian@ceremade.dauphine.fr

$\theta$  given  $x_0$  is then given by

$$\pi(\theta) \propto \frac{1}{Z(\theta)} e^{E(x_0, \theta)} \mu(\theta). \quad (1)$$

When  $Z(\theta)$  cannot be easily evaluated, Monte Carlo simulation from this posterior distribution is problematic even using Markov Chain Monte Carlo (MCMC). [14] uses the term *doubly intractable distribution* to refer to posterior distributions of the form (1). Current Monte Carlo sampling methods do not allow one to deal with such models in a Bayesian framework. For example, a Metropolis-Hastings algorithm with proposal kernel  $Q$  and target distribution  $\pi$ , would have acceptance ratio  $\min\left(1, \frac{e^{E(x_0, \theta')}}{e^{E(x_0, \theta)}} \frac{Z(\theta)}{Z(\theta')} \frac{Q(\theta', \theta)}{Q(\theta, \theta')}\right)$  which cannot be computed as it involves the intractable normalizing constant  $Z$  evaluated at  $\theta$  and  $\theta'$ .

An early attempt to deal with this problem is the pseudo-likelihood approximation of Besag ([4]) which approximates the model  $e^{E(x, \theta)}$  by a more tractable model. Pseudo-likelihood inference provides a first approximation but typically performs poorly (see e.g. [5]). Maximum likelihood inference is possible. MCMC-MLE, a maximum likelihood approach using MCMC has been developed in the 90's ([7, 8]). Another related approach to find MLE estimates is Younes' algorithm ([18]) based on stochastic approximation. An interesting simulation study comparing these three methods is presented in [10].

Comparatively little work has been done to develop asymptotically exact methods for the Bayesian approach to this problem. But various approximate algorithms exist in the literature, often based on path sampling ([6]). Recently, [12] have shown that if exact sampling of  $X$  from  $e^{E(x, \theta)}/Z(\theta)$  (as a density in  $(\mathcal{X}, \mathcal{B})$ ) is possible then a valid MCMC algorithm to sample from (1) can be developed. See also [14] for some improvements. Their approach uses a clever auxiliary variable algorithm. But intractable normalizing constants often occur in models for which exact sampling of  $X$  is not possible or is very expensive. Another recent development to the problem is the approximate Bayesian computation schemes of Plagnol-Tavaré ([15]) but which sample only approximately from the posterior distribution.

In this paper, we propose an adaptive Monte Carlo approach to sample from (1). Our algorithm generates a stochastic process (not Markov in general)  $\{(X_n, \theta_n), n \geq 0\}$  in  $\mathcal{X} \times \Theta$  such that as  $n \rightarrow \infty$ , the marginal distribution of  $\theta_n$  converges to (1). It is clear that any method to sample from (1) will have to deal with the intractable normalizing constant  $Z(\theta)$ . In the auxiliary variable method of [12], computing  $Z(\theta)$  is replaced in a sense by perfect sampling from  $e^{E(x, \theta)}/Z(\theta)$ . This

strategy works well so long as perfect sampling is feasible and inexpensive. In the present work, we take another approach building on the idea of estimating the entire function  $Z$  from a single Monte Carlo sampler. The starting point of the method is importance sampling. Suppose that for some  $\theta^{(0)} \in \Theta$ , we can sample (perhaps by MCMC) from the density  $e^{E(x, \theta^{(0)})}/Z(\theta^{(0)})$  in  $(\mathcal{X}, \mathcal{B})$ . Using this sample, we can certainly estimate  $Z(\theta)/Z(\theta^{(0)})$  for any  $\theta \in \Theta$ . This is the same idea behind the MCMC-MLE algorithm of [8]. But it is well known that these estimates are typically very poor as  $\theta$  gets far from  $\theta^{(0)}$ . Now, suppose that instead of a single point  $\theta^{(0)}$ , we generate a population  $\{\theta^{(i)}, i = 1, \dots, d\}$  in  $\Theta$  and that we can sample from  $\Lambda^*(x, i) \propto e^{E(x, \theta^{(i)})}/Z(\theta^{(i)})$  on  $\mathcal{X} \times \{1, \dots, d\}$ . Then we show that in principle, efficient estimation for  $Z(\theta)$  is possible for any  $\theta \in \Theta$ . Building on [3] and the ideas sketched above, we propose an algorithm that generates a random process  $\{(X_n, \theta_n), n \geq 0\}$  such that the marginal distribution of  $X_n$  converges to  $\Lambda^*$  and the marginal distribution of  $\theta_n$  converges to (1). This random process is not a Markov chain in general but we show (from first principle) that  $\{\theta_n\}$  has limiting distribution  $\pi$  and satisfies a strong law of large numbers.

The paper is organized as follows. A full description of the method including practical implementation details is given in Section 2. We illustrate the algorithm with three examples. The Ising model, a Bayesian image segmentation example and a Bayesian modeling of social networks. The examples are presented in Section 4. Some theoretical aspects of the method are discussed in Section 3 with the proofs postponed to 6.

## 2. Sampling from posterior distributions with intractable normalizing constants

Throughout, we fix the sample space  $(\mathcal{X}, \mathcal{B}, \lambda)$  and the parameter space  $(\Theta, \Xi)$ . The problem of interest is to sample from the posterior distribution (1) with

$$Z(\theta) = \int_{\mathcal{X}} e^{E(x, \theta)} \lambda(dx). \quad (2)$$

Let  $\{\theta^{(i)}, i = 1, \dots, d\}$  be a sequence of  $d$  points in  $\Theta$ . Let  $\Lambda^*$  be the probability measure on  $\mathcal{X} \times \{1, \dots, d\}$  given by:

$$\Lambda^*(x, i) = \frac{e^{E(x, \theta^{(i)})}}{dZ(\theta^{(i)})}, \quad x \in \mathcal{X}, i \in \{1, \dots, d\}. \quad (3)$$

Let  $\kappa(\theta, \theta')$  be a similarity kernel on  $\Theta \times \Theta$  such that  $\sum_{i=1}^d \kappa(\theta, \theta^{(i)}) = 1$  for all  $\theta \in \Theta$ . The starting point of the algorithm is the following decomposition of the partition function:

$$\begin{aligned}
Z(\theta) &= \int_{\mathcal{X}} e^{E(x, \theta)} \lambda(dx) \\
&= \sum_{i=1}^d \kappa(\theta, \theta^{(i)}) \int_{\mathcal{X}} e^{E(x, \theta) - E(x, \theta^{(i)})} e^{E(x, \theta^{(i)})} \lambda(dx) \\
&= d \sum_{i=1}^d \kappa(\theta, \theta^{(i)}) Z(\theta^{(i)}) \int_{\mathcal{X}} e^{E(x, \theta) - E(x, \theta^{(i)})} \frac{e^{E(x, \theta^{(i)})}}{dZ(\theta^{(i)})} \lambda(dx) \\
&= \sum_{i=1}^d \int_{\mathcal{X}} \Lambda^*(x, i) h_{\theta}(x, i) \lambda(dx),
\end{aligned} \tag{4}$$

where

$$h_{\theta}(x, i) = d \kappa(\theta, \theta^{(i)}) Z(\theta^{(i)}) e^{E(x, \theta) - E(x, \theta^{(i)})}. \tag{5}$$

The interest of the decomposition (4) is that  $\{Z(\theta^{(i)})\}$  and  $\Lambda^*$  do not depend on  $\theta$ . Therefore, using samples from  $\Lambda^*$ , this decomposition gives an approach to estimate  $Z(\theta)$  for all  $\theta \in \Theta$ . This estimate should be reliable provided  $\theta$  is close to at least one particle  $\theta^{(i)}$ . The problem of sampling from probability measures such as  $\Lambda^*$  has been recently considered by [3] building on the Wang-Landau algorithm of [17]. We follow and improve that approach here. The resulting estimate of  $Z(\theta)$  can then continuously be fed to a second Monte Carlo sampler that carries the simulation with respect to  $\pi$ . This suggests an adaptive Monte Carlo sampler to sample from (1) which we develop next.

For any  $c = (c(1), \dots, c(d)) \in \mathbb{R}^d$ , we define the following density function on  $\mathcal{X} \times \{1, \dots, d\}$ :

$$\Lambda_c(x, i) \propto e^{E(x, \theta^{(i)}) - c(i)}. \tag{6}$$

With  $c = \log(Z)$ ,  $\Lambda_c = \Lambda^*$ . The reader should think of  $c$  as an estimate of  $z$ , with  $z(i) := \log Z(\theta^{(i)})$ . The algorithm will adaptively adjust  $c$  such that the marginal distribution on  $\{1, \dots, d\}$  is approximately uniform. In which case, we should have  $c(i) = \log Z(i)$ . Let  $\{\gamma_n\}$  be a sequence of (possibly random) positive numbers. We propose a non-Markovian adaptive sampler that lives in  $\mathcal{X} \times \{1, \dots, d\} \times \mathbb{R}^d \times \Theta$ . We start from an initial state  $(X_0, I_0, c_0, \theta_0) \in \mathcal{X} \times \{1, \dots, d\} \times \mathbb{R}^d \times \Theta$ , where  $c_0 \in \mathbb{R}^d$  is the initial estimate of  $z$ . For example,  $c_0 \equiv 0$ . At time  $n$ , given  $(X_n, I_n, c_n, \theta_n)$  we first generate  $X_{n+1}$  from  $P_{I_n}(X_n, \cdot)$ , where  $P_i$  is a transition kernel on  $(\mathcal{X}, \mathcal{B})$  with invariant distribution  $e^{E(x, \theta^{(i)})} / Z(\theta^{(i)})$ . Next, we generate  $I_{n+1}$  from the distribution on  $\{1, \dots, d\}$  proportional

to  $e^{E(X_{n+1}, \theta^{(i)}) - c_n(i)}$ . Then we update the current estimate of  $\log(Z)$  to  $c_{n+1}$  given by:

$$c_{n+1}(i) = c_n(i) + \gamma_n \frac{e^{E(X_{n+1}, \theta^{(i)}) - c_n(i)}}{\sum_{j=1}^d e^{E(X_{n+1}, \theta^{(j)}) - c_n(j)}}, \quad i = 1, \dots, d. \quad (7)$$

In view of (4), we can estimate  $Z(\theta)$  by:

$$Z_{n+1}(\theta) = \sum_{i=1}^d \kappa(\theta, \theta^{(i)}) e^{c_{n+1}(i)} \left[ \frac{\sum_{k=1}^{n+1} e^{E(X_k, \theta) - E(X_k, \theta^{(i)})} \mathbf{1}_i(I_k)}{\sum_{k=1}^{n+1} \mathbf{1}_i(I_k)} \right], \quad (8)$$

with the convention that  $0/0 = 0$ . Finally, for any positive function  $\zeta : \Theta \rightarrow (0, \infty)$ , let  $Q_\zeta$  be a transition kernel on  $(\Theta, \Xi)$  with invariant distribution

$$\pi_\zeta(\theta) \propto \frac{1}{\zeta(\theta)} e^{E(x_0, \theta)} \mu(\theta). \quad (9)$$

Given  $(X_{n+1}, I_{n+1}, c_{n+1}, Z_{n+1}, \theta_n)$ , we generate  $\theta_{n+1}$  from  $Q_{Z_{n+1}}(\theta_n, \cdot)$ , where  $Z_{n+1}$  is the function defined by (8).

The algorithm can be summarized as follows.

**Algorithm 2.1.** . Let  $(X_0, I_0, c_0, \theta_0) \in \mathcal{X} \times \{1, \dots, d\} \times \mathbb{R}^d \times \Theta$  be the initial state of the algorithm. Let  $\{\gamma_n\}$  be (a possibly random) sequence of positive numbers. At time  $n$ , given  $(X_n, I_n, c_n, \theta_n)$ :

1. Generate  $X_{n+1}$  from  $P_{I_n}(X_n, \cdot)$  where  $P_i$  is any ergodic kernel on  $(\mathcal{X}, \mathcal{B})$  with invariant distribution  $e^{E(x, \theta^{(i)})} / Z(i)$ .
2. Generate  $I_{n+1}$  by sampling from the distribution on  $\{1, \dots, d\}$  proportional to  $e^{E(X_{n+1}, \theta^{(i)}) - c_n(i)}$ .
3. Compute  $c_{n+1}$ , the new estimate of  $g$  using (7).
4. Using the function  $Z_{n+1}$  defined by (8), generate  $\theta_{n+1}$  from  $Q_{Z_{n+1}}(\theta_n, \cdot)$ .

**Remark 2.1.** 1. The algorithm can be seen as an MCMC-MCMC analog to the MCMC-MLE of [8]. Indeed, with  $d = 1$ , the decomposition (4) becomes

$$Z(\theta) / Z(\theta^{(1)}) = \mathbb{E} \left[ e^{E(\theta, X) - E(X, \theta^{(1)})} \right],$$

where the expectation is taken with respect to the density  $e^{E(x, \theta^{(1)})} / Z(\theta^{(1)})$ . But as discussed in the introduction, when  $E(\theta, X) - E(X, \theta^{(1)})$  has a large variance, the resulting estimate is terribly poor.

2. We introduce  $\kappa$  to serve as a smoothing factor so that the particles  $\theta^{(i)}$ 's close to  $\theta$  contribute more to the estimation of  $Z(\theta)$ . We expect this smoothing step to reduce the variance of

the overall estimate of  $Z(\theta)$ . In the simulations we choose

$$\kappa(\theta, \theta^{(i)}) = \frac{e^{-\frac{1}{2h^2} \|\theta - \theta^{(i)}\|^2}}{\sum_{j=1}^d e^{-\frac{1}{2h^2} \|\theta - \theta^{(j)}\|^2}}.$$

The value of the smoothing parameter  $h$  is set by trials and errors for each example.

3. The implementation of the algorithm requires keeping track of all the samples  $X_k$  that are generated (Equation (8)).  $\mathcal{X}$  can be a very high-dimensional space and we are aware of the fact that in practice, this bookkeeping can significantly slow down the algorithm. But in many cases, the function  $E$  takes the form  $E(x, \theta) = \sum_{l=1}^K S_l(x) \theta_l$  for some real-valued functions  $S_l$ . In these cases, we only need to keep track of the statistics  $\{(S_1(X_n), \dots, S_K(X_n)), n \geq 0\}$ . All the examples discussed in the paper fall in this latter category.
4. As mentioned earlier, the update of  $(X_n, I_n, c_n)$  is essentially the Wang-Landau algorithm of [3] with the following important difference. [3] propose to update  $c_n$  one component per iteration:

$$c_{n+1}(i) = c_n(i) + \gamma_n \mathbf{1}_{\{i\}}(I_{n+1}).$$

We improve on this scheme in (7) by Rao-Blackwellization where we integrate out  $I_{n+1}$ .

5. As mentioned above, and we stress this again, this algorithm is not Markovian in any way. The process  $\{(X_n, I_n, c_n)\}$  is not a Markov chain but a nonhomogeneous Markov chain if we let  $\{\gamma_n\}$  be a deterministic sequence.  $\{\theta_n\}$ , the main random process of interest is not a Markov chain either. Nevertheless, the marginal distribution of  $\theta_n$  will typically converge to  $\pi$ . This is because,  $Q_{Z_n}$ , the conditional distribution of  $\theta_{n+1}$  given  $\mathcal{F}_n$  converges to  $Q_Z$  as  $n \rightarrow \infty$  and  $Q_Z$  is a kernel with invariant distribution  $\pi$ . We make this precise by showing that a strong law of large numbers holds for additive functionals of  $\{\theta_n\}$ .

We now discuss the choice of the parameters of the algorithm.

### 2.1. Choosing $d$ and the particles $\{\theta^{(i)}\}$

We do not have any general approach in choosing  $d$  and  $\{\theta^{(i)}\}$  but we give some guidelines. The general idea is that the particles  $\{\theta^{(i)}\}$  need to cover reasonably well the range of the density  $\pi$  and be such that for any  $\theta \in \Theta$ , the density  $e^{E(x, \theta)}/Z(\theta)$  in  $\mathcal{X}$  can be well approximated by at least one of the densities  $e^{E(x, \theta^{(i)})}/Z(i)$ . One possibility is to sample  $\theta^{(i)}$  independently from the

prior distribution  $\mu$ , some tempered version of it or some other similar distribution. We follow this approach in the examples below. Another possibility is to use a grid of points in  $\Theta$ . The value of  $d$ , the number of particles, should depend on the size of  $\Theta$ . We need to choose  $d$  such that the distributions  $e^{E(x, \theta^{(i)})}/Z(i)$  (in  $\mathcal{X}$ ) overlap. Otherwise, estimating the constants  $\{Z(i)\}$  can be difficult. This implies that  $d$  should not be too small. In the simulation examples below we choose  $d$  between 100 and 500.

## 2.2. Choosing the step-size $\{\gamma_n\}$

It is shown in [3] that the recursion (7) can also be written as a stochastic approximation algorithm with step-size  $\{\gamma_n\}$ , so that in theory, any positive sequence  $\{\gamma_n\}$  such that  $\sum \gamma_n = \infty$  and  $\sum \gamma_n^2 < \infty$  can be used. But the convergence of  $c_n$  to  $\log Z$  is very sensitive to the choice  $\{\gamma_n\}$ . If the  $\gamma_n$ 's are overly small, the recursive equation in (7) will make very small steps. But if these numbers are overly large, the algorithm will have a large variance. In both cases, the convergence to the solution will be slow. Overall, choosing the right step-size for a stochastic approximation algorithm is a difficult problem. Here we follow [3] which has elaborated on a heuristic approach to this problem originally proposed by [17].

The main idea of the method is that typically, the larger  $\gamma_n$ , the easier it is for the algorithm to move around the state space. Therefore, at the beginning  $\gamma_0$  is set at a relatively large value. This value is kept until  $\{I_n\}$  has visited equally well all the points of  $\{1, \dots, d\}$ . Let  $\tau_1$  be the first time where the occupation measure of  $\{1, \dots, d\}$  by  $\{I_n\}$  is approximately uniform. Then we set  $\gamma_{\tau_1+1}$  to some smaller value (for example  $\gamma_{\tau_1+1} = \gamma_{\tau_1}/2$ ) and the process is iterated until  $\gamma_n$  become sufficiently small. As which point we can choose to switch to a deterministic sequence of the form  $\gamma_n = n^{-1/2-\varepsilon}$ . Combining this idea with Algorithm 2.1, we get the following.

**Algorithm 2.2.** . Let  $\gamma > \varepsilon_1 > 0$ ,  $\varepsilon_2 > 0$  be constants and let  $(X_0, I_0, c_0, \theta_0)$  be some arbitrary initial state of the algorithm. Set  $v = 0 \in \mathbb{R}^d$  and  $n = 0$ . While  $\gamma > \varepsilon_1$  and given  $\mathcal{F}_n = \sigma\{(X_k, I_k, c_k, \theta_k), k \leq n\}$ ,

1. Generate  $(X_{n+1}, I_{n+1}, c_{n+1}, \theta_{n+1})$  as in Algorithm 2.1.
2. For  $i = 1, \dots, d$ : set  $v(i) = v(i) + \mathbf{1}_i(I_{n+1})$ .
3. If  $\max_i \left| v(i) - \frac{1}{d} \right| \leq \frac{\varepsilon_2}{d}$ , then set  $\gamma = \gamma/2$  and set  $v = 0 \in \mathbb{R}^d$ .

**Remark 2.2.** We use this algorithm in the examples below with the following specifications. We set the initial  $\gamma$  to 1,  $\varepsilon_2 = 0.2$ . We run  $\{(X_n, I_n, c_n)\}$  until  $\gamma \leq \varepsilon_1 = 0.001$  before starting  $\{\theta_n\}$  and switching to a deterministic sequence  $\gamma_n = \varepsilon_1/n^{0.7}$ .

### 3. Convergence analysis

In this section, we derive a law of large numbers under some verifiable conditions. The process of interest here is  $\{\theta_n\}$ . Let  $(\Omega, \mathcal{F}, \Pr)$  be the reference probability triplet. We equip  $(\Omega, \mathcal{F}, \Pr)$  with the filtration  $\{\mathcal{F}_n\}$ , where  $\mathcal{F}_n = \sigma\{(X_{k+1}, I_{k+1}, c_{k+1}, \theta_k), k \leq n\}$ . Note that  $\mathcal{F}_n$  includes  $(X_{n+1}, I_{n+1}, c_{n+1})$  since these random variables are used in generating  $\theta_{n+1}$ . From the definition of the algorithm, we have:

$$\Pr(\theta_{n+1} \in A | \mathcal{F}_n) = Q_{Z_{n+1}}(\theta_n, A), \quad \Pr - a.s.. \quad (10)$$

We see from (10) that  $\{\theta_n, \mathcal{F}_n\}$  is an adaptive Monte Carlo algorithm with varying target distribution. In analyzing  $\{\theta_n, \mathcal{F}_n\}$  here, we do not strive for the most general result but restrict ourselves to conditions that can be easily checked in the examples considered in the paper. We assume that  $\Theta$  is a compact subset of  $\mathbb{R}^q$ , the  $q$ -dimensional Euclidean space equipped with its Borel  $\sigma$ -algebra and the Lebesgue measure. Firstly, we assume that the function  $E$  is bounded:

**(A1):** *There exist  $m, M \in \mathbb{R}$  such that:*

$$m \leq E(x, \theta) \leq M, \quad x \in \mathcal{X}, \theta \in \Theta. \quad (11)$$

In many applications, and this is the case for the examples discussed below,  $\mathcal{X}$  is a finite set (typically very large) and  $\Theta$  is a compact set. In these cases, (A1) is easily checked. In order to proceed any further, we need some notations. A transition kernel on  $(\mathcal{X}, \mathcal{B})$  operates on measurable real-valued functions  $f$  as  $Pf(x) = \int P(x, dy)f(y)$ , and the product of two transition kernels  $P_1$  and  $P_2$  is the transition kernel defined as  $P_1P_2(x, A) = \int P_1(x, dy)P_2(y, A)$ . We can then define recursively  $P^n = PP^{n-1}$ ,  $n \geq 1$ ,  $P^0(x, A) = \mathbf{1}_A(x)$ . For two probability measures  $\mu, \nu$ , the total variation distance between  $\mu$  and  $\nu$  is defined as  $\|\mu - \nu\|_{TV} := \sup_A |\mu(A) - \nu(A)|$ . We say that a transition kernel  $P$  is geometrically ergodic if  $P$  is  $\phi$ -irreducible, aperiodic and has an invariant distribution  $\pi$  such that:

$$\|P^n(x, \cdot) - \pi\|_{TV} \leq M(x)\rho^n, \quad n \geq 0$$



for some  $\rho \in (0, 1)$  and some function  $M : \mathcal{X} \rightarrow (0, \infty]$ .

Our next assumption involves the transition kernel  $Q_\zeta$ .

**(A2):** For  $\zeta : \Theta \rightarrow (0, \infty)$ ,  $Q_\zeta$  is a Metropolis kernel with invariant distribution  $\pi_\zeta$  and proposal kernel density  $p$ . There exist  $\varepsilon > 0$  and an integer  $n_0 \geq 1$  such that for all  $\theta, \theta' \in \Theta$ :

$$p^{n_0}(\theta, \theta') \geq \varepsilon. \quad (12)$$

**Remark 3.1.** 1. The condition (12) clearly holds for most symmetric proposal kernels  $p(\theta, \theta')$ , provided that  $p(\theta, \theta')$  remains bounded away from 0 on some ball centered at  $\theta$ .  
2. (12) often implies that  $Q_\zeta$  is uniformly ergodic:

$$\begin{aligned} Q_\zeta(\theta, A) &\geq \int_A \min \left( 1, e^{E(x_0, \theta') - E(x_0, \theta)} \frac{\zeta(\theta')}{\zeta(\theta)} \right) p(\theta, \theta') d\theta' \\ &\geq e^{m-M} \inf_{\theta, \theta' \in \Theta} \left( \frac{\zeta(\theta)}{\zeta(\theta')} \right) \int_A p(\theta, \theta') d\theta'. \end{aligned}$$

Therefore, provided  $\inf_{\theta, \theta' \in \Theta} \left( \frac{\zeta(\theta)}{\zeta(\theta')} \right) > 0$ , if (12) hold then  $Q_\zeta^{n_0}(\theta, A) \geq \varepsilon' \mu_{Leb}(A)$  for some  $\varepsilon' > 0$ .

**(A3):**  $\{\gamma_n\}$  is a random sequence, adapted to  $\{\mathcal{F}_n\}$  such that  $\gamma_n > 0$ ,  $\sum \gamma_n = \infty$  and  $\sum \gamma_n^2 < \infty$  Pr-a.s.

**Theorem 3.1.** Assume (A1)-(A3). Assume also that each kernel  $P_i$  on  $(\mathcal{X}, \mathcal{B})$  is geometrically ergodic. Let  $h : (\Theta, \Sigma) \rightarrow \mathbb{R}$  such that  $|h| \leq 1$ . Then:

$$\frac{1}{n} \sum_{k=1}^n h(\theta_k) \rightarrow \pi(h), \text{ Pr-a.s.} \quad (13)$$

*Proof.* See Section 6. □

## 4. Examples

### 4.1. Ising model

We test the algorithm with the Ising model on a rectangular lattice. This is a simulated example. The model is given by  $e^{\theta E(x)} / Z(\theta)$  where

$$E(x) = \sum_{i=1}^m \sum_{j=1}^{n-1} x_{ij} x_{i,j+1} + \sum_{i=1}^{m-1} \sum_{j=1}^n x_{ij} x_{i+1,j}, \quad (14)$$

and  $x_i \in \{1, -1\}$ . We use  $m = n = 64$ . We generate the data  $x_0$  from  $e^{\theta E(x)}/Z(\theta)$  with  $\theta = 0.40$  by perfect sampling using the Propp-Wilson algorithm. Using  $x_0$  and postulating the model  $e^{\theta E(x)}/Z(\theta)$ , we would like to infer on  $\theta$ . We use the prior  $\mu(\theta) = \mathbf{1}_{(0,1)}(\theta)$ . We set  $d = 100$  and generate the points  $\{\theta^{(i)}\}$  independently and uniformly in  $(0, 1)$ . As described in Section 2.2, we use the flat histogram approach in selecting  $\{\gamma_n\}$  with an initial value  $\gamma_0 = 1$ , until  $\gamma_n$  becomes smaller than 0.001. Then we start feeding the adaptive chain  $\{\theta_n\}$  which is run for 10,000 iterations. In updating  $\theta_n$ , we use a Random Walk Metropolis sampler with proposal distribution  $\mathcal{U}(\theta_n - b, \theta_n + b)$  (with reflexion at the boundaries) for some  $b > 0$ . We adaptively update  $b$  so as to reach an acceptance rate of 30% (see e.g. [2]). We discard the first 1,999 points as a burn-in period. The results are plotted on Figure 1. As we can see from these plots, the sampler appears to have converged to the posterior distribution  $\pi$ . The mixing rate of the algorithm as inferred from the autocorrelation function seems fairly good. In addition, the algorithm yields an estimate of the partition function  $\log Z(\theta)$  which can be re-used in other sampling problems.

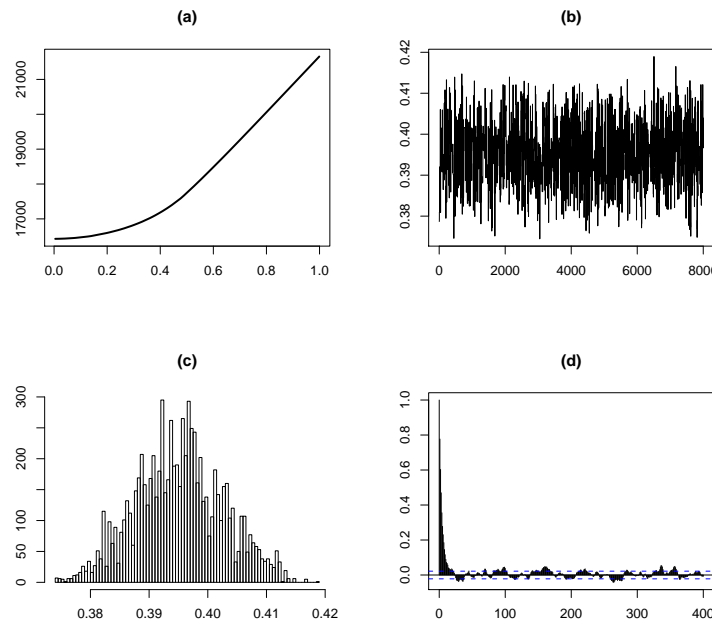


Figure 1: Output for the Ising model  $\theta = 0.40$ ,  $m = n = 64$ . (a): estimation of  $\log Z(\theta)$  up to an additive constant; (b)-(d): Trace plot, histogram and autocorrelation function of the adaptive sampler  $\{\theta_n\}$ .

#### 4.2. An application to image segmentation

We use the Ising model above to illustrate an application of the methodology to image segmentation. In image segmentation, the goal is the reconstruction of images from noisy observations (see e.g. [9, 10]). We represent the image by a vector  $x = \{x_i, i \in \mathcal{S}\}$  where  $\mathcal{S}$  is a  $m \times n$  lattice and  $x_i \in \{1, \dots, K\}$ . Each  $i \in \mathcal{S}$  represents a pixel, and  $x_i$  is the color of the pixel  $i$ .  $K$  is the number of colors. Here we assume that  $K = 2$  and  $x_i \in \{-1, 1\}$  is either black or white. In the image segmentation problem, we do not observe  $x$  but a noisy approximation  $y$ . We assume that:

$$y_i | x, \sigma^2 \stackrel{\text{ind}}{\sim} \mathcal{N}(x_i, \sigma^2), \quad (15)$$

for some unknown parameter  $\sigma^2$ . Even though (15) is a continuous model, it has been shown to provide a relatively good framework for image segmentation problems with multiple additive sources of noise ([10]).

We assume that the true image  $x$  is generated from an Ising model (see Section 4.1) with interaction parameter  $\theta$ . We assume that  $\theta$  follows a uniform prior distribution on  $(0, 1)$  and that  $\sigma^2$  has a prior distribution that is proportional to  $1/\sigma^2 \mathbf{1}_{(0, \infty)}(\sigma^2)$ . The posterior distribution  $(\theta, \sigma^2, x)$  is then given by:

$$\pi(\theta, \sigma^2, x | y) \propto \left(\frac{1}{\sigma^2}\right)^{\frac{|\mathcal{S}|}{2}+1} \frac{e^{\theta E(x)}}{Z(\theta)} e^{-\frac{1}{2\sigma^2} \sum_{s \in \mathcal{S}} (y(s) - x(s))^2} \mathbf{1}_{(0,1)}(\theta) \mathbf{1}_{(0,\infty)}(\sigma^2), \quad (16)$$

where  $E$  is as in (14).

We sample from this posterior distribution using the adaptive chain  $\{(y_n, i_n, c_n, \theta_n, \sigma_n^2, x_n)\}$ . The chain  $\{(y_n, i_n, c_n)\}$  is updated following Steps (1)-(3) of Algorithm 2.1. It is used to form the adaptive estimate of  $Z(\theta)$  as given by (8) (with  $\{y_n, i_n\}$  replacing  $\{X_n, I_n\}$ ). These estimates are used to update  $(\theta_n, \sigma_n^2, x_n)$  using a Metropolis-within-Gibbs scheme. More specifically, given  $\sigma_n^2, x_n$ , we sample  $\theta_{n+1}$  with a Random Walk Metropolis with proposal  $\mathcal{U}(\theta_n - b, \theta_n + b)$  (with reflexion at the boundaries) and target proportional to  $\frac{e^{\theta E(x_n)}}{Z_n(\theta)}$ . Given  $\theta_{n+1}, x_n$ , we generate  $\sigma_{n+1}^2$  by sampling from the Inverse-Gamma distribution with parameters  $(\frac{|\mathcal{S}|}{2}, \frac{1}{2} \sum_{s \in \mathcal{S}} (y(s) - x(s))^2)$ . And given  $(\theta_{n+1}, \sigma_{n+1}^2)$ , we sample each  $x_{n+1}(s)$  from its conditional distribution given  $\{x(u), u \neq s\}$ . This conditional distribution is given by

$$p(x(s) = a | x(u), u \neq s) \propto \exp\left(\theta a \sum_{u \sim s} x(u) - \frac{1}{2\sigma^2} (y(s) - a)^2\right), \quad a \in \{-1, 1\}.$$

Here  $u \sim v$  means that pixels  $u$  and  $v$  are neighbors.

To test this algorithm, we generate a simulated dataset  $y$  according to model (15) with  $x$  generated from  $e^{\theta E(x)}/Z(\theta)$  by perfect sampling. We use  $m = n = 64$ ,  $\theta = 0.40$  and  $\sigma = 0.5$ . For the implementation details of the algorithm, we make exactly the same choices as in Example 4.1 above. In particular we choose  $d = 100$  and generate  $\{\theta^{(i)}\}$  uniformly in  $(0, 1)$ . The results are given in Figure 2. Once again, the sample path obtained from  $\{\theta_n\}$  clearly suggests that the distribution of  $\theta_n$  has converged to  $\pi$  with a good mixing rate, as inferred from the autocorrelation plots.

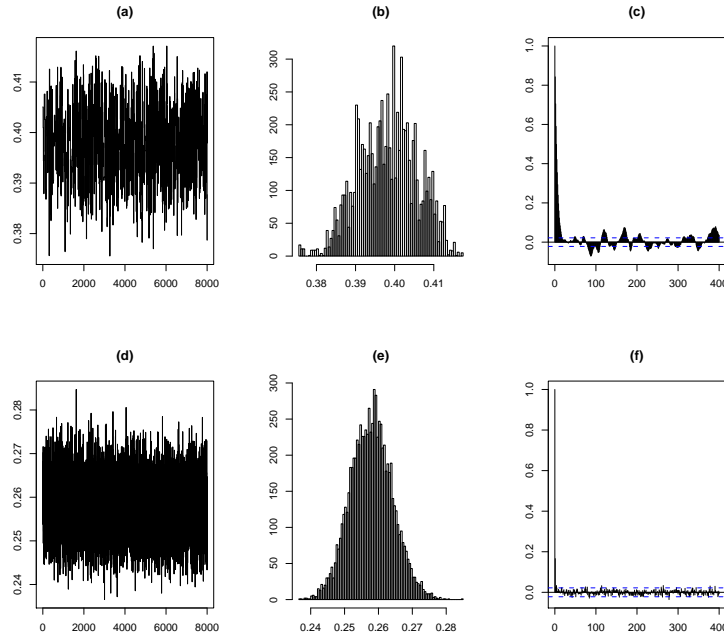


Figure 2: Output for the image segmentation model. (a)-(c): plots of  $\{\theta_n\}$ ; (d)-(f): plots of  $\{\sigma_n\}$ .

#### 4.3. Social network modeling

We now give an application of the method to a Bayesian analysis of social networks. Statistical modeling of social network is a growing subject in social science (See e.g. [16] and the references therein for more details). The set up is the following. We have  $n$  actors  $I = \{1, \dots, n\}$ . For each pair  $(i, j) \in I \times I$ , define  $y_{ij} = 1$  if actor  $i$  has ties with actor  $j$  and  $y_{ij} = 0$  otherwise. In the example below, we only consider the case of a symmetric relationship where  $y_{ij} = y_{ji}$  for all  $i, j$ . The ties referred to here can be of various natures. For example, we might be interested in

modeling how friendships build up between co-workers or how research collaboration takes place between colleagues. Another interesting example from political science is modeling co-sponsorship ties (for a given piece of legislation) between members of a house of representatives or parliament.

In this example we study the Medici business network dataset taken from [16] which describes the business ties between 16 Florentine families. Numbering arbitrarily the family from 1 to 16, we plot the observed social network in Figure 3. The dataset contains relatively few ties between families and even fewer transitive ties.




Figure 3: Business Relationships between 16 Florentine families.

One of the most popular models for social networks is the class of exponential random graph models. In these models, we assume that  $\{y_{ij}\}$  is a sample generated from the distribution

$$p(y|\theta_1, \dots, \theta_K) \propto \exp\left(\sum_{i=1}^K \theta_i S_i(y)\right), \quad (17)$$

for some parameters  $\theta_1, \dots, \theta_K$ ; where  $S_i(y)$  is a statistic used to capture some aspect of the network. For this example, and following [16], we consider a 4-dimensional model with statistics

$$\begin{aligned} S_1(y) &= \sum_{i < j} y_{ij}, & \text{the total number of ties,} \\ S_2(y) &= \sum_{i < j < k} y_{ik} y_{jk}, & \text{the number of two-stars,} \\ S_3(y) &= \sum_{i < j < k < l} y_{il} y_{jl} y_{kl}, & \text{the number of three-stars,} \end{aligned}$$

$$S_4(y) = \sum_{i < j < k} y_{ik} y_{jk} y_{ij}, \quad \text{the number of transitive ties.}$$

We assume a uniform prior distribution on  $D = (-50, 50)^4$  for  $\theta = (\theta_1, \theta_2, \theta_3, \theta_4)$  and the posterior distribution writes:

$$\pi(\theta|y) \propto \frac{1}{Z(\theta)} \exp\left(\sum_{k=1}^4 \theta_k S_k(y)\right) \mathbf{1}_D(\theta). \quad (18)$$

We use Algorithm 2.1 to sample from (18). For this example, we use 400 particles  $\{\theta^{(l)}\}$  generated from a  $N(0, 5)$  the normal distribution with mean 0 and variance 5. We use the same parametrization as in the previous examples to update  $(X_n, I_n, c_n)$ . For the adaptive chain  $\{\theta_n\}$  we use a slightly different strategy. It turns out that some of the components of the target distribution  $\pi$  are strongly related. Therefore we sample from  $\pi$  in one block, using a Random Walk Metropolis with a Gaussian kernel  $N(0, \sigma^2 \Sigma)$  (restricted to  $D$ ) for some  $\sigma > 0$  and a positive definite matrix  $\Sigma$ . We adaptively set  $\sigma$  so as to reach an acceptance rate of 30%. Ideally, we would like to choose  $\Sigma = \Sigma_\pi$  the variance-covariance of  $\pi$  which of course, is not available. We adaptively estimate  $\Sigma_\pi$  during the simulation as in [2]. As before, we run  $(X_n, I_n, c_n)$  until  $\gamma_n < 0.001$ . Then we start  $\{\theta_n\}$  and run the full chain  $(X_n, I_n, c_n, \theta_n)$  for a total of 25,000 iterations. The posterior distributions of the parameters are given in Figures 4a-4d. In Table 1, we give the sample posterior mean together with the 2.5% and 97.5% quantiles of the posterior distribution. Overall, these results are consistent with the maximum likelihood estimates obtained by [16] using MCMC-MLE. The main difference appears in  $\theta_4$  which we find here not to be significant. As a by-product, the sampler gives an estimate of the variance-covariance matrix of the posterior distribution  $\pi$ :

$$\Sigma_\pi = \begin{bmatrix} 1.67 & -0.41 & 0.27 & -0.07 \\ -0.41 & 1.83 & -0.47 & -0.02 \\ 0.27 & -0.47 & 1.78 & -0.03 \\ -0.07 & -0.02 & -0.03 & 1.65 \end{bmatrix}. \quad (19)$$

Parameters	Post. mean	Post. quantiles
$\theta_1$	-2.14	(-3.32, -0.81)
$\theta_2$	0.94	(-0.43, 2.49)
$\theta_3$	-1.06	(-2.72, 0.04)
$\theta_4$	0.09	(-1.39, 1.07)

TABLE 1

Summary of the posterior distribution of the parameters. Posterior means, 2.5% and 97.5% quantiles




Figure 4a: The adaptive MCMC output from (18). (a)-(c): Plots for  $\{\theta_1\}$ . Based on 25,000 iterations.




Figure 4b: The adaptive MCMC output from (18). (a)-(c): Plots for  $\{\theta_2\}$ . Based on 25,000 iterations.




Figure 4c: The adaptive MCMC output from (18). (a)-(c): Plots for  $\{\theta_3\}$ . Based on 25,000 iterations.




Figure 4d: The adaptive MCMC output from (18). (a)-(c): Plots for  $\{\theta_4\}$ . Based on 25,000 iterations.

## 5. Conclusion

Sampling from posterior distributions with intractable normalizing constants is a difficult computational problem. Thus far, all methods proposed in the literature but one entail approximations that do not vanish asymptotically. And the only exception ([12]) requires exact sampling in the data space, which is only possible for very specific cases. In this work, we propose an approach that both is more general than [12] and satisfies a strong law of large numbers with limiting distribution equal to the target distribution. The few applications we have presented here suggest that the method is promising. It remains to be determined how the method will scale with the dimensionality and with the size of the problems, although in this respect, adaptations of the method involving annealing/tempering schemes are easy to imagine, which would allow large problems to be analysed properly.

## Acknowledgements

The research of the third author had been partly supported by the Agence Nationale de la Recherche (ANR, 212, rue de Bercy 75012 Paris) through the 2006-2008 project Adap'MC.

## 6. Proof of Theorem 3.1

*Proof.* Throughout the proof,  $C$  will denote a finite constant but whose actual value can change from one equation to the next. The convergence of the Wang-Landau algorithm has been studied in [3]. It is shown in this work that under the condition of Theorem 3.1,  $\min_i \sum_{k=1}^{\infty} \mathbf{1}_{\{i\}}(I_k) = \infty$  and



more importantly,  $e^{c_n(i)} / \sum_{j=1}^d e^{c_n(j)}$  converges almost surely to a  $Z(\theta^{(i)})$  (up to a multiplicative constant).

Define

$$\omega_n(i) = \frac{e^{c_n(i)}}{\sum_{j=1}^d e^{c_n(j)}},$$

$$v_{n,i}(\theta) = \frac{\sum_{k=1}^n e^{E(X_k, \theta) - E(X_k, \theta^{(i)})} \mathbf{1}_i(I_k)}{\sum_{k=1}^n \mathbf{1}_i(I_k)}$$

and

$$\tilde{Z}_n(\theta) := \frac{Z_n(\theta)}{\sum_{j=1}^d e^{c_n(j)}} = \sum_{i=1}^d \omega_n(i) v_{n,i}(\theta). \quad (20)$$

Instead of  $Z_n$ , we work with  $\tilde{Z}_n$ . This is equivalent because  $\sum_{j=1}^d e^{c_n(j)}$  does not depend on  $\theta$  and  $Z_n$  always appears in  $Q_{Z_n}$  as a ratio. We have:

$$\inf_{\theta \in \Theta} \tilde{Z}_n(\theta) \geq e^{m-M}. \quad (21)$$

$$\inf_{\theta, \theta' \in \Theta} \left( \frac{\tilde{Z}_n(\theta)}{\tilde{Z}_n(\theta')} \right) \geq e^{2(m-M)}. \quad (22)$$

Combining (22) and (12) and part 2 of Remark 3.1, we deduce that there exists  $\varepsilon_0 > 0$  such that for all  $n, j \geq 0$

$$\sup_{|h| \leq 1} \left| Q_{Z_n}^j h(\theta) - \pi_{Z_n}(h) \right| \leq 2(1 - \varepsilon_0)^j, \quad \text{Pr-a.s.} \quad (23)$$

We introduce the notation  $\bar{Q}_n = Q_{\tilde{Z}_n} - \pi_{\tilde{Z}_n}$ . It follows from (23) that for any  $n \geq 1$  the following function  $g_n$  is well defined:

$$g_n(\theta) = \sum_{j=1}^{\infty} \bar{Q}_n^j h(\theta).$$

Moreover  $|g_n(\theta)| \leq 2/\varepsilon_0$  for all  $\theta \in \Theta$ .  $g_n$  satisfies Poisson's equation for  $\bar{Q}_n$  and  $h - \pi_{\tilde{Z}_n}(h)$ :

$$g_n(\theta) - \bar{Q}_n g_n(\theta) = h - \pi_{\tilde{Z}_n}(h). \quad (24)$$

Using this we can rewrite  $\sum_{k=1}^n h(\theta_k) - \pi_{Z_k}(h)$  as:

$$\begin{aligned} \frac{1}{n} \sum_{k=1}^n (h(\theta_k) - \pi_{Z_k}(h)) &= \frac{1}{n} \sum_{k=1}^n (g_k(\theta_k) - \bar{Q}_k g_k(\theta_{k-1})) + \frac{1}{n} \sum_{k=1}^n (\bar{Q}_k g_k(\theta_{k-1}) - \bar{Q}_{k-1} g_{k-1}(\theta_{k-1})) \\ &\quad + \frac{1}{n} (\bar{Q}_0 g_0(\theta_0) - \bar{Q}_n g_n(\theta_n)). \end{aligned} \quad (25)$$

Since  $\sup_{\theta \in \Theta} |g_n(\theta)| \leq 2/\varepsilon_0$ , a similar bound hold for  $\bar{Q}_n g_n$  and we conclude that  $\frac{1}{n} (\bar{Q}_0 g_0(\theta_0) - \bar{Q}_n g_n(\theta_n))$  actually converges almost surely to 0 as  $n \rightarrow \infty$ . Writing  $D_k = g_k(\theta_k) - \bar{Q}_k g_k(\theta_{k-1})$ , it is easily seen that  $\{D_k, \mathcal{F}_k\}$  is a martingale difference with bounded increment and we deduce from martingales theory that  $\frac{1}{n} \sum_{k=1}^n (g_k(\theta_k) - \bar{Q}_k g_k(\theta_{k-1}))$  converges almost surely to 0 as  $n \rightarrow \infty$ .

Since  $Q_Z$  is a Metropolis kernel and using the fact that  $|\min(1, ax) - \min(1, ay)| \leq a|x - y|$  for all  $a, x, y \geq 0$  we deduce that for any function  $h : \Theta \rightarrow \mathbb{R}$  such that  $|h| \leq 1$ ,

$$\begin{aligned} |(Q_{\tilde{Z}_n} - Q_{\tilde{Z}_{n-1}})h(\theta)| &\leq \int \left| \frac{\tilde{Z}_n(\theta)}{\tilde{Z}_n(\theta')} - \frac{\tilde{Z}_{n-1}(\theta)}{\tilde{Z}_{n-1}(\theta')} \right| e^{E(x_0, \theta') - E(x_0, \theta)} p(\theta, \theta') |h(\theta') - h(\theta)| d\theta' \\ &\leq 2e^{M-m} \sup_{\theta, \theta' \in \Theta} \left| \frac{\tilde{Z}_n(\theta)}{\tilde{Z}_n(\theta')} - \frac{\tilde{Z}_{n-1}(\theta)}{\tilde{Z}_{n-1}(\theta')} \right| \\ &\leq C |\tilde{Z}_n(\theta) - \tilde{Z}_{n-1}(\theta)|, \quad \text{using (21 - 22),} \end{aligned} \quad (26)$$

for some finite constant. Combining (23 and 26) we have the following well-known consequence: there exists  $C < \infty$  such that for all  $n \geq 1$ :

$$\sup_{|h| \leq 1} |\pi_{\tilde{Z}_n}(h) - \pi_{\tilde{Z}_{n-1}}(h)| \leq C \sup_{\theta \in \Theta} |\tilde{Z}_n(\theta) - \tilde{Z}_{n-1}(\theta)|. \quad (27)$$

The stability of Poisson's equation for geometrically ergodic transition kernels is well known (see e.g. [1, 3]). Combining (23), (26) and (27), we can find a finite constant  $C$  such that for all  $k \geq 1$ :

$$|(\bar{Q}_k g_k(\theta_{k-1}) - \bar{Q}_{k-1} g_{k-1}(\theta_{k-1}))| \leq C \sup_{\theta, \theta' \in \Theta} |\tilde{Z}_k(\theta) - \tilde{Z}_{k-1}(\theta)|. \quad (28)$$

Given the expression of  $\tilde{Z}_n(\theta)$  in (20) it is not very hard to show there exists  $C < \infty$  such that:

$$\sup_{\theta \in \Theta} |\tilde{Z}_k(\theta) - \tilde{Z}_{k-1}(\theta)| \leq C \left( d\gamma_k + \frac{1}{\min_i \sum_{l=1}^k \mathbf{1}_{\{i\}}(I_l)} \right) \rightarrow 0, \quad (29)$$

as  $k \rightarrow \infty$  as discussed above. It follows indeed that  $\frac{1}{n} \sum_{k=1}^n (h(\theta_k) - \pi_{Z_k}(h))$  converges a.s. to 0.

Given that  $\tilde{Z}_n(\theta) \rightarrow CZ(\theta)$  almost surely for some finite constant  $C$ ,

$$\pi_{Z_n}(h) = \frac{\int \frac{e^{E(\theta, x_0)}}{Z_n(\theta)} h(\theta) d\theta}{\int \frac{e^{E(\theta, x_0)}}{Z_n(\theta)} d\theta} \longrightarrow \frac{\int \frac{e^{E(\theta, x_0)}}{Z(\theta)} h(\theta) d\theta}{\int \frac{e^{E(\theta, x_0)}}{Z(\theta)} d\theta} = \pi(h),$$

as  $n \rightarrow \infty$  by Lebesgue's dominated convergence.  $\square$

## References

- [1] ANDRIEU, C. and MOULINES, É. (2006). On the ergodicity properties of some adaptive MCMC algorithms. *Ann. Appl. Probab.*, **16** 1462–1505.
- [2] ATCHADE, Y. F. (2006). An adaptive version for the Metropolis adjusted Langevin algorithm with a truncated drift. *Methodol Comput Appl Probab*, **8** 235–254.
- [3] ATCHADE, Y. F. and LIU, J. S. (2004). The Wang-Landau algorithm for Monte Carlo computation in general state spaces. *Technical Report*.
- [4] BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *J. Roy. Statist. Soc. Ser. B*, **36** 192–236. With discussion by D. R. Cox, A. G. Hawkes, P. Clifford, P. Whittle, K. Ord, R. Mead, J. M. Hammersley, and M. S. Bartlett and with a reply by the author.
- [5] CUCALA, L., MARIN, J.-M., ROBERT, C. and TITTERINGTON, D. (2008). A Bayesian reassessment of nearest-neighbour classification. Tech. rep., CEREMADE, Université Paris Dauphine. [arXiv:0802.1357](https://arxiv.org/abs/0802.1357).
- [6] GELMAN, A. and MENG, X.-L. (1998). Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statist. Sci.*, **13** 163–185.
- [7] GEYER, C. J. (1994). On the convergence of Monte Carlo maximum likelihood calculations. *J. Roy. Statist. Soc. Ser. B*, **56** 261–274.
- [8] GEYER, C. J. and THOMPSON, E. A. (1992). Constrained Monte Carlo maximum likelihood for dependent data. *J. Roy. Statist. Soc. Ser. B*, **54** 657–699. With discussion and a reply by the authors.
- [9] HURN, M., HUSBY, O. and RUE, H. (2003). A tutorial on image analysis. *Lecture notes in Statistics*, **173** 87–141.
- [10] IBANEZ, M. V. and SIMO, A. (2003). Parameter estimation in Markov random field image modeling with imperfect observations. a comparative study. *Pattern recognition letters*, **24** 2377–2389.
- [11] KLEINMAN, A., RODRIGUE, N., BONNARD, C. and PHILIPPE, H. (2006). A maximum likelihood framework for protein design. *BMC Bioinformatics*, **7**.
- [12] MØLLER, J., PETTITT, A. N., REEVES, R. and BERTHELSEN, K. K. (2006). An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants.

*Biometrika*, **93** 451–458.

- [13] MØLLER, J. and WAAGEPETERSEN, R. P. (2004). *Statistical inference and simulation for spatial point processes*, vol. 100 of *Monographs on Statistics and Applied Probability*. Chapman & Hall/CRC, Boca Raton, FL.
- [14] MURRAY, I., GHAHRAMANI, Z. and MACKAY, D. (2006). MCMC for doubly-intractable distributions. *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [15] PLAGNOL, V. and TAVARÉ, S. (2004). Approximate Bayesian computation and MCMC. In *Monte Carlo and quasi-Monte Carlo methods 2002*. Springer, Berlin, 99–113.
- [16] ROBINS, G., P., P., KALISH, Y. and LUSHER, D. (2007). An introduction to exponential random graph models for social networks. *Social Networks*, **29** 173–191.
- [17] WANG, F. and LANDAU, D. P. (2001). Efficient, multiple-range random walk algorithm to calculate the density of states. *Physical Review Letters*, **86** 2050–2053.
- [18] YOUNES, L. (1988). Estimation and annealing for Gibbsian fields. *Annales de l’Institut Henri Poincaré. Probabilité et Statistiques*, **24** 269–294.