



Title	Efficient MCMC for Gibbs Random Fields using pre-computation
Authors(s)	Boland, Aidan, Friel, Nial, Marie, Florian
Publication date	2018-05-31
Publication information	Boland, Aidan, Nial Friel, and Florian Marie. "Efficient MCMC for Gibbs Random Fields Using Pre-Computation." The Institute of Mathematical Statistics and the Bernoulli Society, May 31, 2018. https://doi.org/10.1214/18-EJS1504 .
Publisher	The Institute of Mathematical Statistics and the Bernoulli Society
Item record/more information	http://hdl.handle.net/10197/10866
Publisher's version (DOI)	10.1214/18-EJS1504

Downloaded 2026-05-01 23:37:21

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Efficient MCMC for Gibbs Random Fields using pre-computation

Aidan Boland^{1,2}, Nial Friel^{1,2}, and Florian Maire^{1,2}

¹School of Mathematics and Statistics, University College Dublin

²Insight Centre for Data Analytics

Abstract

Bayesian inference of Gibbs random fields (GRFs) is often referred to as a doubly intractable problem, since the likelihood function is intractable. The exploration of the posterior distribution of such models is typically carried out with a sophisticated Markov chain Monte Carlo (MCMC) method, the exchange algorithm (Murray et al., 2006), which requires simulations from the likelihood function at each iteration. The purpose of this paper is to consider an approach to dramatically reduce this computational overhead. To this end we introduce a novel class of algorithms which use realizations of the GRF model, simulated offline, at locations specified by a grid that spans the parameter space. This strategy speeds up dramatically the posterior inference, as illustrated on several examples. However, using the pre-computed graphs introduces a noise in the MCMC algorithm, which is no longer exact. We study the theoretical behaviour of the resulting approximate MCMC algorithm and derive convergence bounds using a recent theoretical development on approximate MCMC methods.

1 Introduction

The focus of this study is on Bayesian inference of Gibbs random fields (GRFs), a class of models used in many areas of statistics, such as the autologistic model (Besag, 1974) in spatial statistics, the exponential random graph model in social network analysis (Robins et al., 2007), etc. Unfortunately, for all but trivially small graphs, GRFs suffer from intractability of the likelihood function making standard analysis impossible. Such models are often referred to as *doubly-intractable* in the Bayesian literature, since the normalizing constant of both the likelihood function and the posterior distribution form a source of intractability. In the recent past there has been considerable research activity in designing Bayesian algorithms which overcome this intractability all of which rely on simulation from the intractable likelihood. Such methods include Approximate Bayesian Computation initiated by Pritchard et al. (1999) (see e.g. Marin et al. (2012)

for an excellent review) and Pseudo-Marginal algorithms (Andrieu and Roberts, 2009). Perhaps the most popular approach to infer a doubly-intractable posterior distribution is the exchange algorithm (Murray et al., 2006). The exchange algorithm is a Markov chain Monte Carlo (MCMC) method that extends the Metropolis-Hastings (MH) algorithm (Metropolis et al., 1953) to situations where the likelihood is intractable. Compared to MH, the exchange uses a different acceptance probability and this has two main implications:

- theoretically: the exchange chain is less efficient than the MH chain, in terms of mixing time and asymptotic variance (see Peskun (1973) and Tierney (1998) for a discussion on the optimality of the MH chain)
- computationally: at each iteration, the exchange requires *exact* and *independent* draws from the likelihood model at the current state of the Markov chain to calculate the acceptance probability, a step that may substantially impact upon the computational performance of the algorithm

For many likelihood models, it is not possible to simulate *exactly* from the likelihood function. In those situations, Cucala et al. (2009) and Caimo and Friel (2011) replace the exact sampling step in the exchange algorithm with the simulation of an auxiliary Markov chain targeting the likelihood function, whereby inducing a noise process in the main Markov chain. This approximation was extended further by Alquier et al. (2016) who used multiple samples to speed up the convergence of the exchange algorithm.

This short literature review of the exchange algorithm and its variants shows that simulations from the likelihood function, either exactly or approximately, is central to those methods. However, this simulation step often compromises their practical implementation, especially for large graph models. Indeed, for a realistic run time, a user may end up with a limited number of draws from the posterior as most of the computational budget is dedicated to obtaining likelihood realizations. In addition, note that since the likelihood draws are conditioned on the Markov chain states, those simulation steps are intrinsically incompatible with parallel computing (Friel et al., 2016).

Intuitively, there is a redundancy of simulation. Indeed, should the Markov chain return to an area previously visited, simulation of the likelihood is nevertheless carried out as it had never been done before. This is precisely the point we address in this paper. We propose a novel class of algorithms where likelihood realizations are generated and then subsequently re-used at in an online inference phase. More precisely, a regular grid spanning the parameter space is specified and draws from the likelihood at locations given by the vertices of this grid are obtained offline in a parallel fashion. The grid is tailored to the posterior topology using estimators of the gradient and the Hessian matrix to ensure that the pre-computation sampling covers the posterior areas of high probability. However, using realizations of the likelihood at pre-specified grid points instead of at the actual Markov chain state introduces a noise process in the algorithm. This leads us to study the theoretical behaviour of the resulting approximate MCMC algorithm and to derive quantitative convergence bounds using the noisy MCMC framework developed in Alquier et al. (2016). Essentially, our results allow one to quantify how the noise induced by the pre-computing step propagates through to

the stationary distribution of the approximate chain. We find an upper bound on the bias between this distribution and the posterior of interest, which depends on the pre-computing step parameters i.e. the distance between the grid points and the number of graphs drawn at each grid point. We also show that the bias vanishes asymptotically in the number of simulated graphs at each grid point, regardless of the grid structure.

Note that [Moore et al. \(2015\)](#) suggested a similar strategy to speed-up ABC algorithms by learning about the sufficient statistics of simulated data through an estimated mapping function that uses draws from the likelihood function at a pre-defined set of parameter values. This method was shown to be computationally very efficient but its suitability for models with more than one parameter can be questioned. Finally, we note that a related approach has been presented by [Everitt et al. \(2017\)](#) which also relies on previously sampled likelihood draws in order to estimate the intractable ratio of normalising constants. However this approach falls within a sequential Monte Carlo framework.

The paper is organised as follows. Section 2 introduces the intractable likelihood that we focus on and details our class of approximate MCMC schemes which uses pre-computed likelihood simulations. We also detail how we automatically specific the grid of parameter values. In Section 3, we establish some theoretical results for noisy MCMC algorithms making use of a pre-computation step. In Section 4, the inference of a number of GRFs is carried out using both pre-computed algorithms and exact algorithms such as the exchange. Results show a dramatic improvement of our method over exact methods in time normalized experiments. Finally, this paper concludes with some related open problems.

2 Pre-computing Metropolis algorithms

2.1 Preliminary notation

We frame our analysis in the setting of Gibbs random fields (GRFs) and we denote by $y \in \mathcal{Y}$ the observed graph. A graph is identified by its adjacency matrix and \mathcal{Y} is taken as $\mathcal{Y} := \{0, 1\}^{p \times p}$ where p is the number of nodes in the graph. The likelihood function of y is parameterized by a vector $\theta \in \Theta \subset \mathbb{R}^d$ and is defined as

$$f(y|\theta) = \frac{q_\theta(y)}{Z(\theta)} = \frac{\exp\{\theta^T s(y)\}}{Z(\theta)},$$

where $s(y) \in \mathcal{S} \subset \mathbb{R}_+^d$ is a vector of statistics which are sufficient for the likelihood. The normalizing constant,

$$Z(\theta) = \sum_{y \in \mathcal{Y}} \exp\{\theta^T s(y)\},$$

depends on θ and is intractable for all but trivially small graphs. The aim is to infer the parameters θ through the posterior distribution

$$\pi(\theta | y) \propto \frac{q_\theta(y)}{Z(\theta)} p(\theta),$$

where p denotes the prior distribution of θ . In absence of ambiguity, a distribution and its probability density function will share the same notation.

2.2 Computational complexity of MCMC algorithms for doubly intractable distributions

In Bayesian statistics, Markov chain Monte Carlo methods (MCMC, see *e.g.* [Gilks et al. \(1995\)](#) for an introduction) remain the most popular way to explore π . MCMC algorithms proceed by creating a Markov chain whose invariant distribution has a density equal to the posterior distribution. One such algorithm, the Metropolis-Hastings (MH) algorithm [Metropolis et al. \(1953\)](#), creates a Markov chain by sequentially drawing candidate parameters from a proposal distribution $\theta' \sim h(\cdot | \theta)$ and accepting the proposed new parameter θ' with probability

$$\alpha(\theta, \theta') := 1 \wedge a(\theta, \theta'), \quad a(\theta, \theta') := \frac{q_{\theta'}(y)p(\theta')h(\theta|\theta')}{q_{\theta}(y)p(\theta)h(\theta'|\theta)} \times \frac{Z(\theta)}{Z(\theta')}. \quad (1)$$

This acceptance probability depends on the ratio $Z(\theta)/Z(\theta')$ of the intractable normalising constants and cannot therefore be calculated in the case of GRFs. As a result, the MH algorithm cannot be implemented to infer GRFs.

As detailed in the introduction section, a number of variants of the MH algorithm bypass the need to calculate the ratio $Z(\theta)/Z(\theta')$, replacing it in Eq. (1) by an unbiased estimator

$$\varrho_n(\theta, \theta', x) = \frac{1}{n} \sum_{k=1}^n \frac{q_{\theta}(x_k)}{q_{\theta'}(x_k)}, \quad x_1, x_2, \dots \sim_{iid} f(\cdot | \theta'). \quad (2)$$

Perhaps surprisingly, when $n = 1$ the resulting algorithm, known as the *exchange* algorithm ([Murray et al., 2006](#)), is π -invariant. The general implementation using $n > 1$ auxiliary draws was proposed in [Alquier et al. \(2016\)](#) and referred therein as the *noisy exchange* algorithm. It is not π -invariant but the asymptotic bias in distribution was studied in ([Alquier et al., 2016](#)). We note however that when n is large, the resulting algorithm bears little resemblance with the exchange algorithm and really aims at approximating the MH acceptance ratio (1). For clarity, we will therefore refer to the exchange algorithm whenever $n = 1$ draw of the likelihood is needed at each iteration and to the noisy Metropolis-Hastings whenever $n > 1$.

From Eq. (2), we see that those modified MH algorithms crucially rely on the ability to sample efficiently from the likelihood distribution ($X \sim f(\cdot | \theta)$ for any $\theta \in \Theta$). While perfect sampling is possible for certain GRFs, for example for the Ising model ([Propp and Wilson, 1996](#)), it can be computationally expensive in some cases, including large Ising graphs. For some GRFs such as the exponential random graph model, perfect sampling does not even exist yet. [Cucala et al. \(2009\)](#) and [Caimo and Friel \(2011\)](#) substituted the iid sampling in Eq. (2) with $n = 1$ draw from a long auxiliary Markov chain that admits $f(\cdot | \theta)$ as stationary distribution. Convergence of this type of approximate exchange algorithm was studied in [Everitt \(2012\)](#) under certain assumptions on the main Markov chain. The computational bottleneck of those methods is clearly the simulation step, a drawback which is amplified when n is large and inference is on high-dimensional data such as large graphs.

Intuitively, obtaining a likelihood sample at each step independently of the past history of the chain seems to be an inefficient strategy. Indeed, the Markov chain may return to areas of the state space previously visited. As a result, realizations from the likelihood function are simulated at similar parameter values multiple times, throughout the algorithm. Under general assumptions on the likelihood function, data simulated at similar parameter values will share similar statistical features. Hence, repeated sampling without accounting for previous likelihood simulations seems to lead to an inefficient use of computational time. However, the price to pay to use information from the past history of the chain to speed up the simulation step is the loss of the Markovian dynamic of the chain, leading to a so-called adaptive Markov chain (see *e.g.* [Andrieu and Thoms \(2008\)](#)). We do not pursue this approach in this paper, essentially since convergence results for adaptive Markov chains depart significantly from the theoretical arguments supporting the validity of the exchange and its variants.

In a different context, [Moore et al. \(2015\)](#) addressed the computational expense of repeated simulations of Gibbs random fields used within an Approximate Bayesian Computation algorithm (ABC). The authors defined a pre-processing step designed to learn about the distribution of the summary statistics of simulated data. Part of the total computational budget is spent offline, simulating data from parameter values across the parameter space Θ . Those pre-simulated data are interpolated to create a mapping function $\Theta \rightarrow \mathcal{S}$ that is then used during the course of the ABC algorithm to assign an (estimated) sufficient statistics vector to any parameter θ for which simulation would be otherwise needed. [Moore et al. \(2015\)](#) examined a particular GRF, the single parameter hidden Potts model. They combined the pre-processing idea with path sampling ([Gelman and Meng, 1998](#)) to estimate the ratio of intractable normalising constants. The method presented in [Moore et al. \(2015\)](#) is suitable for single parameter models but the interpolation step remains a challenge when the dimension of the parameter space is greater than 1.

Inspired by the efficiency of a pre-computation step, we develop a novel class of MCMC algorithms, *Pre-computing Metropolis-Hastings*, which uses pre-computed data simulated offline to estimate each normalizing constant ratio $Z(\theta)/Z(\theta')$ in Eq. (1). This makes the extension to multi-parameter models straightforward. The steps undertaken during the pre-computing stage are now outlined.

2.3 Pre-computation step

Firstly, a set of parameter values, referred to as a grid, $\mathfrak{G} := (\hat{\theta}_1, \dots, \hat{\theta}_M)$ must be chosen from which to sample graphs from. \mathfrak{G} should cover the full state space and especially the areas of high probability of π . Finding areas of high probability is not straightforward as this requires knowledge of the posterior distribution. Fortunately, for GRFs we can use Monte Carlo methods to obtain estimates of the gradient and the Hessian matrix of the log posterior at different values of the parameters, which will allow to build a meaningful grid. For a GRF, the well known identity

$$\nabla_{\theta} \log \pi(\theta|y) = s(y) - \mathbb{E}_{f(\cdot|\theta)} s(X) + \nabla_{\theta} \log p(\theta)$$

allows the derivation of the following unbiased estimate of the gradient of the log posterior at a parameter $\theta \in \Theta$:

$$\mathcal{G}(\theta, y) := s(y) - \frac{1}{N} \sum_{i=1}^N s(X_i) + \nabla_{\theta} \log p(\theta), \quad X_1, X_2 \dots \sim_{iid} f(\cdot | \theta). \quad (3)$$

Similarly, the Hessian matrix of the log posterior at a parameter $\theta \in \Theta$ can be unbiasedly estimated by:

$$\mathcal{H}(\theta) := \frac{1}{N-1} \sum_{i=1}^N \{s(X_i) - \bar{s}\} \{s(X_i) - \bar{s}\}^T + \nabla^2 \log p(\theta),$$

$$X_1, X_2 \dots \sim_{iid} f(\cdot | \theta), \quad (4)$$

where \bar{s} is the average vector of simulated sufficient statistics.

The grid specification begins by estimating the mode of the posterior θ^* . This is achieved by mean of a stochastic approximation algorithm (*e.g.* the Robbins-Monro algorithm (Robbins and Monro, 1951)), using the log posterior gradient estimate \mathcal{G} defined at Eq. (3).

The second step is to estimate the Hessian matrix of the log posterior at θ^* using Eq. (4), in order to get an insight of the posterior curvature at the mode. We denote by $V := [v_1, \dots, v_d]$ the matrix whose columns are the eigenvectors v_i of the inverse Hessian at the mode and by $\Lambda := \text{diag}(\lambda_1, \dots, \lambda_d)$ the diagonal matrix filled with its eigenvalues. The idea is to construct a grid that preserves the correlations between the variables. It is achieved by taking regular steps in the uncorrelated space i.e. the space spanned by $[v_1, \dots, v_n]$, starting from θ^* and until subsequent estimated gradients are close to each other. The idea is that, for regular models, once the estimated gradients of two successive parameters are similar, the grid has hit the posterior distribution support boundary. Two tuning parameters are required: a threshold parameter for the gradient comparison $m > 0$ and an exploratory magnitude parameter $\varepsilon > 0$. The grid specification is rigorously outlined in Algorithm 1. Note that in Algorithm 1, we have used the notation δ_j for the d -dimensional indicator vector of direction j i.e. $\{\delta_j\}_{\ell} = \mathbb{1}_{j=\ell}$

The left panel of Figure 1 shows an example of a naively chosen grid built following standard coordinate directions for a two dimensional posterior distribution. The grid on the right hand side is adapted to the topology of the posterior distribution as described above. This method can be extended to higher dimensional models, but the number of sample grid points would then increase exponentially with dimension. In this paper we do not look beyond two dimensions.

Hereafter, we denote by $\{\theta_m, m \leq M\}$ the parameters constituting the grid \mathfrak{G} , assuming M grid points in total. The second step of the pre-computing step is to sample for each $\theta_m \in \mathfrak{G}$, n iid random variables (X_m^1, \dots, X_m^n) from the likelihood function $f(\cdot | \theta_m)$. Note that this step is easily parallelised and samples can therefore be obtained from several grid points simultaneously. Parallel processing can be used to reduce considerably the time taken to sample from every pre-computed grid value. Essentially, these draws allow to form unbiased estimators for any ratio of the type

Algorithm 1 Grid specification

- 1: **require** θ^* , V , Λ , m and ε .
 - 2: Initialise the grid with $\mathfrak{G} = \{\theta^*\}$
 - 3: **for** $i \in \{1, \dots, d\}$ **do**
 - 4: **for all** $\theta \in \mathfrak{G}$ **do**
 - 5: Set $j = 0$ and $\theta_0 = \theta$
 - 6: Calculate $\tilde{\theta} = \theta_0 + \varepsilon V \Lambda^{1/2} \delta_i$
 - 7: **while** $\|\mathcal{G}(\tilde{\theta}) - \mathcal{G}(\theta_j)\| > m$ **do**
 - 8: Set $j = j + 1$, $\theta_j = \tilde{\theta}$ and $\mathfrak{G} = \mathfrak{G} \cup \{\theta_j\}$
 - 9: Calculate $\tilde{\theta} = \theta_j + \varepsilon V \Lambda^{1/2} \delta_i$
 - 10: **end while**
 - 11: **end for**
 - 12: **end for**
 - 13: Obtain a second grid \mathfrak{G}' by repeating steps (2)–(12), but moving in the negative direction
i.e. $\tilde{\theta} = \theta - \varepsilon V \Lambda^{1/2} \delta_i$.
 - 14: **return** $\mathfrak{G} = \mathfrak{G} \cup \mathfrak{G}'$
-

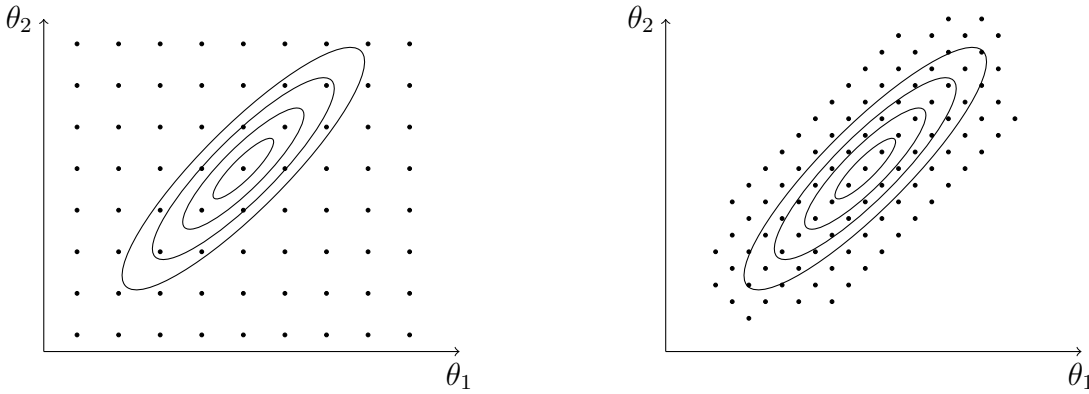


Figure 1: Example of a naive (left panel) and informed (right) grid for a two dimensional posterior distribution. The informed grid was obtained using the process described in Algorithm 1.

$Z(\theta)/Z(\dot{\theta}_m)$:

$$\frac{\widehat{Z(\theta)}}{Z(\dot{\theta}_m)_n} := \frac{1}{n} \sum_{k=1}^n \frac{q_\theta(X_m^k)}{q_{\dot{\theta}_m}(X_m^k)} = \frac{1}{n} \sum_{k=1}^n \exp(\theta - \dot{\theta}_m)^T s(X_m^k). \quad (5)$$

Note that those estimators depend on the simulated data only through the sufficient statistics $\mathfrak{s}_m^k := s(X_m^k)$. As a consequence, only the sufficient statistics $\mathfrak{S} := \{\mathfrak{s}_m^k\}_{m,k}$ need to be saved, as opposed to the actual collection of simulated graphs at each grid point. In the following we denote by $\mathfrak{U} := \{\mathfrak{S}, \mathfrak{G}\}$ the collection of the pre-computing data comprising of the grid \mathfrak{G} and the simulated sufficient statistics \mathfrak{S} .

2.4 Estimators of the ratio of normalising constants

We now detail several pre-computing version of the Metropolis-Hastings algorithm. The central idea is to replace the ratio of normalizing constants in the Metropolis-Hastings acceptance probability (1) by an estimator based on \mathfrak{U} . As a starting point this can be done by observing that for all $(\theta, \theta', \dot{\theta}) \in \Theta^3$,

$$\frac{Z(\theta)}{Z(\theta')} = \frac{Z(\theta)}{Z(\dot{\theta})} \frac{Z(\dot{\theta})}{Z(\theta')} = \frac{Z(\theta)}{Z(\dot{\theta})} \bigg/ \frac{Z(\theta')}{Z(\dot{\theta})}, \quad (6)$$

and in particular for any grid point $\dot{\theta} \in \mathfrak{G}$. We thus consider a general class of estimators of $Z(\theta)/Z(\theta')$ written as

$$\rho_n^X(\theta, \theta', \mathfrak{U}) := \frac{\Psi_n^X(\theta, \theta', \mathfrak{U})}{\Phi_n^X(\theta, \theta', \mathfrak{U})}, \quad (7)$$

where Ψ_n and Φ_n are unbiased estimators of the numerator and the denominator of the right hand side of (6), respectively, based on \mathfrak{U} . In (7), X simply denotes the different type of estimators considered. To simplify notations and in absence of ambiguity, the dependence of ρ_n , Ψ_n and Φ_n on θ , θ' , \mathfrak{U} and X is made implicit and we stress that given $(\theta, \theta', \mathfrak{U}, X)$, the estimators Ψ_n and Φ_n are deterministic.

We first note that ρ_n as defined in (7) is not an unbiased estimator of $Z(\theta)/Z(\theta')$. In fact, resorting to biased estimators of the normalizing constants ratio is the price to pay for using the pre-computed data. This represents a significant departure compared to the algorithms designed in the noisy MCMC literature (Alquier et al., 2016; Medina-Aguayo et al., 2016). Nevertheless, as we shall see in the next Section, this does not prevent us from controlling the distance between the distribution of the pre-computing Markov chain and π .

We propose a number of different estimators of Ψ_n and Φ_n . Those estimators share in common the idea that, given the current chain location θ and an attempted move θ' , a path of grid point(s) $\{\dot{\theta}_{\tau_1}, \dot{\theta}_{\tau_2}, \dots, \dot{\theta}_{\tau_C}\} \subset \mathfrak{G}$ connects θ to θ' .

The simplest path consists of the singleton $\{\dot{\theta}_\tau\}$, where $\dot{\theta}_\tau$ is any grid point. Since only one grid point is used, we refer to this estimator as the *One Pivot* estimator. Following (6), the estimators Ψ_n and Φ_n are defined as

$$\begin{cases} \Psi_n^{\text{OP}}(\theta, \theta', \mathfrak{U}) := 1/n \sum_{k=1}^n q_\theta(X_\tau^k)/q_{\dot{\theta}_\tau}(X_\tau^k), \\ \Phi_n^{\text{OP}}(\theta, \theta', \mathfrak{U}) := 1/n \sum_{k=1}^n q_{\theta'}(X_\tau^k)/q_{\dot{\theta}_\tau}(X_\tau^k). \end{cases} \quad (8)$$

However, for some $(\theta, \theta', \dot{\theta}_\tau) \in \Theta^2 \times \mathfrak{G}$, the variance of Ψ_n or Φ_n defined in Eq. (8) may be large. This is especially likely when $\|\theta - \dot{\theta}_\tau\| \gg 1$ or $\|\theta' - \dot{\theta}_\tau\| \gg 1$. The following Example illustrates this situation.

Example 1. Consider the Erdős-Renyi graph model, where all graphs $y \in \mathcal{Y}$ with the same number of edges $s(y)$ are equally likely. More precisely, the dyads are independent and connected with a probability $\varrho(\theta) := \text{logit}^{-1}(\theta)$ for any $\theta \in \mathbb{R}$. The likelihood function is given for any $\theta \in \mathbb{R}$ by $f(y|\theta) \propto \exp\{\theta s(y)\}$. For this model, the normalizing constant is tractable. In particular, $Z(\theta) = \{1 + \exp(\theta)\}^{\bar{p}}$ where $\bar{p} = \binom{p}{2}$ and p is the number of nodes in the graph.

For all $\theta \in \mathbb{R}$, consider estimating the ratio $Z(\theta')/Z(\theta)$ with $\theta' = \theta + h$ for some $h > 0$ using the estimator

$$\left. \frac{\widehat{Z(\theta+h)}}{Z(\theta)} \right|_n = \frac{1}{n} \sum_{k=1}^n \frac{q_{\theta+h}(X_k)}{q_\theta(X_k)} = \frac{1}{n} \sum_{k=1}^n \exp\{hs(X_k)\}, \quad X_k \sim_{iid} f(\cdot|\theta).$$

Then, when h increases, the variance v_n of this estimator diverges exponentially i.e.

$$nv_n(h) \sim \exp(2h\bar{p})\nu(\theta), \quad (9)$$

where \sim denotes here the asymptotic equivalence notation and $\nu(\theta) = \varrho(\theta)^{\bar{p}}(1 - \varrho(\theta)^{\bar{p}})$ is a constant. Remarkably, $\nu(\theta)$ can be interpreted as the variance of the Bernoulli trial with the full graph and its complementary event as outcomes.

Proof. By straightforward algebra, we have

$$v_n(h) = \frac{1}{n} \left\{ \frac{1 + \exp(2h + \theta)}{1 + \exp(\theta)} \right\}^{\bar{p}} \{1 - R(\theta, h)\},$$

where

$$R(\theta, h) = \frac{\{1 + \exp(\theta + h)\}^{2\bar{p}}}{\{1 + \exp(2h + \theta)\}^{\bar{p}} \{1 + \exp(\theta)\}^{\bar{p}}}.$$

Asymptotically in h , we have

$$R(\theta, h) \sim \frac{\exp(\bar{p}\theta)}{\{1 + \exp(\theta)\}^{\bar{p}}} = \varrho(\theta)^{\bar{p}}$$

and noting that

$$\left\{ \frac{1 + \exp(2h + \theta)}{1 + \exp(\theta)} \right\}^{\bar{p}} \sim \exp(2h\bar{p}) \frac{\exp\{\bar{p}\theta\}}{\{1 + \exp(\theta)\}^{\bar{p}}} = \exp(2h\bar{p})\varrho(\theta)^{\bar{p}}$$

concludes the proof. \square

This is a concern since as we shall see in the next Section, the noise introduced by the pre-computing step in the Markov chain is intimately related to the variance of the estimator of $Z(\theta)/Z(\theta')$. In particular, the distance between the pre-computing chain distribution and π can only be controlled when the variance of Ψ_n and Φ_n is bounded. Example 1 shows that this is not necessarily the case, for some Gibbs random fields at least. The following Proposition hints at the possibility to control the variance of Ψ_n and Φ_n when $\|\theta - \theta'\| \ll 1$.

Proposition 1. For any Gibbs random field model and all $(\theta, \theta') \in \Theta^2$, the variance of the normalizing constant estimator

$$\left. \frac{\widehat{Z(\theta)}}{Z(\theta')} \right|_n := \frac{1}{n} \sum_{k=1}^n \frac{q_\theta(X_k)}{q_{\theta'}(X_k)}, \quad X_k \sim_{iid} f(\cdot | \theta')$$

decreases when $\|\theta - \theta'\| \downarrow 0$ and more precisely

$$\text{var} \left. \frac{\widehat{Z(\theta)}}{Z(\theta')} \right|_n = \mathcal{O}(\|\theta - \theta'\|^2). \quad (10)$$

Proposition 1 motivates the consideration of estimators that may have smaller variability than the One Pivot estimator.

- (1) *Direct Path* estimator: the path between θ and θ' consists now of two grid points $\{\dot{\theta}_1, \dot{\theta}_2\}$ defined such that $\dot{\theta}_1 = \arg \min_{\dot{\theta} \in \mathfrak{G}} \|\dot{\theta} - \theta\|$ and $\dot{\theta}_2 = \arg \min_{\dot{\theta} \in \mathfrak{G}} \|\dot{\theta} - \theta'\|$. We therefore extend (6) and write

$$\frac{Z(\theta)}{Z(\theta')} = \frac{Z(\theta)}{Z(\dot{\theta}_1)} \frac{Z(\dot{\theta}_1)}{Z(\dot{\theta}_2)} \frac{Z(\dot{\theta}_2)}{Z(\theta')} = \frac{Z(\theta)}{Z(\dot{\theta}_1)} \frac{Z(\dot{\theta}_1)}{Z(\dot{\theta}_2)} \bigg/ \frac{Z(\theta')}{Z(\dot{\theta}_2)}.$$

This leads to two estimators Ψ_n and Φ_n defined as

$$\begin{cases} \Psi_n^{\text{DP}}(\theta, \theta', \mathfrak{U}) := 1/n \sum_{k=1}^n q_\theta(X_1^k)/q_{\dot{\theta}_1}(X_1^k) \times 1/n \sum_{k=1}^n q_{\dot{\theta}_1}(X_2^k)/q_{\dot{\theta}_2}(X_2^k), \\ \Phi_n^{\text{DP}}(\theta, \theta', \mathfrak{U}) := 1/n \sum_{k=1}^n q_{\theta'}(X_2^k)/q_{\dot{\theta}_2}(X_2^k). \end{cases} \quad (11)$$

- (2) *Full Path* estimator: the path between θ and θ' consists now of adjacent grid points $\mathfrak{p}(\theta, \theta') := \{\dot{\theta}_1, \dot{\theta}_2, \dots, \dot{\theta}_C\}$, where $C > 1$ is a number that depends on θ and θ' . Note that given (θ, θ') , there is not only one path such as \mathfrak{p} connecting θ to θ' . However, for any possible path, two adjacent points $\{\dot{\theta}_i, \dot{\theta}_{i+1}\} \subset \mathfrak{p}(\theta, \theta')$ always satisfy the following identity (in the basis given by the eigenvector of $\mathcal{H}(\theta^*)$):

$$\exists j \in \{1, \dots, d\}, \quad V^T (\dot{\theta}_i - \dot{\theta}_{i+1}) = \pm \varepsilon \delta_j,$$

where δ_j refers to the d -dimensional indicator vector of direction j i.e. $\{\delta_j\}_\ell = \mathbb{1}_{j=\ell}$. As before, we extend (6) to accommodate this situation and write

$$\frac{Z(\theta)}{Z(\theta')} = \frac{Z(\theta)}{Z(\dot{\theta}_1)} \frac{Z(\dot{\theta}_1)}{Z(\dot{\theta}_2)} \times \dots \times \frac{Z(\dot{\theta}_{C-1})}{Z(\dot{\theta}_C)} \frac{Z(\dot{\theta}_C)}{Z(\theta')} = \frac{Z(\theta)}{Z(\dot{\theta}_1)} \frac{Z(\dot{\theta}_1)}{Z(\dot{\theta}_2)} \times \dots \times \frac{Z(\dot{\theta}_{C-1})}{Z(\dot{\theta}_C)} \bigg/ \frac{Z(\theta')}{Z(\dot{\theta}_C)}.$$

This then lead to consider two estimators Ψ_n and Φ_n defined as

$$\begin{cases} \Psi_n^{\text{FP}}(\theta, \theta', \mathfrak{U}) := 1/n \sum_{k=1}^n q_\theta(X_1^k)/q_{\dot{\theta}_1}(X_1^k) \times 1/n \sum_{k=1}^n q_{\dot{\theta}_1}(X_2^k)/q_{\dot{\theta}_2}(X_2^k) \\ \quad \times \dots \times 1/n \sum_{k=1}^n q_{\dot{\theta}_{C-1}}(X_{C-1}^k)/q_{\dot{\theta}_C}(X_C^k), \\ \Phi_n^{\text{FP}}(\theta, \theta', \mathfrak{U}) := 1/n \sum_{k=1}^n q_{\theta'}(X_C^k)/q_{\dot{\theta}_C}(X_C^k). \end{cases} \quad (12)$$

Variants of the Direct Path and Full Path estimators exist. For the Direct Path, Ψ_n could be estimating $Z(\theta)/Z(\dot{\theta}_{\tau_1})$ and Φ_n the ratio $Z(\dot{\theta}_{\theta'})/Z(\theta_{\tau_1})$. For the Full Path, defining $\dot{\theta}_{\tau_m}$ as a middle point of $\mathbf{p}(\theta, \theta')$, Φ_n and Ψ_n could respectively be defined as estimators of $Z(\theta)/Z(\dot{\theta}_{\tau_m})$ and $Z(\theta')/Z(\dot{\theta}_{\tau_m})$ using the same number of grid points in both estimators. However, our experiments have shown that these alternative estimators have very similar behaviour with those defined in Eqs. (11) and (12). In particular, the variance of an estimator does not vary much when path points are removed from the numerator estimator and added to the denominator estimator, or conversely. As hinted by Proposition 1, the discriminant feature between those estimators is the distance between grid points constituting the path. In this respect, the variance of the Full Path estimator was always found to be lower than that of the Direct Path or One Pivot estimators. Even though establishing a rigorous comparison result between those estimators is a challenge on its own, a reader might be interested in the following result that somewhat formalizes our empirical observations.

Proposition 2. *Let $(\theta, \theta') \in \Theta$ and consider the Direct Path and Full Path estimators of $Z(\theta)/Z(\theta')$ defined at (11) and (12). Denoting by $\{\dot{\theta}_1, \dots, \dot{\theta}_C\}$ a full path connecting θ to θ' , we define for $i \in \{2, \dots, C\}$ R_n^i as the estimator of $Z(\dot{\theta}_{i-1})/Z(\dot{\theta}_i)$ and R_n^{2C} as the estimator of $Z(\dot{\theta}_1)/Z(\dot{\theta}_C)$ i.e.*

$$R_n^i = \frac{1}{n} \sum_{k=1}^n \frac{q_{\dot{\theta}_{i-1}}(X_i^k)}{q_{\dot{\theta}_i}(X_i^k)}, \quad R_n^{2C} = \frac{1}{n} \sum_{k=1}^n \frac{q_{\dot{\theta}_1}(X_C^k)}{q_{\dot{\theta}_C}(X_C^k)}, \quad X_i^k \sim_{iid} f(\cdot | \dot{\theta}_i). \quad (13)$$

Let v_n^{FP} and v_n^{DP} be the variance of the Full Path and Direct Path estimators using n pre-computed sufficient statistics are drawn at each grid point.

Assume Φ_n and Ψ_n are independent. Then, there exists a positive constant $\gamma < \infty$ such that

$$v_n^{\text{DP}} - v_n^{\text{FP}} = \gamma \left\{ \text{var}(R_n^{2C}) - \text{var}(R_n^2 \times \dots \times R_n^C) \right\}. \quad (14)$$

Moreover,

$$\text{var}(R_n^{2C}) = \frac{1}{n} \text{var} \exp \left\{ (\dot{\theta}_1 - \dot{\theta}_C)^T s(X_C) \right\} \quad (15)$$

and for large n and C and small ε we have

$$\text{var}(R_n^2 \times \dots \times R_n^C) = \frac{\varepsilon^4}{n} \sum_{i=2}^C \left\{ v_i \frac{Z(\dot{\theta}_i)Z(\dot{\theta}_1)}{Z(\dot{\theta}_{i-1})Z(\dot{\theta}_C)} \right\}^2 + o(\varepsilon^4/n), \quad (16)$$

where $\{v_1, v_2, \dots\}$ is a sequence of finite numbers such that $v_i \in \mathcal{O}(\varepsilon)$.

Proposition 2 shows that for a large enough number of pre-computed draws n , a long enough path and a dense grid i.e. $\varepsilon \ll 1$, the variance of the Full Path estimator is several order of magnitude less than that of the Direct Path estimator. In particular, unlike the Full Path estimator, the grid refinement does not help to reduce the variance of the Direct Path estimator. Proposition 2 coupled with the observation made at Example 1 helps to understand the variance reduction achieved with the Full Path estimator compared to the Direct Path estimator.

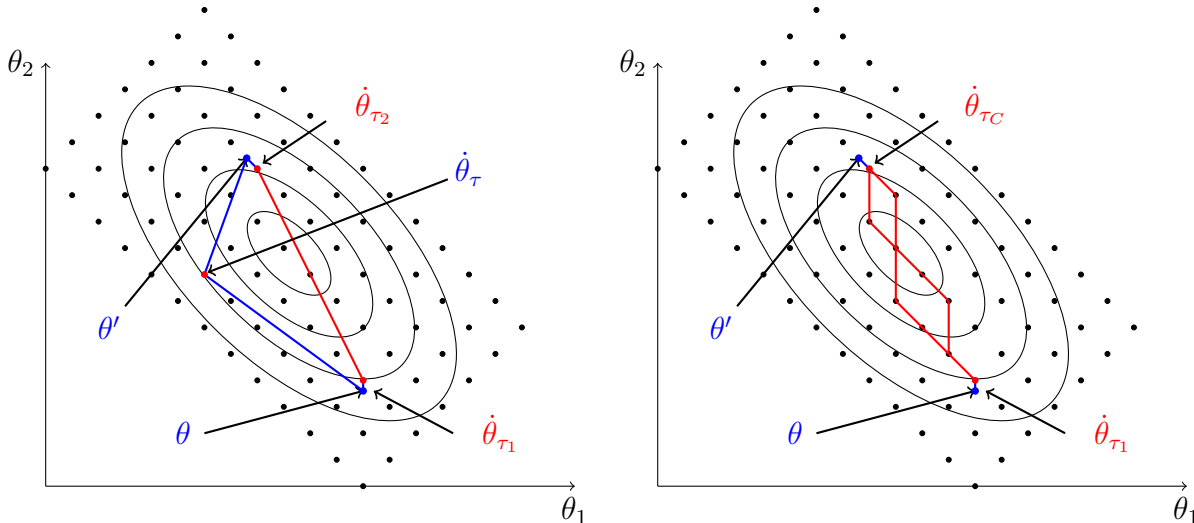


Figure 2: Example of paths between two parameters (θ, θ') in a two-dimensional space Θ . The solid black lines represent level lines of the target distribution and the black dots represent the grid vertices $\mathfrak{G} = \{\theta_1, \dots, \theta_M\}$. The thick lines show the paths $\mathbf{p}(\theta, \theta')$ used by the different estimators introduced in Eqs. (8), (11), (12): an example of a One Pivot path $\mathbf{p}(\theta, \theta') = \{\dot{\theta}\}$ (blue) and a Direct Path $\mathbf{p}(\theta, \theta') = \{\theta_1, \theta_2\}$ (red) are shown on the left panel. Two examples of Full Paths $\mathbf{p}(\theta, \theta') = \{\dot{\theta}_1, \dots, \dot{\theta}_C\}$ (red) are illustrated on the right panel: multiple possible full paths between θ and θ' could be used to average a number of Full Path estimators.

Note that when the parameter space is two-dimensional or higher, there is more than one choice of path connecting θ to θ' . The right panel of Figure 2 shows two different paths. In this situation, one could simply average the Full Path estimators obtained through each (or a number of) possible path. The different steps included in the Pre-computing Metropolis algorithm are summarized in Algorithm 2.

3 Asymptotic analysis of the pre-computing Metropolis-Hastings algorithms

In this section, we investigate the theoretical guarantees for the convergence of the Markov chain $\{\theta_k, k \in \mathbb{N}\}$ produced by the pre-computing Metropolis algorithm (Alg. 2) to the posterior distribution π . The Markov transition kernels considered in this section are conditional probability distributions on the measurable space (Θ, ϑ) where ϑ is the σ -algebra taken as the Borel set on Θ . We will use the following transition kernels:

Algorithm 2 Pre-computing Metropolis algorithm

(1)-Pre-computing

Require: Grid refinement parameter $\varepsilon > 0$ and number of draws $n \in \mathbb{N}$

- 1: Apply Algorithm 1 to define the grid $\mathfrak{G} = \{\dot{\theta}_1, \dots, \dot{\theta}_M\}$.
- 2: Initiate the collection of sufficient statistics to $\mathfrak{S} = \{\emptyset\}$.
- 3: **for** $m = 1$ to M **do**
- 4: **for** $k = 1$ to n **do**
- 5: Draw $X_m^k \sim_{iid} f(\cdot | \dot{\theta}_j)$
- 6: Calculate the vector of sufficient statistics $\mathfrak{s}_m^k = s(X_m^k)$
- 7: Append the pre-computed sufficient statistics set $\mathfrak{S} = \{\mathfrak{S} \cup \mathfrak{s}_m^k\}$
- 8: **end for**
- 9: **end for**

Return: The pre-computed data $\mathfrak{U} = \{\mathfrak{G}, \mathfrak{S}\}$

(2)-MCMC sampling

Require: Initial distribution μ and proposal kernel h , pre-computed data \mathfrak{U} and a type of estimator ρ_n^X , $X \in \{\text{OP}, \text{DP}, \text{FP}\}$

- 1: Initiate the Markov chain with $\theta_0 \sim \mu$
- 2: Identify the closest grid point from θ_0 , say $\dot{\theta}_i$, and calculate

$$Z_0 := \frac{1}{n} \sum_{k=1}^n \exp \left\{ (\theta_0 - \dot{\theta}_i)^T \mathfrak{s}_i^k \right\} .$$

- 3: **for** $i = 1, 2, \dots$ **do**
- 4: Draw $\theta' \sim h(\cdot | \theta_{i-1})$
- 5: Identify the closest grid point from θ' , say $\dot{\theta}_i$, and calculate

$$Z' := \frac{1}{n} \sum_{k=1}^n \exp \left\{ (\theta' - \dot{\theta}_i)^T \mathfrak{s}_i^k \right\} .$$

- 6: Using Z_{i-1} , Z' and \mathfrak{S} , calculate the normalizing ratio estimator ρ_n^X , depending on the type of estimator X using Eq. (8), (11) or (12).
- 7: Set $\theta_i = \theta'$ and $Z_i = Z'$ with probability

$$\bar{\alpha}(\theta_{i-1}, \theta', \mathfrak{U}) := 1 \wedge \bar{a}(\theta_{i-1}, \theta', \mathfrak{U}) ,$$

$$\bar{a}(\theta_{i-1}, \theta', \mathfrak{U}) = \frac{q_{\theta'}(y)p(\theta')h(\theta_i|\theta')}{q_{\theta_i}(y)p(\theta_i)h(\theta'|\theta_i)} \times \rho_n^X(\theta_{i-1}, \theta', \mathfrak{U}) \quad (17)$$

and else set $\theta_i = \theta_{i-1}$ and $Z_i = Z_{i-1}$.

- 8: **end for**

Return: The Markov chain $\{\theta_1, \theta_2, \dots\}$.

- Let P be the Metropolis-Hastings (MH) transition kernel defined as:

$$P(\theta, A) = \int_A h(d\theta' | \theta) \alpha(\theta, \theta') + \delta_\theta(A) r(\theta),$$

$$r(\theta) = 1 - \int_{\Theta} h(d\theta' | \theta) \alpha(\theta, \theta'), \quad (18)$$

where δ_θ is the dirac mass at θ and α the (intractable) MH acceptance probability defined at Eq. (1).

- Let $\bar{P}_{\mathfrak{U}}$ be the pre-computing Metropolis transition kernel, conditioned on the pre-computing data \mathfrak{U} and defined as:

$$\bar{P}_{\mathfrak{U}}(\theta, A) = \int_A h(d\theta' | \theta) \bar{\alpha}(\theta, \theta', \mathfrak{U}) + \delta_\theta(A) \bar{r}(\theta, \mathfrak{G}),$$

$$\bar{r}(\theta, \mathfrak{U}) = 1 - \int_{\Theta} h(d\theta' | \theta) \bar{\alpha}(\theta, \theta', \mathfrak{U}), \quad (19)$$

where $\bar{\alpha}$ is the pre-computing Metropolis acceptance probability defined at Eq. (17).

We recall that the MH Markov chain is π -invariant, a property which is lost by the pre-computing Metropolis algorithm. In what follows, we regard \bar{P} as a noisy version of the MH kernel P and $\bar{\alpha}$ as an approximation of the intractable quantity α . In terms of notations, we will use the following: for any $i \in \mathbb{N}$, P^i is the transition kernel P iterated i times and for any measure μ on (Θ, ϑ) , μP is the probability measure on (Θ, ϑ) defined as $\mu P(A) := \int \mu(d\theta) P(\theta, A)$.

Using the theoretical framework, developed in [Alquier et al. \(2016\)](#), we show that under certain assumptions, the distance between the distribution of the pre-computing Metropolis Markov chain and π can be made arbitrarily small, in function of the grid refinement and the number of auxiliary draws. The metric used on the space of probability distributions is the total variation distance, defined for two distributions (ν, μ) that admit a density function with respect to the Lebesgue measure as

$$\|\nu - \mu\| := (1/2) \int_{\Theta} |\nu(\theta) - \mu(\theta)| d\theta.$$

3.1 Noisy Metropolis-Hastings

We first recall the main result from [Alquier et al. \(2016\)](#) that will be used to analyse the pre-computing Metropolis algorithm.

Proposition 3 (Corollary 2.3 in [Alquier et al. \(2016\)](#)).

Let us assume that,

- **(H1)** A MH Markov chain with transition kernel P (Eq. 18), proposal kernel h and acceptance probability α (Eq. 1) is uniformly ergodic i.e. there are constants $B > 0$ and $\rho < 1$ such that

$$\forall i \in \mathbb{N}, \quad \sup_{\theta_0 \in \Theta} \|\delta_{\theta_0} P^i - \pi\| \leq B \rho^i.$$

- **(H2)** There exists an approximation of the Metropolis acceptance ratio a , $\hat{a}(\theta, \theta', X)$ that satisfies for all $(\theta, \theta') \in \Theta^2$

$$\mathbb{E} |\hat{a}(\theta, \theta', X) - a(\theta, \theta')| \leq \epsilon(\theta, \theta'),$$

where the expectation is with respect to the noise random variable X .

Then, denoting by \hat{P} the noisy Metropolis-Hastings kernel (Eq. 19), we have for any starting point $\theta_0 \in \Theta$ and any integer $i \in \mathbb{N}$:

$$\|\delta_{\theta_0} P^i - \delta_{\theta_0} \hat{P}^i\| \leq \left(\lambda - \frac{B\rho^\lambda}{1-\rho} \right) \sup_{\theta \in \Theta} \int d\theta' h(\theta'|\theta) \epsilon(\theta, \theta'), \quad (20)$$

where $\lambda = \left(\frac{\log(1/B)}{\log(\rho)} \right)$.

An immediate consequence of Proposition 3 is that if ϵ is uniformly bounded, i.e. there exists some $\bar{\epsilon} > 0$ such that for all $(\theta, \theta') \in \Theta^2$, $\epsilon(\theta, \theta') \leq \bar{\epsilon} < \infty$, then

$$\forall i \in \mathbb{N}, \quad \|\delta_{\theta_0} P^i - \delta_{\theta_0} \hat{P}^i\| \leq \bar{\epsilon} \left(\lambda - \frac{B\rho^\lambda}{1-\rho} \right). \quad (21)$$

Moreover, defining $\hat{\pi}_i$ as the distribution of the i -th state of the noisy chain yields

$$\lim_{i \rightarrow \infty} \|\pi - \hat{\pi}_i\| \leq \bar{\epsilon} \left(\lambda - \frac{B\rho^\lambda}{1-\rho} \right). \quad (22)$$

3.2 Convergence of the pre-computing Metropolis algorithm

In preparation to apply Proposition 3, we make the following assumptions:

- **(A1)** there is a constant c_p such that for all $\theta \in \Theta$, $1/c_p \leq p(\theta) \leq c_p$.
- **(A2)** there is a constant c_h such that for all $(\theta, \theta') \in \Theta^2$, $1/c_h \leq h(\theta'|\theta) \leq c_h$.

Assumptions **(A1)** and **(A2)** are typically satisfied when Θ is a bounded set and p and $h(\cdot|\theta)$ are dominated by the Lebesgue measure. Under similar assumptions, Proposition 3 was applied to the noisy Metropolis algorithm (Alquier et al., 2016) that uses the unbiased estimator ϱ_n (Eq. 2). More precisely, it was shown that the distance between π and $\hat{\pi}_i$ satisfies $\|\pi - \hat{\pi}_i\| \leq \kappa/\sqrt{n}$, where $\kappa > 0$ is a positive constant, asymptotically in i .

Establishing an equivalent result for the pre-computing Metropolis algorithms is not straightforward. The main difficulty is that the acceptance ratio $\tilde{a}(\theta, \theta', \mathfrak{U})$ (Eq. 17) is a biased estimator of the MH acceptance ratio $a(\theta, \theta')$ (Eq. 1). The following Proposition only applies to the pre-computing Metropolis algorithm involving the approximation of the normalizing constant ratio using the full path estimator. Weaker results can be obtained using similar arguments for the One Pivot and Direct Path estimators.

Proposition 4. Assume that **(H1)**, **(A1)** and **(A2)** hold and for any $(\theta, \theta') \in \Theta^2$ define by C the shortest path $\mathbf{p}(\theta, \theta')$ length. Then, there exists a sequence of functions $u_n : \mathbb{N} \rightarrow \mathbb{R}^+$ and a function $v : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ satisfying

$$u_n(C) = \frac{\sqrt{C}}{\sqrt{n}} + o(n^{-1/2}), \quad v(\varepsilon) = 2\sqrt{d\psi_1\varepsilon} + o(\varepsilon^{1/2}), \quad (23)$$

such that the pre-computing Metropolis acceptance ratio $\bar{a}(\theta, \theta', \mathfrak{S})$ (Eq. 17) satisfies

$$\mathbb{E} |\bar{a}(\theta, \theta', \mathfrak{S}) - a(\theta, \theta')| \leq 2c_p^2 c_h^2 K_1 K_2^{C+2d-1}(\varepsilon) \{u_n(C) + v(\varepsilon)\}. \quad (24)$$

In Eq. (23), $\psi_1 < \infty$ is a constant, n is the number of pre-computed GRF realizations for each grid point and ε is the distance between grid points. In Eq. (24), K_1 and $K_2(\varepsilon)$ are finite constants such that $K_2(\varepsilon) \rightarrow 1$ when $\varepsilon \downarrow 0$.

Corollary 1. Define by $\bar{\pi}_i$ the distribution of the i -th iteration of the pre-computing Metropolis algorithm implemented with the Full Path estimator. Under Assumptions **(H1)**, **(A1)** and **(A2)**, we have

$$\lim_{i \rightarrow \infty} \|\pi - \bar{\pi}_i\| \leq \bar{\kappa} K_2^{2d-1}(\varepsilon) \sum_{c=1}^M K_2^c(\varepsilon) \{u_n(c) + v(\varepsilon)\} p_c, \quad (25)$$

where $p_c = \mathbb{P}\{C = c\}$ is the probability distribution of the path length and

$$\bar{\kappa} = \left(\lambda - \frac{B\rho^\lambda}{1-\rho} \right) 2c_p^2 c_h^2 K_1.$$

In Eq. (25), u_n and v are defined in Eq. (23).

Corollary 1 states that the asymptotic distance between the pre-computing Markov chain distribution and π admits an upper bound that has two main components:

- $u_n(c) \sim \sqrt{c/n}$ which is related to the variance of each estimator of a normalizing constant ratio estimator,
- $v(\varepsilon) \sim 2\sqrt{d\psi_1\varepsilon}$ that arises from using a fixed step size grid.

This provides useful guidance as to how to tune the pre-computing parameters n and ε . In particular, n should increase with the proposal kernel h variance and ε should decrease with the dimension of Θ , that is d . When $\varepsilon \rightarrow 0$ the upper bound of $\|\pi - \bar{\pi}_i\|$ is in $1/\sqrt{n}$ which is in line with the noisy Metropolis rate of (Alquier et al., 2016). Interestingly, when $\varepsilon \ll 1$, we believe that our bound is tighter thanks to the lower variability of the Full Path estimator compared to the unbiased estimator ϱ_n (Eq. 2) used in the noisy Metropolis algorithm. Indeed, their bound is in $o(K_1^4/\sqrt{n})$ which, given the crude definition of K_1 , is much looser compared to our $o(K_1/\sqrt{n})$ bound.

The following Proposition shows that when the number of data n simulated at the pre-computing step tends to infinity then $\mathbb{E} |\bar{a}(\theta, \theta', \mathfrak{S}) - a(\theta, \theta')|$ vanishes. This result is somewhat reassuring as it suggests that the pre-computing algorithm will converge to the true distribution, asymptotically in n , regardless of the grid specification. However, it is not possible to embed this result in the framework developed in Alquier et al. (2016) as the convergence comes without a rate.

Proposition 5. *For any pre-computing Metropolis acceptance ratio that use an estimator of the normalizing constants ratio of the form specified at Eq. (7):*

$$\mathbb{E} |\bar{a}(\theta, \theta', \mathfrak{U}) - a(\theta, \theta')| \leq \int |\psi\phi - \alpha| f_n(d\psi | \phi) (g_n(\phi) - g(\phi)) d\phi + \frac{1}{\mathbb{E}(\Phi_1)\sqrt{n}} \left\{ \sqrt{\text{var}(\Psi_1)} + \sqrt{\text{var}(\Phi_1)} \frac{\mathbb{E}(\Psi_n | \zeta)}{\mathbb{E}(\Phi_1)} \right\},$$

where $\zeta \sim \mathcal{N}(0, 1)$, f_n , g_n and g are probability density functions such that g_n converges weakly to g .

3.3 Toy Example

We consider in this section the toy example used to illustrate the Exchange algorithm in (Murray et al., 2006, Section 5). More precisely, the experiment consists of sampling from the posterior distribution of the precision parameter θ arising from the following model:

$$f(\cdot | \theta) = \mathcal{N}(0, 1/\theta), \quad p = \text{Gamma}(1, 1),$$

using one observation $y = 2$ and pretending that the normalizing constant of the likelihood, namely $Z(\theta) = \int \exp(-\theta y^2/2) dy = \sqrt{2\pi/\theta}$ is intractable. The grid is set as $\mathfrak{G} = \{\hat{\theta}_m = m\varepsilon, 0 < m \leq \lfloor 10/\varepsilon \rfloor\}$. Our objective is to quantify the bias in distribution generated by the pre-computing algorithms.

We consider the situation where the interval between the grid points is $\varepsilon = 0.1$ and $n = 10$ data are simulated per grid points. Table 1 reports the bias and the variance of the three estimators, i.e. the One Pivot, Direct Path and Full Path, of the ratio $Z(\theta)/Z(\theta')$ for three couples (θ, θ') . This shows that the Full Path estimators enjoys a greater stability than the two other estimators, even when n is relatively small. This is completely in line with the results developed in Propositions 1 and 2.

Figure 3 illustrates the convergence of the three pre-computing Markov chains by reporting the estimated total variation distance between $\bar{\pi}_i$ and π . We also report the convergence of the exchange Markov chain: this serves as a ground truth since π is the stationary distribution of this algorithm. For each algorithm, the total variation distance was estimated by simulating 100,000 iid copies of the Markov chain of interest and calculating at each iteration the occupation measure. This measure is then compared to π which is, in this example, fully tractable. In view of Table 1, the chains implemented with the One Path and Direct Path estimators converge, as expected, further away from π than the Full Path chain.

Interestingly, it can be noted that the Full Path pre-computing chain converges faster than the exchange algorithm. This is an illustration of the observation stated in the introduction regarding the theoretical efficiency of the exchange, compared to that of the plain MH algorithm. Indeed, the pre-computing algorithms aim at approximating MH, and not the exchange algorithm, and should as such inherits MH's fast rate of convergence, provided that the variance of the estimator is controlled.

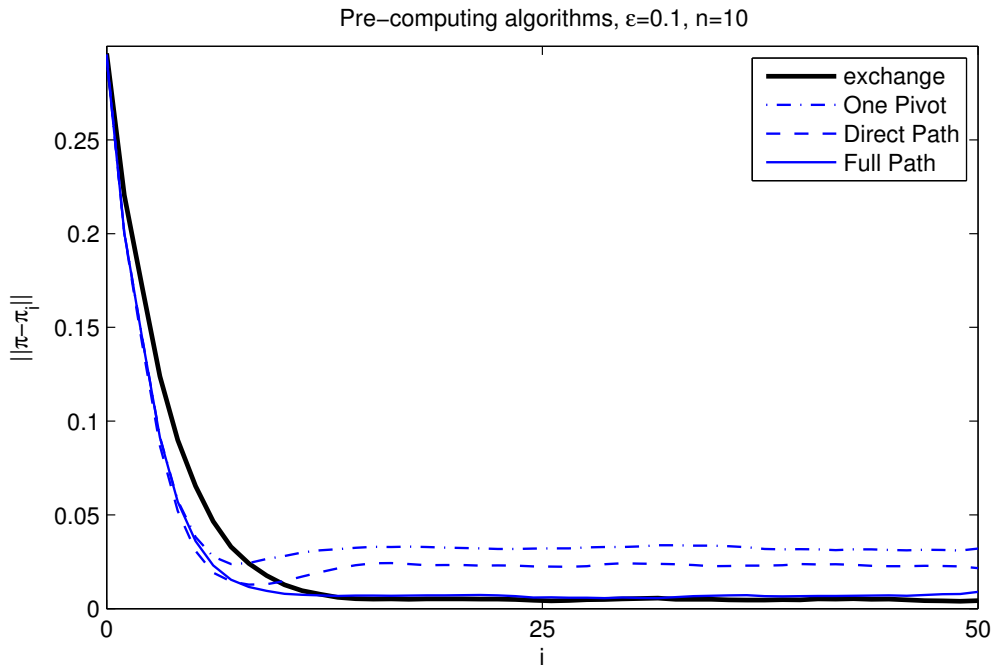


Figure 3: Convergence of the pre-computing Metropolis algorithms distribution. Results were obtained from 100,000 iid copies of the Markov chains initiated with $\mu = p$. All the chains were implemented with the same proposal kernel, namely $\theta' = \theta \exp \sigma \zeta$, $\zeta \sim \mathcal{N}(0, 1)$ and run for 50 iterations. The pre-computing parameters were set to $\varepsilon = 0.1$ and $n = 10$. Comparing the convergence of the pre-computing chains to that of the exchange (which theoretically converges to π), we see that the Full Path estimator has a negligible bias. This is not the case for the One Pivot and Direct Path implementations.

Table 1: Bias and variance of the different estimators of the ratio $Z(\theta)/Z(\theta')$ for various couples (θ, θ') in the setup of Figure 3. The bias and variance were estimated by simulating 10,000 independent realisations of each estimators for each couple (θ, θ') .

	$(\theta, \theta') = (1.01, 2.06)$		$(\theta, \theta') = (3.02, 0.55)$		$(\theta, \theta') = (0.12, 0.94)$	
	bias	var.	bias	var.	bias	var.
FP	.0007	.005	.0004	.001	.01	1.42
DP	.003	.208	.003	.013	.27	99.02
OP	.004	.199	.003	.014	.32	129.81

4 Results

This section illustrates our algorithm. A simulation study using the Ising model demonstrates the application to a ‘large’ dataset for a single parameter model. More challenging examples are provided with application to a multi-parameter autologistic and Exponential Random Graph Model (ERGM). In the single parameter example we use the estimates of the normalizing constant from Equations (11) and (12), denoted Full Path and Direct Path respectively. For the single parameter example we compare the pre-computing Metropolis algorithm with the standard exchange algorithm (Murray et al., 2006) and also with a version of the methods in Moores et al. (2015). Rather than the Sequential Monte Carlo ABC used in Moores et al. (2015), we implemented their pre-computation approach with a MCMC-ABC algorithm (Majoram et al., 2003). This allowed a fair comparison of expected total variation distance and effective sample size.

MCMC-ABC

Moores et al. (2015) used a pre-computing step with Sequential Monte Carlo ABC (see e.g. Del Moral et al. (2006)) to explore the posterior distribution. However, Sequential Monte Carlo has a stopping criterion which results in a finite sample size of values from the posterior distribution. To establish a fair comparison between algorithms whose sample size consistently increases over time, we implemented a modified version of the method proposed in Moores et al. (2015) using the MCMC-ABC algorithm. The modification made to the MCMC-ABC algorithm amounts to replace a draw $y' \sim f(\cdot|\theta)$ by a distribution that uses the pre-computed data. More precisely, sufficient statistics of a graph at a particular value θ are sampled from a normal distribution

$$s \sim \mathcal{N}(\mu(\theta, \mathfrak{U}), \sigma^2(\theta, \mathfrak{U})) .$$

The parameters $\mu(\cdot, \mathfrak{U})$ and $\sigma^2(\cdot, \mathfrak{U})$ are interpolated using the mean and variance of the pre-computed sufficient statistics obtained at the grid points. This pre-computing version of ABC-MCMC is described in Algorithm 3.

Algorithm 3 Pre-computing MCMC-ABC sampler

Require: Initial distribution ν , a proposal kernel h and ABC tolerance parameter $\epsilon > 0$

- 1: Apply the pre-computing step detailed in [Moore et al. \(2015\)](#)
 \rightsquigarrow pre-computed data \mathfrak{U}' .
- 2: Draw $\theta_0 \sim \nu$
- 3: **for** $i = 1, 2, \dots$ **do**
- 4: Draw $\theta' \sim h(\cdot | \theta_{i-1})$
- 5: Calculate the mean μ' and variance σ'^2 using the interpolation method in [Moore et al. \(2015\)](#) and the pre-computed data \mathfrak{U}' for the parameter θ'
- 6: Simulate the sufficient statistic $s' \sim \mathcal{N}(\mu, \sigma'^2)$
- 7: Set $\theta_i = \theta'$ with probability

$$\alpha_{\text{ABC}}(\theta, \theta', \mathfrak{U}) := 1 \wedge \frac{\pi(\theta')h(\theta_{i-1}|\theta')}{\pi(\theta_{i-1})h(\theta|\theta_{i-1})} \times \mathbb{1}_{|s'-s(y)|<\epsilon}(s')$$

and else set $\theta_{i+1} = \theta_i$.

8: **end for**

Return: The Markov chain $\{\theta_1, \theta_2, \dots\}$.

In the multi-parameter example we only compare results of the pre-computing Metropolis with the standard exchange algorithm since the method of [Moore et al. \(2015\)](#) cannot be implemented in higher dimensions.

4.1 Ising simulation study

The Ising model is defined on a rectangular lattice or grid. It is used to model the spatial distribution of binary variables, taking values -1 and 1. The joint density of the Ising model can be written as

$$f(y|\theta) = \frac{1}{Z(\theta)} \exp \left(\theta \sum_{j=1}^M \sum_{i \sim j} y_i y_j \right),$$

where $i \sim j$ denotes that i and j are neighbours and $Z(\theta) = \sum_y \exp \left(\sum_{j=1}^M \sum_{i \sim j} y_i y_j \theta \right)$. The normalizing constant is rarely available analytically since this relies on taking summation over all different possible realisations of the lattice. For a lattice with M nodes this equates to $2^{\frac{M(M-1)}{2}}$ different possible lattice formations.

In this study, 24 lattices of size 80×80 were simulated. The true distribution of the graphs were estimated using a long run (24 hours) of the exchange algorithm. Each of the algorithms was run for just over 60 minutes. The pre-computation step of choosing the parameter grid and estimating the ratios for every pair of grid values took approximately 13 minutes. For each of the algorithms we estimated the total variation distance using numerical integration across the kernel density estimates. The

values obtained give an indication of which of the chain best matches the long run of the exchange algorithm. The graph in Figure 4 is the average of the total variation for each algorithm over all 24 lattices.

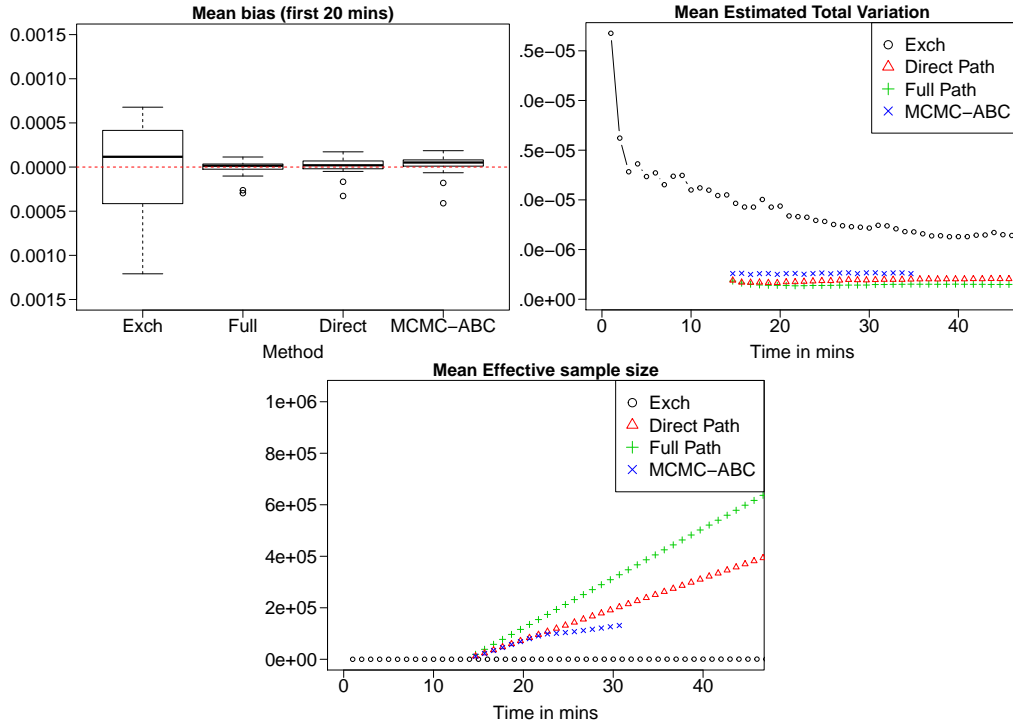


Figure 4: Results for the Ising study. The boxplots on the top left show the mean bias of the 24 graphs after the first 20 minutes of computation time: the pre-computing Metropolis algorithm performs the best. The plot on the top right shows the mean estimated total variation of the 24 graphs over time, the pre-computing Metropolis and the MCMC-ABC algorithm both outperform the standard exchange algorithm. The bottom plot shows the effective sample size over time, the pre-computing Metropolis algorithm, implemented with the Full Path estimator performs the best followed by the Direct Path estimator.

The results shown in Figure 4 illustrate how the pre-computing Metropolis algorithms (full path and direct path) outperforms the exchange algorithm over time. As more iterations can be calculated per second, the pre-computing Metropolis algorithm converges quicker. In this simulation, for fairness of comparison, the pre-computing data \mathcal{U} were re-simulated for each individual graph. Indeed, since all the graphs are on the same state space, only one single pre-computation step for a large set of parameter values over the full state space could have been sufficient. When analysis is required for many graphs which lie on the same state space, we only need to carry out the pre-computation step once. We stress that in practice, this situation is common and the speed-up factor obtained by using the pre-computing algorithm would be even more striking.



Figure 5: Presence (red) and absence (black) of red deer in the Grampian region of Scotland.

4.2 Autologistic Study

For the second illustration, we extend the Ising model to the autologistic model. The autologistic model is a GRF model for spatial binary data. The likelihood of the autologistic model is given by,

$$\begin{aligned} f(y|\theta) &\propto \exp(\theta^T s(y)) \\ &= \exp(\theta_1 s_1(y) + \theta_2 s_2(y)), \end{aligned}$$

where $s_1(y) = \sum_{i=1}^N y_i$ and $s_2(y) = \sum_{i \sim j} y_i y_j$ with $i \sim j$ denoting node i and node j are neighbours. θ_1 controls the relative abundance of -1 and $+1$ values while θ_2 controls the level of spatial aggregation. We implement the autologistic model using red deer census data, presence or absence of deer by 1km square in the Grampian region of Scotland (Augustin et al., 1996). Figure 5 shows the observed data, a red square indicates the presence of deer, while a black square indicates the absence of deer.

A long run (4 hours) of the exchange algorithm was used to set a 'ground truth'. The pre-computing grid points (top left of Figure 6) were chosen using the method described in Algorithm 1. A total of 124 parameter values were chosen as the values to pre-sampled from. It took just over 45 seconds to choose the grid and calculate the ratios for all pairs of parameter values. The pre-computing Metropolis algorithms all outperform the exchange algorithm as they converge much quicker, as shown at the top right panel of Figure 6. In this example, the two different choices of paths yield very similar results in terms of estimate total variation distance. The pre-computing Metropolis algorithms result in a more accurate mean and variance parameter estimates when compared to the exchange algorithm run for the same amount of time ; see Table 2. When the chains are run for longer, it takes the exchange algorithm 34 minutes to reach the same estimated total variation distance that the pre-computing Metropolis algorithms takes to reach in 200 seconds. This illustrates the substantial time saving resulting from the pre-computing Metropolis algorithms.

Table 2: Posterior means and variances for the deer data. The table shows that the mean and variance estimates of the noisy exchange are closer to the 'ground truth' long exchange run.

	θ_1		θ_2	
	Mean	Variance	Mean	Variance
Exchange (long)	-0.1435429	0.00028611	0.1516334	0.00016096
Exchange	-0.1424322	0.00026794	0.1530567	0.00014771
Full Path	-0.1434566	0.00026373	0.1519860	0.00015384
Direct Path	-0.1436186	0.00028256	0.1515273	0.00016495

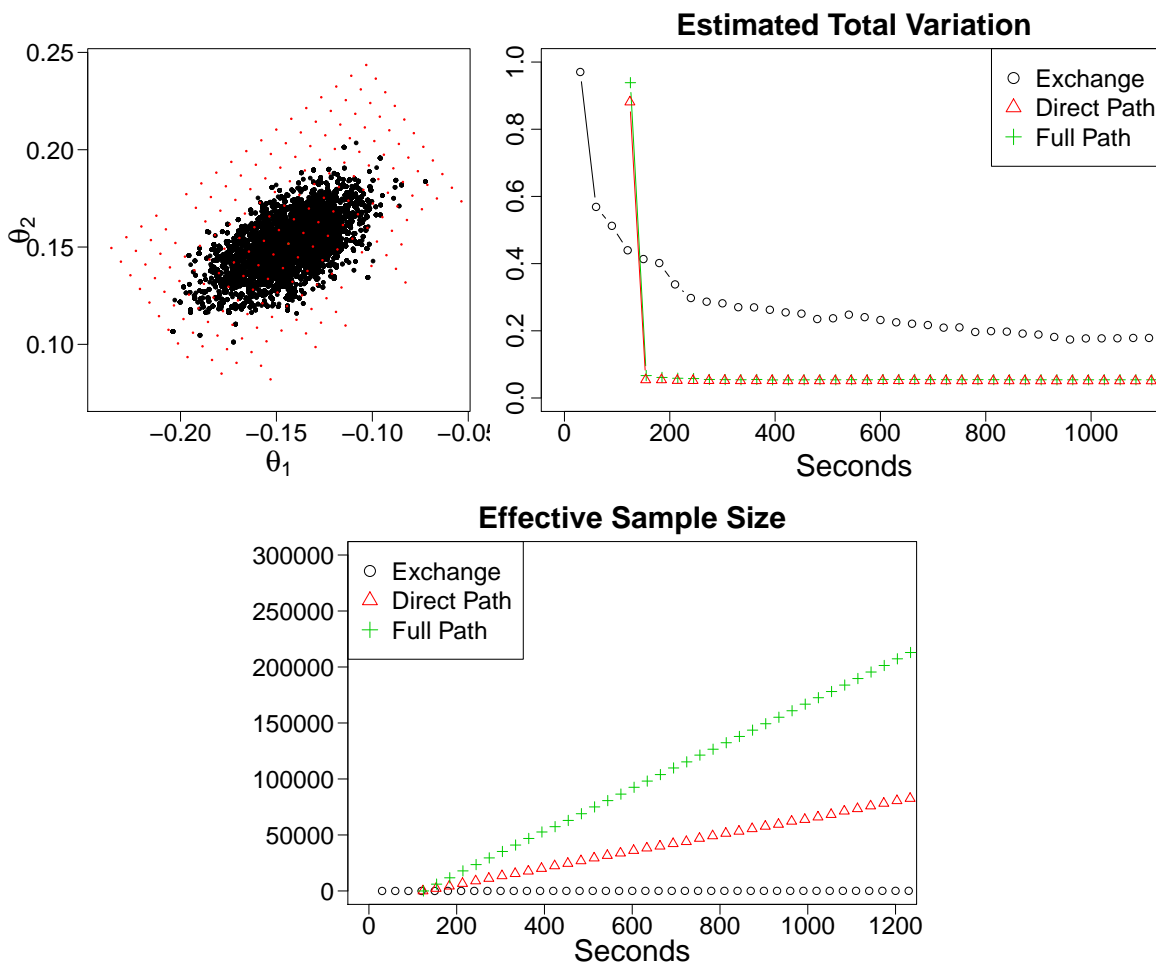


Figure 6: Grid for pre-computing (top left) and estimated total variation over time (right). The plot on the top right shows that when the estimated total variation distance between the algorithms and the long exchange is compared, the pre-computing Metropolis algorithms outperform the exchange algorithm. The two versions of the pre-computing Metropolis algorithm also outperform the exchange in terms of effective sample size.

4.3 ERGM study

We now show how our algorithms may be applied to the Exponential Random Graph model (ERGM) (Robins et al., 2007), a model which is widely used in social network analysis. An ERGM is defined on a random adjacency matrix \mathcal{Y} of a graph on p nodes (or actors) and a set of edges (dyadic relationships) $\{Y_{ij} : i = 1, \dots, M; j = 1, \dots, M\}$ where $Y_{ij} = 1$ if the pair (i, j) is connected by an edge, and $Y_{ij} = 0$ otherwise. An edge connecting a node to itself is not permitted so that $Y_{ii} = 0$. The dyadic variables may be undirected, whereby $Y_{ij} = Y_{ji}$ for each pair (i, j) , or directed, whereby a directed edge from node i to node j is not necessarily reciprocated.

The likelihood of an observed network $y \in \mathcal{Y}$ is modelled in terms of a collection of sufficient statistics $\{s_1(y), \dots, s_d(y)\}$, each with corresponding parameter vector $\theta = \{\theta_1, \dots, \theta_d\}$,

$$f(y|\theta) = \frac{q_\theta(y)}{Z(\theta)} = \frac{\exp\{\sum_{l=1}^m \theta_l s_l(y)\}}{Z(\theta)}.$$

Typical statistics include the observed number of edges and the observed number of two-stars, which is the number of configurations of pairs of edges which share a common node. Those statistics are usually defined as

$$s_1(y) := \sum_{i < j} y_{ij}, \quad s_2(y) := \sum_{i < j < k} y_{ik} y_{jk}.$$

It is also possible to consider statistics which count the number of triangle configurations, that is, the number of configurations in which nodes $\{i, j, k\}$ are all connected to each other.

4.3.1 Karate dataset

We consider Zachary's karate club (Zachary, 1977) which represents the undirected social network graph of friendships between 34 members of a karate club at a US university in the 1970s.

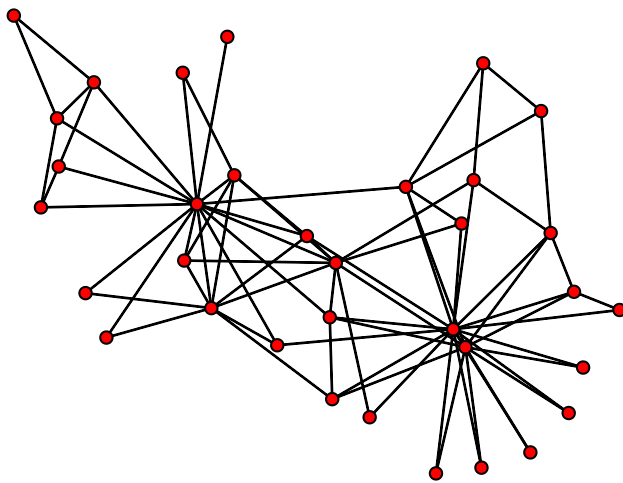


Figure 7: Karate club data.

We consider the following two-dimensional model,

$$f(y|\theta) = \frac{1}{Z(\theta)} \exp \{ \theta_1 s_1(y) + \theta_2 s_2(y) \},$$

where $s_1(y)$ is the number of edges in the graph and $s_2(y)$ is the number of triangles in the graph.

A long run of the exchange algorithm was again used to set a ‘ground truth’. The pre-computing step took roughly 30 seconds to set the $M = 191$ parameter values constituting the grid and to calculate the estimated normalizing ratio between each pair of parameter values using $n = 1,000$ simulated graphs. The mean and variance of the parameter estimates for the noisy exchange algorithms using the two different paths and a short run of the exchange algorithm are compared in Table 3. Figure 8 shows the choice of parameter for pre-processing (left) and the estimated total variation distance over time (right). Some grid points lie beyond the posterior distribution high density region, indicating that some graphs sampled from the tail regions could have been avoided. In practice however, it was found that allowing the grid to span beyond the posterior distribution high density regions gave much better results. The two versions of the pre-computing Metropolis algorithm outperform the exchange algorithm in the estimated total variation distance over time.

Table 3: Posterior means and variances for the karate data.

	Edge		Triangle	
	Mean	Var	Mean	Var
Exchange (long)	-2.0471	0.0962	0.3807	0.0306
Exchange	-2.1758	0.0739	0.4670	0.0254
Full Path	-2.3328	0.0991	0.4922	0.0210
Direct Path	-2.1645	0.1095	0.4518	0.0454

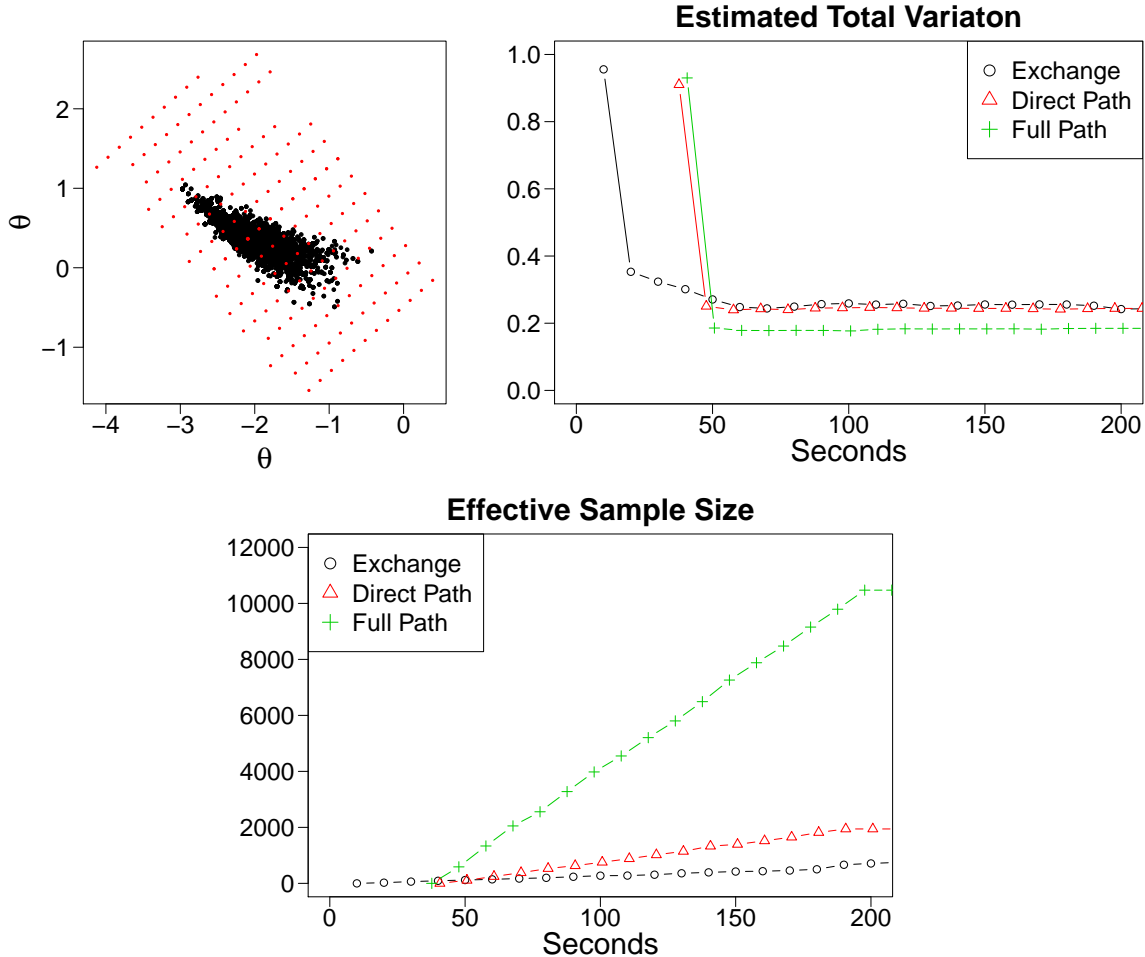


Figure 8: Grid for pre-computing (left) and estimated total variation distance over time (right). The pre-computing Metropolis algorithms outperform the exchange in terms of estimated total variation distance. The effective sample size of the pre-computing algorithms is much higher than the exchange.

5 Conclusion

This paper considers including an offline, easily parallelizable, pre-computing step as a way to overcome the computational bottleneck of certain variants of the Metropolis algorithm. In particular, we show how such a strategy can be efficient when inferring a doubly-intractable distribution, a situation that typically arises in the study of Gibbs random fields. The pre-computing Metropolis algorithms that we develop in this paper somewhat borrow from previous pre-computing algorithms (see *e.g.* Moores et al. (2015)) but scale better to higher dimensional settings. We however note that our method would be impractical for very high dimensions. Yet, the limit on the number of dimensions is similar to the limit on the INLA method (Rue et al., 2009), which has

seen widespread use in many areas.

The pre-computing Metropolis algorithms are noisy MCMC algorithms in the sense that the posterior of interest is not the invariant distribution of the Markov chain. However, we establish, under certain conditions, some theoretical results showing that the pre-computing Metropolis distribution converges into a ball centered on the true posterior distribution. Interestingly, the ball radius can be made arbitrarily small according to the pre-computing parameters, namely the space between grid points and the number of auxiliary data simulated per grid points. Our main contribution to the theoretical analysis of approximate Markov chains is twofold:

- In contrast to estimators of the Metropolis acceptance ratio that have been used in the approximate MCMC literature (see *e.g.* [Alquier et al. \(2016\)](#), [Medina-Aguayo et al. \(2016\)](#) and [Bardenet et al. \(2014\)](#)), the different estimators considered in this paper (i.e. the One Pivot, the Direct Path and the Full Path) are all biased. We stress that, when computational time is not an issue, there is no particular gain in efficiency using biased estimators but biasedness is an inevitable by-product when estimators make use of pre-computed data exclusively.
- A recurrent outcome from the research on approximate MCMC methods highlights the importance of using estimators of the Metropolis acceptance ratio with *small* variance. We refer for instance to the aforementioned works and [Bardenet et al. \(2015\)](#), [Quiroz et al. \(2017\)](#) and [Stoehr et al. \(2017\)](#). In the context of estimating a ratio of normalizing constants, we argue that the pre-computing step allows to specify low variance estimators, yet biased, at low computational cost by considering intermediate grid points, an idea that has been exploited by the Full Path estimator.

The empirical results show that in time normalized experiments, the pre-computing Metropolis algorithms provide accurate and efficient inference that outperform existing techniques such as the exchange algorithm ([Murray et al., 2006](#)).

Focus for future research will examine alternative methods that would allow inference of higher dimensional models. As it stands, the curse of dimensionality implies an exponential growth of the number of grid points, which makes our pre-computing step far too computationally intensive to be implemented in this setting. A way to overcome this challenge would be to design the grid adaptively, i.e. as the Markov chain is being simulated, in order to avoid unnecessary simulations at grid points whose vicinity is never visited by the Markov chain. Even though such a strategy is straightforward to implement, the theoretical analysis of the resulting algorithm is more involved. Indeed, it calls for results on ergodicity of approximate adaptive Markov chain, a research topic which is for now essentially unexplored.

Acknowledgements

The Insight Centre for Data Analytics is supported by Science Foundation Ireland under Grant Number SFI/12/RC/2289. Nial Friel’s research was also supported by a Science Foundation Ireland grant: 12/IP/1424.

References

- Alquier, P., N. Friel, R. Everitt, and A. Boland (2016). Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. Statistics and Computing 26(1), 29–47.
- Andrieu, C. and G. Roberts (2009). The pseudo-marginal approach for efficient Monte Carlo computations. The Annals of Statistics 37(2), 697–725.
- Andrieu, C. and J. Thoms (2008). A tutorial on Adaptive MCMC. Statistics and computing 18(4), 343–373.
- Augustin, N. H., M. A. Muggleston, and S. T. Buckland (1996). An autologistic model for the spatial distribution of wildlife. Journal of Applied Ecology 33(2), pp. 339–347.
- Bardenet, R., A. Doucet, and C. Holmes (2014). Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In Proceedings of the 31st International Conference on Machine Learning (ICML-14), pp. 405–413.
- Bardenet, R., A. Doucet, and C. Holmes (2015). On Markov chain Monte Carlo methods for tall data. arXiv preprint arXiv:1505.02827.
- Besag, J. E. (1974). Spatial interaction and the statistical analysis of lattice systems. Journal of the Royal Statistical Society 36, 192–236.
- Caimo, A. and N. Friel (2011). Bayesian inference for exponential random graph models. Social Networks 33(1), 41–55.
- Cucala, L., J.-M. Marin, C. P. Robert, and D. Titterton (2009). A Bayesian re-assessment of nearest-neighbour classification. Journal of the American Statistical Association 104, 263–273.
- Del Moral, P., A. Doucet, and A. Jasra (2006). Sequential Monte Carlo samplers. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68(3), 411–436.
- Everitt, R. (2012). Bayesian parameter estimation for latent Markov random fields and social networks. Journal of Computational and Graphical Statistics 21(4), 940–960.
- Everitt, R., D. Prangle, P. Maybank, and M. Bell (2017). Marginal sequential monte carlo for doubly intractable models. arXiv.
- Friel, N., A. Mira, C. J. Oates, et al. (2016). Exploiting multi-core architectures for reduced-variance estimation with intractable likelihoods. Bayesian Analysis 11(1), 215–245.
- Gelman, A. and X.-L. Meng (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. Statistical science, 163–185.

- Gilks, W. R., S. Richardson, and D. Spiegelhalter (1995). Markov chain Monte Carlo in practice. CRC press.
- Majoram, P., J. Molitor, V. Plagnol, and S. Tavaré (2003). Markov chain Monte Carlo without likelihoods. Proceedings of the National Academy of Sciences 100(26), 324–328.
- Marin, J.-M., P. Pudlo, C. P. Robert, and R. J. Ryder (2012). Approximate Bayesian computational methods. Statistics and Computing, 1–14.
- Medina-Aguayo, F. J., A. Lee, and G. O. Roberts (2016). Stability of noisy Metropolis-Hastings. Statistics and Computing 26(6), 1187–1211.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. The journal of chemical physics 21(6), 1087–1092.
- Moores, M., C. Drovandi, K. Mengersen, and C. Robert (2015). Pre-processing for approximate Bayesian computation in image analysis. Statistics and Computing 25(1), 23–33.
- Moores, M. T., A. N. Pettitt, and K. Mengersen (2015). Scalable Bayesian inference for the inverse temperature of a hidden Potts model. arXiv preprint arXiv:1503.08066.
- Murray, I., Z. Ghahramani, and D. MacKay (2006). MCMC for doubly-intractable distributions. In Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence UAI06. Arlington, Virginia, AUAI Press.
- Peskun, P. H. (1973). Optimum Monte Carlo sampling using Markov chains. Biometrika 60(3), 607–612.
- Pritchard, J., M. Seielstad, A. Perez-Lwzaun, and M. Feldman (1999). Population growth of human y chromosomes: a study of y chromosome microsatellites. Molecular Biology and Evolution 16, 1791–1798.
- Propp, J. and D. Wilson (1996). Exactly sampling with coupled Markov chains and applications to statistical mechanics. Random Structures and Algorithms 9, 223–252.
- Quiroz, M., M.-N. Tran, M. Villani, and R. Kohn (2017). Speeding up MCMC by delayed acceptance and data subsampling. Journal of Computational and Graphical Statistics (just-accepted).
- Robbins, H. and S. Monro (1951). A stochastic approximation method. The Annals of Mathematical Statistics 22(3), 400–407.
- Robins, G., P. Pattison, Y. Kalish, and D. Lusher (2007). An introduction to exponential random graph models for social networks. Social Networks 29(2), 169–348.

Rue, H., S. Martino, and N. Chopin (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 71(2), 319–392.

Stoehr, J., A. Benson, and N. Friel (2017). Noisy Hamiltonian Monte Carlo for doubly-intractable distributions. arXiv preprint arXiv:1706.10096.

Tierney, L. (1998). A note on Metropolis-Hastings kernels for general state spaces. Annals of applied probability, 1–9.

Zachary, W. W. (1977). An information flow model for conflict and fission in small groups. Journal of anthropological research, 452–473.

Appendix

Variance of the estimators

Proof of Proposition 1. Denoting by v_n the variance in Eq. (10), it comes

$$nv_n(\theta, \theta') = \mathbb{E}_{\theta'} \exp 2(\theta - \theta')^T s(y) - \left(\frac{Z(\theta)}{Z(\theta')} \right)^2$$

Showing that $nv_n(\theta, \theta') \in \mathcal{O}(\|\theta - \theta'\|^2)$ follows from Taylor expanding the function around θ :

$$nv_n(\theta, \theta') = 1 + 2h^T \mathbb{E}_{\theta'} s(y) + \mathbb{E}_{\theta'} f(h, y) - \frac{1}{Z(\theta')^2} (Z(\theta')^2 + 2Z(\theta')h^T \nabla Z(\theta') + \mathcal{O}(\|h\|^2))$$

where we have introduced $h = \theta - \theta'$ and $f = \mathcal{O}(\|h\|^2)$. Noting that $f = \mathcal{O}(\|h\|^2) \Rightarrow \mathbb{E}_{\theta'} f(h, y) = \mathcal{O}(\|h\|^2)$ and $\nabla Z(\theta')/Z(\theta') = \mathbb{E}_{\theta'} S(y)$ yields:

$$nv_n(\theta, \theta') = \mathcal{O}(\|h\|^2),$$

which concludes the proof. □

Proof of Proposition 2. Note that

$$\begin{aligned} v_n^{\text{DP}} - v_n^{\text{FP}} &= \mathbb{E} \left(\frac{R_n^1}{\Psi_n} \right)^2 \left\{ (R_n^{2C})^2 - (R_n^2 \times \dots \times R_n^C)^2 \right\} \\ &\quad - \left\{ \mathbb{E} \frac{R_n^1}{\Phi_n} R_n^{2C} \right\}^2 + \left\{ \mathbb{E} \frac{R_n^1}{\Phi_n} R_n^2 R_n^3 \times \dots \times R_n^C \right\}^2, \end{aligned}$$

which under the assumption that Φ_n and Ψ_n are independent yields

$$\begin{aligned} v_n^{\text{DP}} - v_n^{\text{FP}} &= \mathbb{E} \left(\frac{R_n^1}{\Psi_n} \right)^2 \left\{ \mathbb{E} (R_n^{2C})^2 - \mathbb{E} (R_n^2 \times \dots \times R_n^C)^2 \right\} \\ &\quad - \left(\mathbb{E} \frac{R_n^1}{\Psi_n} \right)^2 \left\{ (\mathbb{E} R_n^{2C})^2 - (\mathbb{E} R_n^2 \times \dots \times R_n^C)^2 \right\}. \end{aligned}$$

Equation (14) holds with $\gamma = \mathbb{E} (R_n^1/\Psi_n)^2$ as a result of $\mathbb{E}R_n^{2C} = \mathbb{E}R_n^2 \times \cdots \times R_n^C = Z(\dot{\theta}_1)/Z(\dot{\theta}_C)$.

For simplicity of notation, define $R_n := R_n^2 \times \cdots \times R_n^C$ and $X_n = \log R_n$. For large n , R_n^i can be approximate by a truncated normal (in the positive range) $\bar{R}_n^i \sim \mathcal{N}_+(\mu_i, (1/n)\sigma_i^2)$, where $\mu_i := Z(\dot{\theta}_{i-1})/Z(\dot{\theta}_i)$ and $\sigma_i = \text{var}\{\exp(\dot{\theta}_{i-1} - \dot{\theta}_i)^T s(X_i)\}$. It can be noted that, upon reparameterization of the sufficient statistics vector (in the space spanned by the matrix V column vectors), we have $\sigma_i = \text{var}\{\exp \varepsilon s_i\}$ where s_i is the projection on the only one dimension where $\dot{\theta}_{i-1}$ and $\dot{\theta}_i$ are not equal of the sufficient statistics $s(X_i)$, $X_i \sim f(\cdot | \dot{\theta}_i)$. Applying the delta method yields that X_i can be approximate by

$$\bar{X}_n^i := \log \bar{R}_n^i \sim \mathcal{N} \left(\log \mu_i, \frac{\sigma_i^2}{n\mu_i^2} \right). \quad (26)$$

Define $\bar{X}_{n,C} := \sum_{i=1}^{C-1} \bar{X}_n^{i+1}$ and note that the sequence $\{\bar{X}_n^1, \bar{X}_n^2, \dots\}$ satisfies a Lyapunov condition i.e.

$$\lim_{C \rightarrow \infty} \frac{\sum_{i=1}^C \mathbb{E} |\bar{X}_n^i - \mathbb{E} \bar{X}_n^i|^4}{\left\{ \sum_{i=1}^C \text{var}(\bar{X}_n^i) \right\}^4} = 0. \quad (27)$$

Indeed, it can be checked that the fourth central moment of a Gaussian random variable verifies $\mathbb{E} |\bar{X}_n^i - \mathbb{E} \bar{X}_n^i|^4 = 3\text{var}(\bar{X}_n^i)^2$. Moreover since the σ_i 's are bounded, there exists two numbers $0 < m < M < \infty$ such that $m \leq \text{var}(\bar{X}_n^i) \leq M$. This allows to justify (27) since

$$\frac{\sum_{i=1}^C \text{var}(\bar{X}_n^i)^2}{\left\{ \sum_{i=1}^C \text{var}(\bar{X}_n^i) \right\}^4} \leq \frac{\sum_{i=1}^C \text{var}(\bar{X}_n^i)^2}{m^2 C^2 \left\{ \sum_{i=1}^C \text{var}(\bar{X}_n^i) \right\}^2} \leq \frac{1}{C \left\{ \sum_{i=1}^C \text{var}(\bar{X}_n^i) \right\}^2} \left(\frac{M}{m} \right)^2,$$

whose right hand side goes to 0 when $C \rightarrow \infty$. In virtue of (27), a central limit holds for $\bar{X}_{n,C}$ and in particular, asymptotically in C ,

$$\bar{X}_{n,C} \Rightarrow \mathcal{N} \left(\sum_{i=1}^{C-1} \mathbb{E} \bar{X}_n^{i+1}, \sum_{i=1}^{C-1} \text{var} \bar{X}_n^{i+1} \right), \quad (28)$$

which implies that \bar{R}_n is log-normal and, as a consequence,

$$\text{var} \bar{R}_n = \{ \exp(\text{var} \bar{X}_{n,C}) - 1 \} \exp(2\mathbb{E} \bar{X}_{n,C} + \text{var} \bar{X}_{n,C}). \quad (29)$$

First note that combining (26) and (28)

$$\mathbb{E} \bar{X}_{n,C} = \log \frac{Z(\dot{\theta}_1)}{Z(\dot{\theta}_C)}, \quad \text{var} \bar{X}_{n,C} = \frac{1}{n} \sum_{i=2}^C \left\{ \frac{Z(\dot{\theta}_i)}{Z(\dot{\theta}_{i-1})} \right\}^2 \{ \text{var} \exp(\varepsilon s_i) \}^2 \quad (30)$$

and

$$\text{var} \exp(\varepsilon s_i) = \text{var} \left(1 + \varepsilon \sum_{j=1}^{\infty} \frac{\varepsilon^{j-1} s_i^j}{j!} \right) = \varepsilon^2 v_i, \quad v_i := \text{var} \left(\sum_{j=1}^{\infty} \frac{\varepsilon^{j-1} s_i^j}{j!} \right).$$

Putting together with (30), we have:

$$\text{var} \bar{X}_{n,C} = \frac{\varepsilon^4}{n} \sum_{i=2}^C \left\{ v_i \frac{Z(\dot{\theta}_i)}{Z(\dot{\theta}_{i-1})} \right\}^2,$$

which eventually using (29) leads to

$$\text{var} \bar{R}_n = \frac{\varepsilon^4}{n} \sum_{i=2}^C \left\{ v_i \frac{Z(\dot{\theta}_i)Z(\dot{\theta}_1)}{Z(\dot{\theta}_{i-1})Z(\dot{\theta}_C)} \right\}^2 + o(\varepsilon^4/n). \quad (31)$$

□

Remark 1 (On the proof of Proposition 2). *Even though Proposition 1 is established under the assumption that Ψ_n and Φ_n are independent, note that this can be relaxed if the Direct Path estimator includes one more grid point in its path i.e. if Ψ_n^{DP} estimates $Z(\theta)/Z(\dot{\theta}_1) \times Z(\dot{\theta}_1)/Z(\dot{\theta}_{C-1}) \times Z(\dot{\theta}_{C-1})/Z(\dot{\theta}_C)$. When ε is small, we expect that the Direct Path estimator and this alternate version would be highly similar. The result of comparison between the variances of the Full Path estimator and this alternate version of the Direct Path estimator holds without the independence assumption.*

Remark 2 (On the proof of Proposition 2). *Unlike n and ε , the path length C in the Full Path estimator is a random variable that depends on (θ, θ') . Therefore, one can critically comment on the use of a central limit theorem in C that is needed to establish Eq. (16). However, we insist that C could be made arbitrarily as large as needed by using a path connecting θ to θ' that is long enough. This type of path should, however, not use a same grid point twice in order to satisfy the independence assumption of the central limit theorem.*

Convergence of the pre-computing transition kernel We preface the proof of Proposition 4 with the following Lemma.

Lemma 6. *Let $\bar{X}_n^1, \dots, \bar{X}_n^r$ be r iid sample mean estimators i.e. for $j \in \{1, \dots, r\}$, $\bar{X}_n^j = n^{-1} \sum_{k=1}^n X_{j,k}$, $X_{j,1}, \dots, X_{j,n} \sim_{\text{iid}} \pi_j$, where π_j is any distribution. Assume that there exists a positive number $M > 0$ such that for all j , the support of π_j is such that $\text{supp}(\pi_j) \subseteq (0, M)$. Then:*

$$\text{var}(\bar{X}_n^1 \times \dots \times \bar{X}_n^r) \leq M^{2r} \left\{ \left(1 + \frac{1}{n}\right)^r - 1 \right\}.$$

Proof. This follows from the variance of a product of independent random variables. More precisely, $\text{var}(\bar{X}_n^1 \times \dots \times \bar{X}_n^r)$ is a sum of $2^r - 1$ products of r positive factors. Each factor is either a squared expectation $(\mathbb{E} \bar{X}_n^j)^2$ or a variance $\text{var} \bar{X}_n^j$ so that one of the $2^r - 1$ products that contains k variances ($k > 0$) and $r - k$ squared expectations is

$$p_k := \text{var}(\bar{X}_1^n) \text{var}(\bar{X}_2^n) \times \dots \times \text{var}(\bar{X}_k^n) (\mathbb{E} \bar{X}_{k+1}^n)^2 (\mathbb{E} \bar{X}_{k+2}^n)^2 \times \dots \times (\mathbb{E} \bar{X}_r^n)^2.$$

Note that p_k can be reexpressed as

$$p_k = \frac{1}{n^k} \text{var}(X_1) \text{var}(X_2) \times \dots \times \text{var}(X_k) (\mathbb{E} X_{k+1})^2 (\mathbb{E} X_{k+2})^2 \times \dots \times (\mathbb{E} X_r)^2, \quad (32)$$

where for simplicity we have defined $X_j \sim \pi_j$ in Eq. 32. Interestingly, p_k can be uniformly bounded in k as follows:

$$p_k \leq \frac{1}{n^k} \mathbb{E}(X_1^2) \mathbb{E}(X_2^2) \times \cdots \times \mathbb{E}(X_k^2) (\mathbb{E}X_{k+1})^2 (\mathbb{E}X_{k+2})^2 \times \cdots \times (\mathbb{E}X_r)^2 \leq \frac{M^{2r}}{n^k}. \quad (33)$$

Since there are $\binom{r}{k}$ terms that have k variances and $r - k$ squared expectations, their sum \bar{p}_k can be bounded using the uniform bound provided in Eq. (33) so that

$$\bar{p}_k \leq \binom{r}{k} \frac{M^{2r}}{n^k}.$$

The proof is completed by rearranging the sum of $2^r - 1$ products by aggregating those products that have the same number k of variance factors, i.e.

$$\text{var}(\bar{X}_n^1 \times \cdots \times \bar{X}_n^r) = \sum_{k=1}^r \bar{p}_k \leq M^{2r} \sum_{k=1}^r \binom{r}{k} \frac{1}{n^k} = M^{2r} \left(1 + \frac{1}{n}\right)^r - 1.$$

□

Proof of Proposition 4. For notational simplicity, \mathbb{E} is the expectation operator under the distribution of the pre-computed data \mathfrak{U} . Under the assumptions of Lemma 4, the two following constants

$$T := \sup_{\theta \in \Theta} \|\theta\|, \quad \psi_1 := \sup_{x \in \mathcal{Y}} \sup_{j \in \{1, \dots, d\}} s_j(x) \quad (34)$$

are finite. We first state three inequalities that are immediate consequences from the grid geometry:

- Noting that $|(\theta - \theta')^T s(x)| \leq \psi_1 \|\theta - \theta'\|$, for any $(\theta, \theta') \in \Theta^2$ and $x \in \mathcal{Y}$, we have:

$$\exp(-2T\psi_1) \leq \frac{q_\theta(x)}{q_{\theta'}(x)} = \exp(\theta - \theta')^T s(x) \leq \exp(2T\psi_1) := K_1.$$

- For two neighboring points $\dot{\theta}_k$ and $\dot{\theta}_m$ in the pre-computed grid, there exists $j \in \{1, \dots, d\}$ such that $|(\dot{\theta}_k - \dot{\theta}_m)^T s(x)| = \pm \varepsilon s_j(x)$, which yields

$$1/K_2(\varepsilon) \leq \frac{q_{\dot{\theta}_k}(x)}{q_{\dot{\theta}_m}(x)} = \exp(\dot{\theta}_k - \dot{\theta}_m)^T s(x) \leq \exp(\varepsilon\psi_1) := K_2(\varepsilon).$$

- For any $\theta \in \Theta$, there is a point $\dot{\theta} \in \mathfrak{G}$ such that for all $j \in \{1, \dots, d\}$, $|\theta_j - \dot{\theta}_j| < \varepsilon$, we have

$$1/K_2^d(\varepsilon) \leq \frac{q_\theta(x)}{q_{\dot{\theta}_m}(x)} \leq K_2^d(\varepsilon). \quad (35)$$

We will intensively use the result from Lemma 6 on the variance of a product of independent estimators and the fact that for any random variable X

$$\exists M \in \mathbb{R} \quad \text{s.t.} \quad X \leq M \quad \Rightarrow \quad \text{var}X \leq \mathbb{E}X^2 \leq M^2. \quad (36)$$

We recall that in the pre-computing Metropolis algorithm, the normalizing constant ratio $Z(\theta)/Z(\theta')$ is estimated by

$$\rho_n(\theta, \theta', \mathfrak{U}) = \frac{\Psi_n(\theta, \theta', \mathfrak{U})}{\Phi_n(\theta, \theta', \mathfrak{U})}$$

and that

$$\frac{Z(\theta)}{Z(\theta')} = \mathbb{E} \{ \Psi_n(\theta, \theta', \mathfrak{U}) \} / \mathbb{E} \{ \Phi_n(\theta, \theta', \mathfrak{U}) \} . \quad (37)$$

Now for any $(\theta, \theta') \in \Theta^2$, the expectation of the absolute value between the exact and approximate acceptance ratio is

$$\begin{aligned} \mathbb{E} \left| \bar{a}(\theta, \theta', \mathfrak{U}) - a(\theta, \theta') \right| &= \mathbb{E} \left| \frac{h(\theta|\theta')}{h(\theta'|\theta)} \frac{p(\theta')}{p(\theta)} \frac{q_{\theta'}(y)}{q_{\theta}(y)} \left(\rho_n(\theta, \theta', \mathfrak{U}) - \frac{Z(\theta)}{Z(\theta')} \right) \right| \\ &\leq c_p^2 c_h^2 K_1 \mathbb{E} \left| \rho_n(\theta, \theta', \mathfrak{U}) - \frac{Z(\theta)}{Z(\theta')} \right| . \end{aligned} \quad (38)$$

In absence of ambiguity, we let the dependence on $(\theta, \theta', \mathfrak{U})$ of the random variables ρ_n , Φ_n and Ψ_n be implicit. Using (37), we have:

$$\begin{aligned} \mathbb{E} \left| \rho_n - \frac{Z(\theta)}{Z(\theta')} \right| &= \mathbb{E} \left| \Psi_n / \Phi_n - \mathbb{E} \Psi_n / \mathbb{E} \Phi_n \right| , \\ &\leq \mathbb{E} \left| \Psi_n / \Phi_n - \mathbb{E} (\Psi_n / \Phi_n) \right| + \left| \mathbb{E} (\Psi_n / \Phi_n) - \mathbb{E} \Psi_n / \mathbb{E} \Phi_n \right| , \\ &\leq \sqrt{\text{var} \{ \Psi_n / \Phi_n \}} + \left| \text{cov}(\Psi_n, 1/\Phi_n) + \mathbb{E} \Psi_n \mathbb{E} (1/\Phi_n) - \mathbb{E} \Psi_n / \mathbb{E} \Phi_n \right| , \\ &\leq \sqrt{\text{var} \rho_n} + \sqrt{\text{var} \Psi_n \text{var} (1/\Phi_n)} + \mathbb{E} \{ \Psi_n | \mathbb{E} (1/\Phi_n) - 1/\mathbb{E} \Phi_n | \} , \\ &\leq \sqrt{\text{var} \rho_n} + \sqrt{\text{var} \Psi_n \text{var} (1/\Phi_n)} + \mathbb{E} \Psi_n \{ \mathbb{E} (1/\Phi_n) - 1/\mathbb{E} (\Phi_n) \} . \end{aligned} \quad (39)$$

Our objective is now to bound uniformly in (θ, θ') the RHS of Eq. (39). Using Eq. (35), we have that

$$\mathbb{E} \Psi_n \{ \mathbb{E} (1/\Phi_n) - 1/\mathbb{E} (\Phi_n) \} \leq K_2^{C+d-1}(\varepsilon) \left\{ K_2^d(\varepsilon) - \frac{1}{K_2^d(\varepsilon)} \right\} . \quad (40)$$

Defining $\Psi_{1,n} = (1/n) \sum_{k=1}^n \exp(\theta - \hat{\theta}_1)^T s(X_k^1)$ and $\Psi_{2,n} = \Psi_n / \Psi_{1,n}$, note that

$$\begin{aligned} \text{var} \Psi_n &= \text{var} (\Psi_{1,n} \Psi_{2,n}) = \text{var} \Psi_{1,n} \text{var} \Psi_{2,n} + (\mathbb{E} \Psi_{2,n})^2 \text{var} \Psi_{1,n} + (\mathbb{E} \Psi_{1,n})^2 \text{var} \Psi_{2,n} , \\ &= \text{var} \Psi_{2,n} \mathbb{E} \Psi_{1,n}^2 + \text{var} \Psi_{1,n} (\mathbb{E} \Psi_{2,n})^2 . \end{aligned}$$

Applying Lemma 6 to $\Psi_{2,n}$, leads to

$$\text{var} \Psi_{2,n} \leq K_2^{2(C-1)}(\varepsilon) \left\{ \left(1 + \frac{1}{n} \right)^{C-1} - 1 \right\} ,$$

which combined to

- $\mathbb{E} \Psi_{1,n}^2 \leq K_2^{2d}(\varepsilon)$

- $\mathbb{E}\Psi_{2,n} \leq K_2^{C-1}(\varepsilon)$
- $\text{var}\Psi_{1,n} \leq K_2^{2d}(\varepsilon)/n$

yields

$$\text{var}\Psi_n \leq K_2^{2(C+d-1)}(\varepsilon) \left\{ \left(1 + \frac{1}{n}\right)^{C-1} - 1 + \frac{1}{n} \right\}. \quad (41)$$

Finally, combining Eq. 41 and the fact that $\text{var}(1/\Phi_n) \leq K_2^{2d}(\varepsilon)$, we obtain the following bound:

$$\sqrt{\text{var}\Psi_n \text{var}(1/\Phi_n)} \leq K_2^{C+2d-1}(\varepsilon) \sqrt{\left(1 + \frac{1}{n}\right)^{C-1} - 1 + \frac{1}{n}}. \quad (42)$$

Bounding $\text{var}\rho_n$ follows the same technique. Because Φ_n and Ψ_n are not independent, we need to rewrite ρ_n in preparation for applying Lemma 6 as $\rho_n = A_n B_n C_n$ where

$$A_n = \Psi_{1,n}, \quad B_n = \Psi_{2,n} \left/ \frac{1}{n} \sum_{k=1}^n \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X_k^C), \right.$$

$$C_n = \frac{\sum_{k=1}^n \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X_k^C)}{\sum_{k=1}^n \exp(\theta' - \dot{\theta}_C)^T s(X_k^C)}.$$

First note that

$$\begin{aligned} \text{var}\rho_n &= \text{var}A_n B_n \text{var}C_n + \text{var}A_n B_n (\mathbb{E}C_n)^2 + \text{var}C_n (\mathbb{E}A_n B_n)^2 \\ &= \text{var}A_n B_n \mathbb{E}(C_n^2) + (\mathbb{E}A_n B_n)^2 \text{var}C_n. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \text{var}C_n &= \mathbb{E} \left\{ \frac{\sum_{k=1}^n \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X_k^C)}{\sum_{k=1}^n \exp(\theta' - \dot{\theta}_C)^T s(X_k^C)} \right\}^2 - \left\{ \frac{\mathbb{E} \sum_{k=1}^n \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X_k^C)}{\sum_{k=1}^n \exp(\theta' - \dot{\theta}_C)^T s(X_k^C)} \right\}^2, \\ &\leq \frac{K_2^{2d}(\varepsilon)}{n^2} \mathbb{E} \left\{ \sum_{k=1}^n \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X_k^C) \right\}^2 - \frac{1}{K_2^{2d}(\varepsilon)} \left\{ \mathbb{E} \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X^C) \right\}^2, \\ &\leq \frac{K_2^{2d}(\varepsilon)}{n} \text{var} \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X^C) + \left\{ \mathbb{E} \exp(\dot{\theta}_{C-1} - \dot{\theta}_C)^T s(X^C) \right\}^2 \left\{ K_2^{2d}(\varepsilon) - \frac{1}{K_2^{2d}(\varepsilon)} \right\}, \\ &\leq K_2^{2(d+1)}(\varepsilon) \left\{ 1 + \frac{1}{n} - \frac{1}{K_2^{4d}(\varepsilon)} \right\} \quad (43) \end{aligned}$$

and using a similarly technique, we obtain

$$\text{var}A_n B_n \leq K_2^{2(C+d-2)}(\varepsilon) \left\{ \left(1 + \frac{1}{n}\right)^{C-1} - 1 \right\}. \quad (44)$$

Combining Eqs. (43) and (44) with $\mathbb{E}(C_n^2) \leq K_2^{2(d+1)}(\varepsilon)$ and $(\mathbb{E}A_n B_n)^2 \leq K_2^{2(d+C-2)}(\varepsilon)$, we obtain

$$\sqrt{\text{var}\rho_n} \leq K_2^{C+2d-1}(\varepsilon) \sqrt{\left(1 + \frac{1}{n}\right)^{C-1} + \frac{1}{n} - \frac{1}{K_2^{4d}(\varepsilon)}}. \quad (45)$$

Using the bounds derived in Eqs. (40), (42) and (45), Eq. (39) can be written as

$$\begin{aligned} \mathbb{E} \left| \rho_n - \frac{Z(\theta)}{Z(\theta')} \right| &\leq K_2^{C+2d-1}(\varepsilon) \left\{ \sqrt{\left(1 + \frac{1}{n}\right)^{C-1} + \frac{1}{n} - \frac{1}{K_2^{4d}(\varepsilon)}} \right. \\ &\quad \left. + \sqrt{\left(1 + \frac{1}{n}\right)^{C-1} - 1 + \frac{1}{n} + 1 - \frac{1}{K_2^{2d}(\varepsilon)}} \right\}, \\ &\leq 2K_2^{C+2d-1}(\varepsilon) \left\{ \underbrace{\sqrt{\left(1 + \frac{1}{n}\right)^{C-1} + \frac{1}{n} - 1}}_{:=u_n} + \underbrace{\sqrt{1 - \frac{1}{K_2^{4d}(\varepsilon)}}}_{:=v(\varepsilon)} \right\}, \end{aligned}$$

where we have used the fact that for two positive numbers (α, β) , and $\gamma > 1$,

$$\sqrt{\alpha + \beta} \leq \sqrt{\alpha} + \sqrt{\beta} \quad \text{and} \quad 1 - \frac{1}{\gamma} \leq \sqrt{1 - \frac{1}{\gamma^2}}.$$

The proof is completed by noting that $u_n := \sqrt{C/n} + o(n^{-1/2})$ and $v(\varepsilon) = 2\sqrt{d\phi_1}\varepsilon + o(\varepsilon^{1/2})$. \square

Proof of Proposition 5. For n large enough, the delta method shows that the asymptotic distribution of $\{1/\Phi_n\}$ is

$$\frac{1}{\Phi_n} \Rightarrow g := \mathcal{N}\left(\frac{1}{\mathbb{E}(\Phi_1)}, \frac{\text{var}(\Phi_1)}{n\mathbb{E}(\Phi_1)^4}\right). \quad (46)$$

The nice benefit of this observation is that we know that denoting $\{g_n\}_n$ the sequence of distributions of $\{1/\Phi_n\}$, we have that

$$\lim_{n \rightarrow \infty} \int h(x) dg_n(x) = \int h(x) dg(x), \quad (47)$$

for any bounded measurable function h . Defining f_n as the pdf of (Ψ_n, Φ_n) and $\alpha = \mathbb{E}\Psi_n/\mathbb{E}\Phi_n$, the observation (46) motivates rewriting r_n (??) as follows:

$$\begin{aligned} r_n(\theta, \theta') &= \int \left| \frac{\psi}{\phi} - \alpha \right| f_n(d\psi, d\phi) = \int |\psi\phi - \alpha| f_n(d\psi | \phi) g_n(d\phi), \\ &= \underbrace{\int |\psi\phi - \alpha| f_n(d\psi | \phi) (g_n(\phi) - g(\phi)) d\phi}_{r_{n,1}(\theta, \theta')} + \underbrace{\int |\psi\phi - \alpha| f_n(d\psi | \phi) g(d\phi)}_{r_{n,2}(\theta, \theta')}. \quad (48) \end{aligned}$$

The pdfs f_n and g_n implicitly depend on θ and θ' . It is clear that given (47), for any $(\theta, \theta') \in \Theta^2$, $r_{n,1}(\theta, \theta') \rightarrow 0$, although it is not straightforward to obtain a rate of convergence, uniformly in (θ, θ') .

Interestingly, we also have $r_{n,2}(\theta, \theta') \rightarrow 0$ and more precisely defining $W \sim g$ and $\varepsilon \sim \mathcal{N}(0, 1)$, we have:

$$r_{n,2}(\theta, \theta') = \mathbb{E} |\Psi_n W - \alpha| = \mathbb{E} \left| \Psi_n \left\{ \mu_1 + \frac{\sigma_1}{\sqrt{n}} \varepsilon \right\} - \alpha \right|, \quad (49)$$

where we have defined $\mu_1 = 1/\mathbb{E}(\Phi_1)$ and $\sigma_1^2 = \text{var}(\Phi_1)/\mathbb{E}(\Phi_1)^4$. This yields

$$\begin{aligned}
r_{n,2}(\theta, \theta') &= \mathbb{E} \left| \mu_1 \Psi_n + \frac{\sigma_1}{\sqrt{n}} \Psi_n \varepsilon - \alpha \right| \\
&\leq \frac{1}{\mathbb{E}(\Phi_1)} \mathbb{E} |\Psi_n - \mathbb{E}(\Psi_1)| + \frac{\sigma_1}{\sqrt{n}} \mathbb{E}(\Psi_n |\varepsilon|) \\
&\leq \sqrt{\frac{\text{var}(\Psi_1)}{n\mathbb{E}(\Phi_1)^2}} + \sqrt{\frac{\text{var}(\Phi_1)}{n\mathbb{E}(\Phi_1)^4}} \mathbb{E}(\Psi_n |\varepsilon|) \\
&= \frac{1}{\mathbb{E}(\Phi_1)\sqrt{n}} \left\{ \sqrt{\text{var}(\Psi_1)} + \sqrt{\text{var}(\Phi_1)} \frac{\mathbb{E}(\Psi_n |\varepsilon|)}{\mathbb{E}(\Phi_1)} \right\}. \quad (50)
\end{aligned}$$

Summarizing we have the following upper bound for r_n :

$$\begin{aligned}
r_n(\theta, \theta') &\leq \int |\psi\phi - \alpha| f_n(d\psi | \phi)(g_n(\phi) - g(\phi))d\phi + \\
&\quad \frac{1}{\mathbb{E}(\Phi_1)\sqrt{n}} \left\{ \sqrt{\text{var}(\Psi_1)} + \sqrt{\text{var}(\Phi_1)} \frac{\mathbb{E}(\Psi_n |\varepsilon|)}{\mathbb{E}(\Phi_1)} \right\}. \quad (51)
\end{aligned}$$

□