



<b>Title</b>	Noisy Hamiltonian Monte Carlo for Doubly Intractable Distributions
<b>Authors(s)</b>	Stoehr, Juilian, Benson, Alan, Friel, Nial
<b>Publication date</b>	2018-10-29
<b>Publication information</b>	Stoehr, Juilian, Alan Benson, and Nial Friel. "Noisy Hamiltonian Monte Carlo for Doubly Intractable Distributions" 28, no. 1 (October 29, 2018).
<b>Publisher</b>	Taylor & Francis
<b>Item record/more information</b>	<a href="http://hdl.handle.net/10197/10489">http://hdl.handle.net/10197/10489</a>
<b>Publisher's statement</b>	This is an electronic version of an article published in Stoehr, J., Benson, A. and Friel, N. Noisy Hamilton Monte Carlo for Doubly Intractable Distributions. Journal of Computational and Graphical Statistics, 28(1). Journal of Computational and Graphical Statistics is available online at <a href="http://www.tandfonline.com/doi/abs/10.1080/10618600.2018.1506346">www.tandfonline.com/doi/abs/10.1080/10618600.2018.1506346</a>
<b>Publisher's version (DOI)</b>	10.1080/10618600.2018.1506346

Downloaded 2023-03-15T17:09:45Z

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

# NOISY HAMILTONIAN MONTE CARLO FOR DOUBLY-INTRACTABLE DISTRIBUTIONS

JULIEN STOEHR, ALAN BENSON, NIAL FRIEL

School of Mathematics and Statistics, University College Dublin and Insight  
Centre for Data Analytics

March 26, 2018

## Abstract

Hamiltonian Monte Carlo (HMC) has been progressively incorporated within the statistician's toolbox as an alternative sampling method in settings when standard Metropolis-Hastings is inefficient. HMC generates a Markov chain on an augmented state space with transitions based on a deterministic differential flow derived from Hamiltonian mechanics. In practice, the evolution of Hamiltonian systems cannot be solved analytically, requiring numerical integration schemes. Under numerical integration, the resulting approximate solution no longer preserves the measure of the target distribution, therefore an accept-reject step is used to correct the bias. For doubly-intractable distributions – such as posterior distributions based on Gibbs random fields – HMC suffers from some computational difficulties: computation of gradients in the differential flow and computation of the accept-reject proposals poses difficulty. In this paper, we study the behaviour of HMC when these quantities are replaced by Monte Carlo estimates.

*Keywords:* Bayesian inference, Markov random fields, Markov chain Monte Carlo.

## 1 Introduction

Developing satisfactory methodology for the Bayesian analysis of statistical models with intractable likelihood functions is of considerable interest. Such models are motivated by a wide range of applications including spatial statistics, social network analysis, population genetics and image analysis. The challenges raised by such models stem from mathematical reasons – the likelihood function does not admit a closed form as a function of  $\theta$  – or computational reasons – the likelihood function evaluation is time

consuming. Such issues appear in settings where the likelihood function is known up to a parameter dependent normalising constant, that is,

$$\frac{1}{Z(\boldsymbol{\theta})}q(\boldsymbol{\theta}, \mathbf{x}), \quad \text{where} \quad Z(\boldsymbol{\theta}) = \int_{\mathcal{X}} q(\boldsymbol{\theta}, \mathbf{x})\mu(d\mathbf{x}), \quad (1.1)$$

which is the focus of this paper. One such class of models are Markov random fields (MRFs), undirected graphical models such as the Ising or Potts model which are used in a wide range of applications to model the dependency structure of correlated data. Applications may be found in epidemiology (Green and Richardson, 2002), genetic analysis (François et al., 2006), ecology (Augustin et al., 1998), image analysis (*e.g.* Hurn et al., 2003), amongst others. In social network analysis, the exponential random graph model (ERGM) (Robins et al., 2007) can be used to model the structure of a social network represented using a directed or undirected graph where edges in the graph show connections between the nodes, *e.g.*, friendship. For ERGMs, the likelihood function (1.1) is constructed over binary adjacency matrices  $\mathbf{x} \in \mathcal{X}$  which represent the graph and the likelihood function is intractable since the number of graphs in a network with  $n$  actors grows as  $2^{\binom{n}{2}}$  for the undirected edge networks and  $2^{n(n-1)}$  in the directed edge networks.

Dealing with intractable likelihood functions has led to the development of important theoretical and methodological advancements in Bayesian statistics. A first approach to overcome the model’s intractable bottleneck is to replace the true model with a pseudo-model selected among a collection of much simpler and more tractable set of probability distributions (*e.g.* variational Bayes (Jaakkola and Jordan, 2000), mean-fields approximation for MRFs (Jordan et al., 1999) or with an easily-normalised full conditional distributions (*e.g.* composite likelihood (Lindsay, 1988) and pseudolikelihood for MRFs (Besag, 1974)). However, these cruder approximations to the true model often miss some of the features of the original intractable model and can lead to unreliable estimates of the model parameters (Friel, 2012, Stoehr and Friel, 2015).

Another point of view arises from sampling methods. Markov Chain Monte Carlo (MCMC) methods are arguably the most popular methodology, but Approximate Bayesian Computation (ABC, Marin et al., 2012), another simulation based approach has recently generated much activity in the literature. ABC deals with the situation where the likelihood function cannot be evaluated, due to intractability or otherwise, but can be simulated from. When performing parameter estimation, the method is particularly well suited for problems where the likelihood function does not admit an algebraic form, a situation where MCMC methods are at a loss but which we don’t explore here. MCMC methods produce ergodic estimates of functionals with respect to the posterior of interest such as the expectation of parameters and so on. While some Bayesian estimators can be efficiently estimated with such methods via the empirical distribution, most methods cannot be used with a parameter dependent intractable likelihood function. Indeed, to produce a Markov chain which is reversible with respect to the posterior distribution, the method performs at each iteration an accept-reject step which requires computation of a ratio of intractable normalising constants. This problem is sometimes referred to as doubly-intractable Bayesian inference. Murray et al. (2006), extending the work

of [Møller et al. \(2006\)](#), derive MCMC methods to handle the presence of intractable normalising terms in the acceptance ratio. The target distribution can be estimated without bias by using auxiliary variables whose proposal distributions have the relevant normalising constant. The latter solution leads to replacing the acceptance probability of the accept-reject step by a single point importance sampling estimate. Such ideas have appeared in the generalised importance Metropolis-Hastings of [Beaumont \(2003\)](#) and have later been extended by [Andrieu and Roberts \(2009\)](#).

For well-known MCMC methods, such as the Metropolis-Hastings algorithm ([Metropolis et al., 1953](#), [Hastings, 1970](#)), transitions of the Markov chain are driven by a random walk and exploring the parameter space in this manner can be quite inefficient. Indeed, it is delicate to propose large transitions across the parameter space that will be accepted with high probability. For small transitions, converging to the target distribution may require an excessive amount of time. Such methods hence can exhibit low acceptance rates, poor mixing and highly correlated samples ([Robert and Casella, 2004](#)). Surged by the development of the software package `Stan` ([Carpenter et al., 2016](#)), Hamiltonian Monte Carlo (HMC, [Duane et al., 1987](#), [Neal, 2011](#)) has emerged as a reliable alternative for sampling in general settings. The method relies on Hamiltonian dynamics to produce large transition paths across the parameter space. In practice, the transition paths are approximated using a gradient based numerical integrator. In order to preserve the target measure, the numerical solutions must be corrected using an accept-reject step. Both constructing transition paths and the accept-reject step form a central issue in this paper as they are unavailable for doubly-intractable distributions. In this paper, we explore the opportunity of using a “noisy” version of the HMC sampling method and apply this “noisy” scheme to the Potts model and to the ERGM.

The paper begins with a review of HMC in Section 2. We then introduce the noisy version of the algorithm in Section 3. The latter relies on Monte Carlo estimates of gradients, used in the numerical integration, and unbiased importance sampling estimates of the intractable ratio involved in the accept-reject step. Our approach contrasts with the finite difference schemes or exact derivatives usually used in HMC, which are not available. Furthermore, we use all the intermediate points visited by the integrator to derive our importance sampling estimator. This leads to a more robust estimate than the point estimates of [Murray et al. \(2006\)](#) and extended by [Alquier et al. \(2016\)](#) in their work on noisy MCMC. We end the paper with a detailed numerical study to intractable likelihood problems in Section 4.

## 2 Background on Hamiltonian Monte Carlo

Consider a probability measure  $\pi$  on  $\Theta \subset \mathbb{R}^d$  with density, also denoted  $\pi$ , with respect to the Lebesgue measure,

$$\pi(\boldsymbol{\theta}) = \frac{\exp\{-V(\boldsymbol{\theta})\}}{\int_{\Theta} \exp\{-V(\tilde{\boldsymbol{\theta}})\} d\tilde{\boldsymbol{\theta}}},$$

where  $V$  is continuously differentiable. Markov chain Monte Carlo provides a very general framework to allow estimation of functionals of the form

$$\int_{\Theta} g(\boldsymbol{\theta})\pi(d\boldsymbol{\theta}),$$

for some function  $g$  by generating a Markov chain  $(\boldsymbol{\theta}_n)_{n \in \mathbb{N}}$  with transition kernel  $P$  which leaves  $\pi$  invariant. Hamiltonian Monte Carlo also belongs to the MCMC toolbox and involves augmenting the target distribution with an auxiliary variable  $\mathbf{r} \in \mathbb{R}^d$ , usually referred to as a momentum variable, whose density is a  $d$ -dimensional normal distribution with mean  $\mathbf{0}_d$  and covariance matrix  $\mathbf{M}$ , denoted  $\mathcal{N}(\cdot \mid \mathbf{0}_d, \mathbf{M})$  in what follows. HMC thus samples from the augmented distribution

$$\tilde{\pi}(\boldsymbol{\theta}, \mathbf{r}) = \pi(\boldsymbol{\theta})\mathcal{N}(\mathbf{r} \mid \mathbf{0}_d, \mathbf{M}),$$

whose marginal chain in  $\boldsymbol{\theta}$  is the distribution of interest. The method originally appeared in statistical physics (Duane et al., 1987) before being more widely applied for statistical inference. We refer the reader to Neal (2011) for a comprehensive review. Consider the unnormalised negative joint log-density

$$H(\boldsymbol{\theta}, \mathbf{r}) = V(\boldsymbol{\theta}) + \frac{1}{2}\mathbf{r}^T\mathbf{M}^{-1}\mathbf{r}.$$

The method consists of generating proposals for  $\boldsymbol{\theta}$  based on the canonical Hamilton's equations which write with respect to a fictitious time  $t$

$$\frac{d\boldsymbol{\theta}}{dt} = \frac{\partial H}{\partial \mathbf{r}} = \mathbf{M}^{-1}\mathbf{r} \quad \text{and} \quad \frac{d\mathbf{r}}{dt} = -\frac{\partial H}{\partial \boldsymbol{\theta}} = -\nabla_{\boldsymbol{\theta}}V(\boldsymbol{\theta}), \quad (2.1)$$

where  $\nabla_{\boldsymbol{\theta}} = [\partial/\partial\theta_1, \dots, \partial/\partial\theta_d]^T$  denotes the gradient operator. The main reason for relying on such a mechanism is to efficiently explore the target density  $\pi$  by proposing a new state far from the current state while preserving the measure  $\tilde{\pi}$  (Neal, 2011). In particular, the marginal Markov chain on  $\Theta$  is invariant with respect to the target distribution  $\pi$ .

In practice, the differential equations (2.1) cannot be solved analytically, requiring numerical integration to approximate the solution flow. The most popular numerical integration scheme, if only for its good tradeoff between accuracy and computational cost, is the second order Störmer-Verlet or leapfrog integrator. Consider a time-step  $\epsilon$  and the following transformations

$$g_{1,\epsilon} : (\boldsymbol{\theta}, \mathbf{r}) \rightarrow (\boldsymbol{\theta} + \epsilon\mathbf{M}^{-1}\mathbf{r}, \mathbf{r}) \quad \text{and} \quad g_{2,\epsilon} : (\boldsymbol{\theta}, \mathbf{r}) \rightarrow (\boldsymbol{\theta}, \mathbf{r} - \epsilon/2\nabla_{\boldsymbol{\theta}}V(\boldsymbol{\theta})).$$

The leapfrog integrator yields a map  $F_{\epsilon} : (\boldsymbol{\theta}, \mathbf{r}) \rightarrow (\boldsymbol{\theta}', \mathbf{r}')$  defined by  $F_{\epsilon} = g_{2,\epsilon} \circ g_{1,\epsilon} \circ g_{2,\epsilon}$ . Put another way, the scheme decomposes into the following three-stage procedure:

$$\tilde{\mathbf{r}} = \mathbf{r} - \frac{\epsilon}{2}\nabla_{\boldsymbol{\theta}}V(\boldsymbol{\theta}), \quad \boldsymbol{\theta}' = \boldsymbol{\theta} + \epsilon\mathbf{M}^{-1}\tilde{\mathbf{r}}, \quad \mathbf{r}' = \tilde{\mathbf{r}} - \frac{\epsilon}{2}\nabla_{\boldsymbol{\theta}}V(\boldsymbol{\theta}').$$

The map  $F_{\epsilon}$  approximates the solution at time  $\epsilon$  and to get the approximated solution at a time  $t$ , one iterates it  $L = \lfloor \frac{t}{\epsilon} \rfloor$  times, referred to as number of leapfrog steps.

**Algorithm 1:** Hamiltonian Monte Carlo (single iteration)

**Input:** the current state of the chain  $\boldsymbol{\theta} := \boldsymbol{\theta}_0$ , a step size  $\epsilon$ , a number of Leapfrog steps  $L$ .

**draw**  $\mathbf{r} := \mathbf{r}_0$  from  $\mathcal{N}(\mathbf{0}_d, \mathbf{M})$ ;

**for**  $\ell \leftarrow 1$  **to**  $L$  **do**

  | **compute**  $\{\boldsymbol{\theta}_\ell, \mathbf{r}_\ell\} = g_{2,\epsilon} \circ g_{1,\epsilon} \circ g_{2,\epsilon}(\boldsymbol{\theta}_{\ell-1}, \mathbf{r}_{\ell-1})$ ;

**end**

**set**  $(\boldsymbol{\theta}', \mathbf{r}') = (\boldsymbol{\theta}_L, -\mathbf{r}_L)$  with probability  $1 \wedge \exp\{H(\boldsymbol{\theta}_0, \mathbf{r}_0) - H(\boldsymbol{\theta}_L, -\mathbf{r}_L)\}$ ;

**set**  $(\boldsymbol{\theta}', \mathbf{r}') = (\boldsymbol{\theta}_0, -\mathbf{r}_0)$  otherwise;

Nevertheless, the approximated flow  $F_\epsilon$  does not preserve the measure  $\tilde{\pi}(d\boldsymbol{\theta}, d\mathbf{r})$ . To correct the bias introduced, an accept-reject step is used (see Algorithm 1) which following a Metropolis-Hastings algorithm, results in a transition from  $(\boldsymbol{\theta}, \mathbf{r})$  to  $(\boldsymbol{\theta}', -\mathbf{r}')$  accepted with probability

$$\rho(\boldsymbol{\theta}, \mathbf{r}, \boldsymbol{\theta}', \mathbf{r}') = 1 \wedge \exp\{H(\boldsymbol{\theta}, \mathbf{r}) - H(\boldsymbol{\theta}', \mathbf{r}')\}.$$

The transition kernel of the Metropolis-Hastings update satisfies detailed balance since the deterministic mapping  $T \circ F_\epsilon$ , where  $T : (\boldsymbol{\theta}, \mathbf{r}) \rightarrow (\boldsymbol{\theta}, -\mathbf{r})$ , is an involution on  $\Theta \times \Theta$  (Tierney, 1998). It is a direct consequence of time reversibility of the approximated flow  $F_\epsilon$  – as each leapfrog step is reversible by negating  $\epsilon$  – and its volume preserving property – as Jacobians of transformations  $g_{1,\epsilon}$  and  $g_{2,\epsilon}$  have unit determinant.

### 3 HMC for doubly-intractable distribution

Consider the target distribution  $\pi$  being a Bayesian posterior distribution expressed as

$$\pi(\boldsymbol{\theta} \mid \mathbf{x}) \propto f(\mathbf{x} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta}), \quad (3.1)$$

where  $p(\boldsymbol{\theta})$  denotes a prior density on the parameter space  $\Theta$  with respect to a reference measure (often the Lebesgue measure of the Euclidean space) and  $f(\mathbf{x} \mid \boldsymbol{\theta})$  denotes the likelihood of the observed data  $\mathbf{x} \in \mathcal{X}$ . Here we are concerned with the situation where the unnormalised posterior distribution, the right-hand-side of (3.1), is intractable. In particular, we focus on likelihood models of the form

$$\mathfrak{P} = \left\{ f(\mathbf{x} \mid \boldsymbol{\theta}) = \frac{\exp\{A(\boldsymbol{\theta}, \mathbf{x})\}}{Z(\boldsymbol{\theta})} := \frac{q(\boldsymbol{\theta}, \mathbf{x})}{Z(\boldsymbol{\theta})} \mid \boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^d, A(\cdot, \mathbf{x}) \in \mathcal{C}^1(\Theta) \right\}, \quad (3.2)$$

where the parameter dependent likelihood normalising constant,  $Z(\boldsymbol{\theta})$ , is intractable. Gibbs random fields represent such a class of intractable likelihood models and are the focus of Section 4. This complication results in what is often termed a doubly-intractable posterior distribution, since the posterior distribution itself is normalised by the evidence (or marginal likelihood) which is typically also intractable. In this context, a direct implementation of HMC is not feasible for two reasons:

1. The mapping  $g_{2,\epsilon}$ , and more precisely the gradient  $\nabla_{\boldsymbol{\theta}} V(\boldsymbol{\theta}) = -\nabla_{\boldsymbol{\theta}} \log f(\mathbf{x} | \boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{\theta})$  is analytically intractable, see Section 3.1,
2. The accept-reject step in Algorithm 1 is unavailable for doubly-intractable Bayesian analysis as it requires an evaluation of a ratio of intractable normalising constants, see Section 3.2.

In what follows, we propose to overcome these two issues by considering Monte Carlo estimates of both the gradient of the log target and the ratio of intractable normalising constants. Moreover, both of these quantities can be estimated by simulating from the likelihood model, as we will now show in detail.

### 3.1 Gradient estimates

Closed-form gradients for complex models are typically out of reach. Computing the gradient in  $g_{2,\epsilon}$  is usually addressed using automatic differentiation as in the software package `Stan` (Carpenter et al., 2016). However point-wise estimation is impossible for the likelihood model described in (3.2) and therefore we require another approach. Here we note that the gradient of the log-posterior distribution (3.1) can be written as

$$\nabla_{\boldsymbol{\theta}} \log \pi(\boldsymbol{\theta} | \mathbf{x}) = \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{x}) - \nabla_{\boldsymbol{\theta}} \log Z(\boldsymbol{\theta}) + \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{\theta}), \quad (3.3)$$

Forward-simulations from the likelihood taken at each leapfrog step can be used to provide a Monte Carlo estimate of the gradient, using the following identity,

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \log Z(\boldsymbol{\theta}) &= \frac{1}{Z(\boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} Z(\boldsymbol{\theta}) \\ &= \frac{1}{Z(\boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} \int_{\mathcal{X}} \exp \{A(\boldsymbol{\theta}, \mathbf{x})\} \mu(d\mathbf{x}) \\ &= \int_{\mathcal{X}} \frac{\exp \{A(\boldsymbol{\theta}, \mathbf{x})\}}{Z(\boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{x}) \mu(d\mathbf{x}) \\ &= \mathbf{E}_{\boldsymbol{\theta}} \{ \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{X}) \}. \end{aligned} \quad (3.4)$$

So far, we have only assumed that  $A(\cdot, \mathbf{x})$  is continuously differentiable on  $\Theta$ . However this identity holds under regularity conditions which allow one to switch the derivative and integral operators (the domain  $\mathcal{X}$  of  $\mathbf{X}$  is assumed to be independent of  $\boldsymbol{\theta}$ ) and under the assumption that  $\nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{X})$  is integrable with respect to  $f(\mathbf{x} | \boldsymbol{\theta}) \mu(d\mathbf{x})$ . Using Monte Carlo samples  $\{\mathbf{u}^{(1,\boldsymbol{\theta})}, \dots, \mathbf{u}^{(N,\boldsymbol{\theta})}\}$  from  $f(\cdot | \boldsymbol{\theta})$ , the expected value (3.4) can be estimated using the empirical mean of the random variable  $\nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{X})$  over the sample. This leads to the following estimate of the gradient of the log-posterior (3.3) at  $\boldsymbol{\theta}$

$$\widehat{\nabla}_{\mathbf{u}^{(1,\boldsymbol{\theta})}, \dots, \mathbf{u}^{(N,\boldsymbol{\theta})}} \log \pi(\boldsymbol{\theta} | \mathbf{x}) := \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{x}) - \frac{1}{N} \sum_{n=1}^N \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}, \mathbf{u}^{(n,\boldsymbol{\theta})}) + \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{\theta}). \quad (3.5)$$

### 3.2 Metropolis-Hastings ratio estimates

The intractability of the likelihood model in (3.2) implies, in particular, that the standard MCMC toolbox is infeasible. For example, a naive implementation of Algorithm 1 when proposing to move from  $(\boldsymbol{\theta}, \mathbf{r})$  to  $(\boldsymbol{\theta}', \mathbf{r}')$  requires the computation of the unknown normalising constants,  $Z(\boldsymbol{\theta})$  and  $Z(\boldsymbol{\theta}')$ ,

$$\rho(\boldsymbol{\theta}, \mathbf{r}, \boldsymbol{\theta}', \mathbf{r}') = 1 \wedge \frac{Z(\boldsymbol{\theta})}{Z(\boldsymbol{\theta}')} \frac{q(\boldsymbol{\theta}', \mathbf{x})}{q(\boldsymbol{\theta}, \mathbf{x})} \frac{\mathcal{N}(\mathbf{r}' | \mathbf{0}_d, \mathbf{M})}{\mathcal{N}(\mathbf{r} | \mathbf{0}_d, \mathbf{M})} \frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})}. \quad (3.6)$$

The exchange algorithm (Murray et al., 2006), extending the work of Møller et al. (2006), is a popular MCMC method to allow sampling from doubly-intractable distributions. Denote  $\nu(\boldsymbol{\theta}' | \boldsymbol{\theta})$  the proposal distribution to move from  $\boldsymbol{\theta}$  to  $\boldsymbol{\theta}'$ , the exchange algorithm samples from an augmented distribution

$$\pi(\boldsymbol{\theta}', \mathbf{u}', \boldsymbol{\theta} | \mathbf{x}) \propto f(\mathbf{x} | \boldsymbol{\theta}) p(\boldsymbol{\theta}) \nu(\boldsymbol{\theta}' | \boldsymbol{\theta}) f(\mathbf{u}' | \boldsymbol{\theta}').$$

whose marginal distribution in  $\boldsymbol{\theta}$  is the posterior distribution of interest. Murray et al. (2006) present a clever Metropolis-within-Gibbs algorithm to sample from this augmented distribution : given the current value  $\boldsymbol{\theta}$ , one iteratively samples  $\boldsymbol{\theta}'$  from its full-conditional  $\nu(\cdot | \boldsymbol{\theta})$ , then  $\mathbf{u}'$  from the intractable distribution  $f(\cdot | \boldsymbol{\theta}')$  and proposes to deterministically swap  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}'$  through a Metropolis-Hastings step. It turns out that the ratio of intractable normalising constants drops out of the Metropolis acceptance probability

$$1 \wedge \frac{\cancel{Z(\boldsymbol{\theta})} \cancel{Z(\boldsymbol{\theta}')} q(\boldsymbol{\theta}, \mathbf{u}')}{\cancel{Z(\boldsymbol{\theta}')} q(\boldsymbol{\theta}', \mathbf{u}') \cancel{Z(\boldsymbol{\theta})}} \frac{q(\boldsymbol{\theta}', \mathbf{x}) \nu(\boldsymbol{\theta} | \boldsymbol{\theta}') p(\boldsymbol{\theta}')}{q(\boldsymbol{\theta}, \mathbf{x}) \nu(\boldsymbol{\theta}' | \boldsymbol{\theta}) p(\boldsymbol{\theta})}.$$

Murray et al. (2006) point out that the fraction  $q(\boldsymbol{\theta}, \mathbf{u}')/q(\boldsymbol{\theta}', \mathbf{u}')$  which appears above, can be considered as an single sample importance estimator of  $Z(\boldsymbol{\theta})/Z(\boldsymbol{\theta}')$  since it holds that

$$\mathbf{E}_{\boldsymbol{\theta}'} \left\{ \frac{q(\boldsymbol{\theta}, \mathbf{U}')}{q(\boldsymbol{\theta}', \mathbf{U}')} \right\} = \frac{Z(\boldsymbol{\theta})}{Z(\boldsymbol{\theta}')}, \quad (3.7)$$

where  $\mathbf{E}_{\boldsymbol{\theta}'}$  is the expectation with respect to  $\mathbf{U}' \sim f(\cdot | \boldsymbol{\theta}')$ . In fact Alquier et al. (2016), consider a generalised exchange algorithm based, at each step of the algorithm, on an improved unbiased estimate of  $Z(\boldsymbol{\theta})/Z(\boldsymbol{\theta}')$  including multiple auxiliary draws with respect to the proposed parameter, namely,

$$\frac{\widehat{Z(\boldsymbol{\theta})}}{Z(\boldsymbol{\theta}')} = \frac{1}{N} \sum_{n=1}^N \frac{q(\boldsymbol{\theta}, \mathbf{u}^{(n, \boldsymbol{\theta}')} )}{q(\boldsymbol{\theta}', \mathbf{u}^{(n, \boldsymbol{\theta}')} )}, \quad (3.8)$$

where the auxiliary variables  $\{\mathbf{u}^{(1, \boldsymbol{\theta}')} , \dots , \mathbf{u}^{(N, \boldsymbol{\theta}')} \}$  are drawn from  $f(\cdot | \boldsymbol{\theta}')$ . However this so-called noisy exchange algorithm no longer leaves the target distribution invariant, nevertheless it is possible to provide convergence guarantees that the resulting Markov chain is close in some sense to the target distribution. An alternative to previous methods



presented and one which we do not explore here but is worth mentioning is Russian Roulette sampling (Lyne et al., 2015) which can be used to get an unbiased estimate of  $1/Z(\boldsymbol{\theta})$ .

Nevertheless, the strategy of using such importance sampling estimates (ISE) in HMC framework is questionable. In particular, the importance sampling weights  $q(\boldsymbol{\theta}, \mathbf{u})/q(\boldsymbol{\theta}', \mathbf{u})$  can lead to unreliable estimates of the intractable ratio  $Z(\boldsymbol{\theta})/Z(\boldsymbol{\theta}')$  for large transitions between  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta}' \sim \nu(\cdot | \boldsymbol{\theta})$ , see Figure 1. In what follows, we develop an alternative importance sampling estimator which is compatible with the integrator path.

### 3.3 Noisy Hamiltonian Monte Carlo

To deal with the different difficulties of doubly-intractable Bayesian analysis, we derive a noisy version of HMC, see Algorithm 2. Consider the set of points  $\{\boldsymbol{\theta} = \boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}' = \boldsymbol{\theta}_L\}$  visited by the symplectic integrator in Algorithm 1. At each leapfrog step  $\ell$ , we perform  $N$  auxiliary draws with respect to the current parameter value  $\boldsymbol{\theta}_\ell$  to compute surrogates of mapping  $g_{2,\epsilon}$  using gradient estimates (3.5). Auxiliary draws can then be reused to compute at not cost the Metropolis-Hastings proposal (3.6) using importance sampling estimator (3.8). The transition kernel from Algorithm 1 is then replaced by an approximated version arising from Algorithm 2 based on stochastic estimators. So far there are no theoretical guarantees regarding the effect of the approximated kernel on the limiting distribution and mixing properties. Nonetheless Alquier et al. (2016) give some theoretical results in the particular case of the Langevin algorithm, that is when  $L = 1$ . In addition, the work of Chen et al. (2014) who establish some results when the gradient of the target distribution is estimated using mini-batches of the data may prove useful.

In Algorithm 2, different strategies can be adopted to estimate the acceptance probability (3.6). An obvious solution would be to plug in the ISE (3.8) which solely uses auxiliary draws with respect to the proposed value  $\boldsymbol{\theta}' = \boldsymbol{\theta}^{(L)}$  and discard all others. Nevertheless, such a solution turns out to be quite inefficient for large transition. For an illustrative purpose, consider the absolute error of the log ratio estimates

$$\text{abs.err}(\boldsymbol{\theta}, \boldsymbol{\theta}') = \left| \log \frac{Z(\boldsymbol{\theta})}{Z(\boldsymbol{\theta}')} - \log \frac{\widehat{Z}(\boldsymbol{\theta})}{\widehat{Z}(\boldsymbol{\theta}')} \right|$$

at each iteration of HMC. We examined the latter for a particular distribution in  $\mathfrak{P}$ , namely a Potts model (see Section 4.1 for details) defined on a regular lattice for which we can compute exactly the log ratio  $Z(\boldsymbol{\theta})/Z(\boldsymbol{\theta}')$  using the R-package GiRaF (Stoehr et al., 2016). Figure 1 shows absolute errors obtained for 3,500 pairs  $(\boldsymbol{\theta}, \boldsymbol{\theta}')$  with respect to the distance in  $L^2$ -norm between the current value  $\boldsymbol{\theta}$  and the proposed one  $\boldsymbol{\theta}'$ . Two conclusions can be drawn from first row of Figure 1. On one hand, a somewhat naive conclusion but consistent with the theory is that the importance sampling estimate is all the more precise that we use multiple auxiliary draws. Table 1 presenting the mean squared error of the estimator emphasises this. On the other hand, the quality of the

**Algorithm 2:** Noisy HMC (single iteration)

**Input:** the current state of the chain  $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ , the observed dataset  $\mathbf{x}$ , a step size  $\epsilon$ , a number of Leapfrog steps  $L$ , a number of auxiliary draws  $N$ .

**draw**  $\mathbf{r}$  from  $\mathcal{N}(\mathbf{0}_d, \mathbf{M})$ ;

**draw** auxiliary sample  $\{\mathbf{u}^{(1, \boldsymbol{\theta}_0)}, \dots, \mathbf{u}^{(N, \boldsymbol{\theta}_0)}\}$  from  $f(\cdot \mid \boldsymbol{\theta}_0)$ ;

**compute**  $\mathbf{r}_0 = \mathbf{r} + \frac{\epsilon}{2} \widehat{\nabla}_{\mathbf{u}^{(1, \boldsymbol{\theta}_0)}, \dots, \mathbf{u}^{(N, \boldsymbol{\theta}_0)}} \log \pi(\boldsymbol{\theta}_0 \mid \mathbf{x})$ ;

**for**  $\ell \leftarrow 1$  **to**  $L$  **do**

**compute**  $\boldsymbol{\theta}_\ell = \boldsymbol{\theta}_{\ell-1} + \epsilon \mathbf{M}^{-1} \mathbf{r}_{\ell-1}$ ;

**draw** auxiliary sample  $\{\mathbf{u}^{(1, \boldsymbol{\theta}_\ell)}, \dots, \mathbf{u}^{(N, \boldsymbol{\theta}_\ell)}\}$  from  $f(\cdot \mid \boldsymbol{\theta}_\ell)$ ;

**if**  $\ell < L$  **then**

**compute**  $\mathbf{r}_\ell = \mathbf{r}_{\ell-1} + \epsilon \widehat{\nabla}_{\mathbf{u}^{(1, \boldsymbol{\theta}_\ell)}, \dots, \mathbf{u}^{(N, \boldsymbol{\theta}_\ell)}} \log \pi(\boldsymbol{\theta}_\ell \mid \mathbf{x})$ ;

**else**

**compute**  $\mathbf{r}_\ell = \mathbf{r}_{\ell-1} + \frac{\epsilon}{2} \widehat{\nabla}_{\mathbf{u}^{(1, \boldsymbol{\theta}_\ell)}, \dots, \mathbf{u}^{(N, \boldsymbol{\theta}_\ell)}} \log \pi(\boldsymbol{\theta}_\ell \mid \mathbf{x})$ ;

**end**

**compute**

$$\frac{Z(\widehat{\boldsymbol{\theta}_{\ell-1}})}{Z(\boldsymbol{\theta}_\ell)} = \frac{1}{N} \sum_{n=1}^N \frac{q(\boldsymbol{\theta}_{\ell-1}, \mathbf{u}^{(n, \boldsymbol{\theta}_\ell)})}{q(\boldsymbol{\theta}_\ell, \mathbf{u}^{(n, \boldsymbol{\theta}_\ell)})}$$

**end**

**compute**

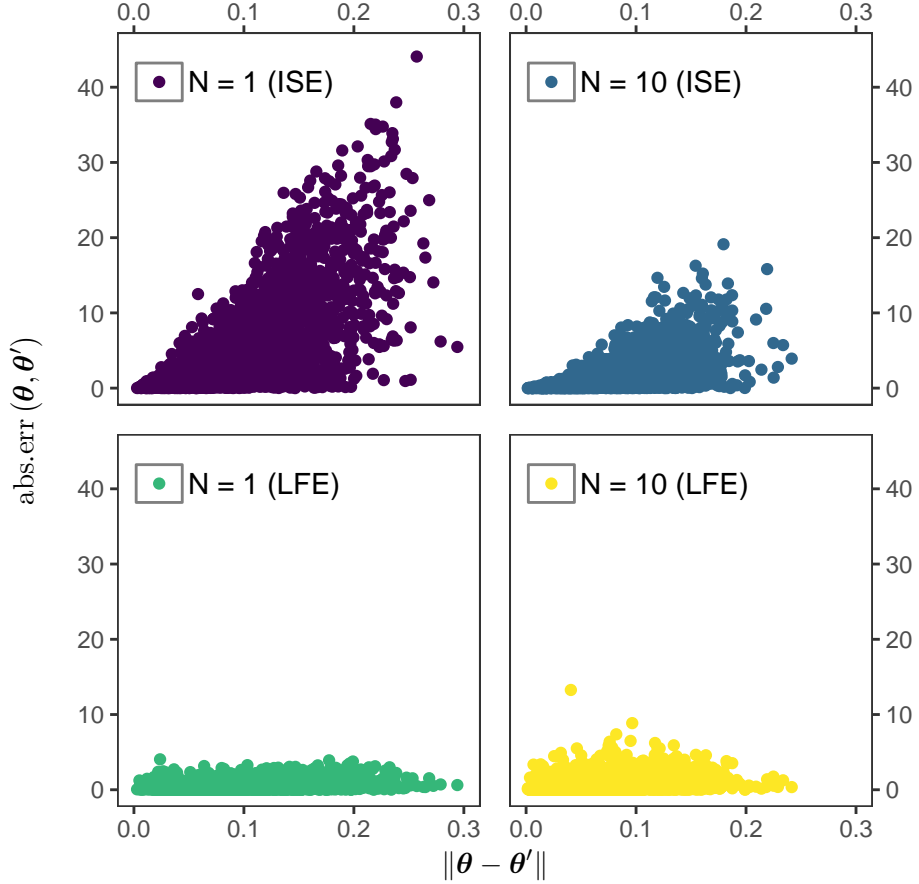
$$\widehat{\rho}(\boldsymbol{\theta}_0, \mathbf{r}, \boldsymbol{\theta}_L, \mathbf{r}_L) = 1 \wedge \frac{q(\boldsymbol{\theta}_L, \mathbf{x}) \mathcal{N}(\mathbf{r}_L \mid \mathbf{0}_d, \mathbf{M}) p(\boldsymbol{\theta}_0)}{q(\boldsymbol{\theta}_0, \mathbf{x}) \mathcal{N}(\mathbf{r} \mid \mathbf{0}_d, \mathbf{M}) p(\boldsymbol{\theta}_0)} \prod_{\ell=1}^L \frac{Z(\widehat{\boldsymbol{\theta}_{\ell-1}})}{Z(\boldsymbol{\theta}_\ell)}$$

**set**  $(\boldsymbol{\theta}', \mathbf{r}') = (\boldsymbol{\theta}_L, -\mathbf{r}_L)$  with probability  $\widehat{\rho}(\boldsymbol{\theta}_0, \mathbf{r}, \boldsymbol{\theta}_L, \mathbf{r}_L)$ ;

**Otherwise** **set**  $(\boldsymbol{\theta}', \mathbf{r}') = (\boldsymbol{\theta}_0, -\mathbf{r})$ ;

importance sampling estimate strongly decreases when the  $L^2$ -norm  $\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$  increases. Such a peculiarity has no impact on Metropolis-Hastings methods based on random walks such as the exchange algorithm. Indeed, the variance of the proposal is picked to ensure an acceptance rate high enough. This results in proposing move whose norm  $\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$  then remains close to zero where the absolute error is moderate. However the deterministic proposals of HMC algorithms are designed to produce large transitions. We thus observe a deterioration of the estimator as shown in Figure 1 when  $\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$  increases and therefore the importance sampling estimator (3.8) cannot be advocated in the present situation.

To overcome such an issue, we take advantage of all the auxiliary draws with respect to intermediate points  $\{\boldsymbol{\theta} = \boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}' = \boldsymbol{\theta}_L\}$ . The ratio of the normalising constant is replaced by an unbiased importance sampling estimate, referred to as the leapfrog



**Figure 1:** Absolute error of the log-ratio estimates ( $y$ -axis) with respect to the distance between the current and the proposed states ( $x$ -axis). The first row refers to importance sampling estimates (ISE) which only use the current and the proposed states. The second row refers to leapfrog estimates (LFE) which make a use of all parameter values involved in the integrator scheme. In terms of absolute error, LFE provides more accurate ratio estimates regardless the distance between the current and the proposed states.

estimator (LFE), based on the product of ratios taken at two consecutive points of the integration path, namely

$$\frac{\widehat{Z}(\boldsymbol{\theta})}{\widehat{Z}(\boldsymbol{\theta}')} = \prod_{\ell=0}^{L-1} \frac{\widehat{Z}(\boldsymbol{\theta}_\ell)}{\widehat{Z}(\boldsymbol{\theta}_{\ell+1})} = \prod_{\ell=0}^{L-1} \frac{1}{N} \sum_{n=1}^N \frac{q(\boldsymbol{\theta}_\ell, \mathbf{u}^{(n, \boldsymbol{\theta}_{\ell+1})})}{q(\boldsymbol{\theta}_{\ell+1}, \mathbf{u}^{(n, \boldsymbol{\theta}_{\ell+1})})}, \quad (3.9)$$

**Table 1:** Mean squared error of the log ratio estimator

	$N = 1$	$N = 10$
Importance sampling estimator (ISE)	55.5	6.08
Leapfrog estimator (LFE)	0.59	1.13

where  $\mathbf{u}^{(\cdot|\theta_\ell)}$  is sampled from  $f(\cdot|\theta_\ell)$ . As it simply reuses the draws involved in gradient estimates (3.5), this new estimator comes at no extra simulation cost but requires one to evaluate the function  $q$   $2N \times L$  times. The solution we advocate is slightly different than the exchange algorithm with bridging (Murray et al., 2006) which relies on an annealed importance sampling technique (Neal, 2001). Here we do not use a sequence of distributions bridging between the original proposal and target distributions for fixed  $(\theta, \theta')$ .

Overall, we observe in Figure 1 that the leapfrog estimators (LFE, second row) are much more accurate than importance sampling estimators (ISE, first row) even for a small number of auxiliary draws and thus a poor estimation of the gradient (see Section 3.1). Mean squared errors presented in Table 1 support this conclusion.

### 3.4 Tuning the noisy HMC algorithm

Tuning HMC often turns out to be a delicate task. In practice, sampling from a density  $\pi$  using HMC is highly sensitive to user-specified parameters: the step size  $\epsilon$ , the number of leapfrog steps  $L$  and the covariance matrix  $\mathbf{M}$ . We refer the reader to Neal (2011) and Hoffman and Gelman (2014) for a more comprehensive discussion. While the tuning issue is not really the focus of this paper, hereafter we present a tuning strategy which we have followed for doubly-intractable problems and which could well be improved upon in various ways.

For standard HMC, too large a step size  $\epsilon$  results in an inaccurate approximated flow  $F_\epsilon$  and subsequently a high rejection rate. Conversely, if  $\epsilon$  is too small, the leapfrog integrator will be precise but will require a significant computational cost to simulate a trajectory. These general considerations should be put in perspective with the quality of the Monte Carlo estimate (3.5) used in place of the gradient of the log-posterior. Following a poor gradient estimate for too long will affect the dynamic and significantly cut down the acceptance rate. Increasing  $N$  will improve the precision in the gradient estimation allowing larger steps but involves increased CPU time sampling from the model. Therefore, a trade-off has also to be found with the number of auxiliary draws  $N$ . Following the guidelines of Hoffman and Gelman (2014) and also Stan (Carpenter et al., 2016), the parameter  $\epsilon$  was chosen using the dual averaging method, that is  $\epsilon$  is adaptively tuned during a burn-in period so that for a given integration time  $t = \epsilon L$  the average acceptance probability of HMC reaches an optimal value  $\delta$ . Beskos et al. (2013) showed that for a given integration time  $t$ , the optimal value of  $\epsilon$  produces a chain with probability  $\delta = 0.65$ , approximately. It is not clear whether it is an appropriate probability to target for the noisy HMC algorithm since, for a fixed  $N$ , the transition kernel of the standard HMC method has been replaced by an approximate version. However since the estimators (3.5) and (3.9) converge, almost surely, as  $N$  goes to infinity, the latter is then the limiting probability associated to the standard kernel and can be used as a rough but reasonable target.

Furthermore  $L$ , or equivalently  $t$ , needs to be large enough to avoid a random walk behaviour and thereby the slow mixing issue which one would like to prevent in the first place. Too long a trajectory is counter-productive since the dynamics retraces its steps bringing the proposed value  $\boldsymbol{\theta}'$  back to a neighbourhood of the current value  $\boldsymbol{\theta}$ . The NO-U-Turn Sampler (NUTS, [Hoffman and Gelman, 2014](#)) was specifically introduced to avoid such a scenario by producing automatically “optimal” trajectory lengths. Whilst NUTS fulfils detailed balance and reversibility, this is marred by a major point for doubly-intractable distribution. Indeed, NUTS introduces a slice variable whose conditional distribution given  $(\boldsymbol{\theta}, \mathbf{r})$  is uniform on  $[0; f(\mathbf{x} | \boldsymbol{\theta})p(\boldsymbol{\theta})\mathcal{N}(\mathbf{r} | \mathbf{0}_d, \mathbf{M})]$  which requires evaluating the intractable normalising constant  $Z(\boldsymbol{\theta})$ . NUTS hence cannot be used to tune  $L$  and we have to rely on our personal expertise using a fixed integration time  $t$  instead based on the following heuristic. Given a parameter  $(\boldsymbol{\theta}, \mathbf{r})$  and a step size  $\epsilon$ , the proposed value  $\boldsymbol{\theta}'$  at the end of Algorithm 1 is written as

$$\boldsymbol{\theta}' = \boldsymbol{\theta} + \frac{\epsilon^2 L \mathbf{M}^{-1}}{2} \nabla_{\boldsymbol{\theta}} \log \pi(\boldsymbol{\theta} | \mathbf{x}) + \underbrace{\epsilon^2 \mathbf{M}^{-1} \sum_{\ell=1}^{L-1} (L - \ell) \nabla_{\boldsymbol{\theta}} \log \pi(\boldsymbol{\theta}^{(\ell)} | \mathbf{x})}_{:=\Delta(\boldsymbol{\theta}, \mathbf{r})} + \epsilon L \mathbf{M}^{-1} \mathbf{r}.$$

Given  $\boldsymbol{\theta}$ , the latter can be seen as a non-linear transformation of the resampled auxiliary variable  $\mathbf{r} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{M})$  whose distribution is typically intractable when  $L > 1$  due to the term  $\Delta(\boldsymbol{\theta}, \mathbf{r})$ . We remark that when  $L = 1$ , the  $\Delta(\boldsymbol{\theta}, \mathbf{r})$  term vanishes and we get the MALA proposal ([Roberts and Tweedie, 1996](#)). In order to calibrate the integration time, we neglect the randomness of the  $\Delta(\boldsymbol{\theta}, \mathbf{r})$  term and approximate the HMC proposal by

$$\boldsymbol{\theta}' | \boldsymbol{\theta} \sim \mathcal{N}\left(\boldsymbol{\theta} + \frac{\epsilon^2 L \mathbf{M}^{-1}}{2} \nabla_{\boldsymbol{\theta}} \log \pi(\boldsymbol{\theta} | \mathbf{x}), \epsilon^2 L^2 \mathbf{M}^{-1}\right). \quad (3.10)$$

The gold standard for dealing with doubly-intractable distributions is the exchange algorithm. For this particular random-walk Metropolis algorithm, the proposal distribution is of the form

$$\boldsymbol{\theta}' | \boldsymbol{\theta} \sim \mathcal{N}\left(\boldsymbol{\theta}, \sigma^2 \mathbf{M}^{-1}\right).$$

For such a class of algorithms, the optimal scaling ([Roberts and Rosenthal, 2001](#)) is obtained for

$$\sigma = \frac{2.38}{\sqrt{d}}. \quad (3.11)$$

In what follows, we set the integration  $t$  to the above (3.11). Such a choice ensure that the scaling of the covariance matrix in HMC proposal (3.10) is exactly the same than for the exchange algorithm. This settings is obviously arguable since it leads to a sub-optimal choice for the step-size  $\epsilon$ . Indeed, the optimal tuning for HMC in high dimension is a step-size scaled as  $\epsilon = \ell \times d^{-1/4}$  for some positive constant  $\ell$  ([Beskos et al., 2013](#)). However the current choice leads a step-size scaled as  $\epsilon = \ell \times d^{-1/2}$  and can obviously be improved. However being in small dimension, we observed a good behaviour for such a setting.

Finally the last tunable parameter is the mass matrix  $\mathbf{M}$ . A mass matrix well suited to the covariance  $\Sigma$  of the posterior, namely  $\mathbf{M} = \Sigma^{-1}$ , can enhance both the speed

and the mixing of HMC. For doubly-intractable distributions it is possible to estimate  $\Sigma$  in a burn-in phase by using a stochastic approximation algorithm which makes use of the gradient of the log posterior distribution defined in (3.5). The method works by estimating the mode  $\boldsymbol{\theta}^*$  of the posterior (3.1) using the gradient (3.5) within a Robbins Monro algorithm (Robbins and Monro, 1951) or Ruppert-Polyak averaging (Ruppert, 1988, Polyak and Juditsky, 1992). Once  $\boldsymbol{\theta}^*$  has been found, consequently to Equation (3.4) the Hessian of the log posterior at  $\boldsymbol{\theta}^*$  can be approximated using the sample covariance of auxiliary draws at the mode

$$\nabla_{\boldsymbol{\theta}}^2 A(\boldsymbol{\theta}^*, \mathbf{x}) + Cov \left[ \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}^*, \mathbf{u}^{(1, \boldsymbol{\theta}^*)}), \dots, \nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}^*, \mathbf{u}^{(N, \boldsymbol{\theta}^*)}) \right] + \nabla_{\boldsymbol{\theta}}^2 \log p(\boldsymbol{\theta}^*) := \widehat{\Sigma}^{-1}. \quad (3.12)$$

where  $\{\mathbf{u}^{(1, \boldsymbol{\theta}^*)}, \dots, \mathbf{u}^{(N, \boldsymbol{\theta}^*)}\}$  are samples from  $f(\cdot \mid \boldsymbol{\theta}^*)$ . This approximation to the Hessian can hence be used as an estimate of  $\mathbf{M}$ . We consider this estimation of  $\mathbf{M}$  for Potts and the exponential random graph model numerical study below.

For a matrix  $\mathbf{M}$  that depends on  $\boldsymbol{\theta}$ , Girolami and Calderhead (2011) provide a fully automated scheme based on the Riemann geometry of the parameter space to adapt  $\mathbf{M}$  along the run. However,  $\mathbf{M}$  then becomes parameter dependent and the leapfrog integrator needs to be replaced by a more sophisticated integration scheme. The latter is somewhat delicate to implement for doubly-intractable distributions and we will not consider it for our numerical study in Section 4.

## 4 Numerical results

### 4.1 Toy example: Potts model study

Consider an undirected graph  $\mathcal{G}$  inducing a topology on a finite set of sites  $\mathcal{S} = \{1, \dots, n\}$ : by definition, sites  $i$  and  $j$  are neighbours (denoted  $i \sim j$ ) if and only if  $i$  and  $j$  are linked by an edge in  $\mathcal{G}$ . One then calls clique a subset of  $\mathcal{S}$  where all elements are mutual neighbours. A discrete Markov random field with respect to undirected graph  $\mathcal{G}$  is a random process  $\mathbf{X} = (X_1, \dots, X_n)$  indexed by  $\mathcal{S}$ , and taking values in  $\mathcal{X} \subset \mathbb{Z}^n$ , whose conditional distribution of  $X_i$ ,  $i \in \mathcal{S}$ , depends only upon its neighbours in  $\mathcal{G}$ . The Hammersley-Clifford theorem states that if the distribution of a Markov random field with respect to a graph  $\mathcal{G}$  is positive for all configurations  $\mathbf{x}$  then it admits a Gibbs representation for the same topology (see *e.g.* Grimmett (1973), Besag (1974) and for a historical perspective Clifford (1990)), namely a density function  $f(\cdot \mid \boldsymbol{\theta})$  on  $\mathcal{X}$  parametrised by  $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^d$  and given with respect to the counting measure by

$$f(\mathbf{x} \mid \boldsymbol{\theta}) = \frac{1}{Z(\boldsymbol{\theta})} \exp \{A_{\mathcal{G}}(\boldsymbol{\theta}, \mathbf{x})\},$$

where  $A_{\mathcal{G}}$  denotes the potential function which can be written as a sum over the set  $\mathcal{C}$  of all cliques of the graph, namely  $A_{\mathcal{G}}(\boldsymbol{\theta}, \mathbf{x}) = \sum_{c \in \mathcal{C}} A_c(\boldsymbol{\theta}, \mathbf{x})$  for all configurations  $\mathbf{x} \in \mathcal{X}$ . The inherent difficulty of all these models arises from the intractable normalising

constant, sometimes called the partition function, defined by

$$Z(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in \mathcal{X}} \exp \{A_{\mathcal{G}}(\boldsymbol{\theta}, \mathbf{x})\}.$$

The latter is a summation over the numerous possible realisations of the random field  $\mathbf{X}$ , which is of combinatorial complexity and cannot be computed directly (except for small grids and states space  $\mathcal{X} = \{0, \dots, K-1\}^n$  with small number of states  $K$ ).

In what follows, we focus on a particular pairwise model representative of the general level of difficulty, namely the Potts model (Potts, 1952) originally used in statistical mechanics to model interacting spins. The function  $A_{\mathcal{G}}$  writes as a sum over the cliques of size 1 (corresponding to the nodes of  $\mathcal{G}$ ) and cliques of size 2 (corresponding to the edges of  $\mathcal{G}$ ). Denote  $\boldsymbol{\theta} = (\alpha_0, \dots, \alpha_{K-1}, \beta)$ ,

$$A_{\mathcal{G}}(\boldsymbol{\theta}, \mathbf{x}) = \sum_{i=1}^n \sum_{k=0}^{K-1} \alpha_k \mathbf{1}\{x_i = k\} + \beta \sum_{i \sim j} \mathbf{1}\{x_i = x_j\},$$

where the above sum  $\sum_{i \sim j}$  ranges the set of edges of the graph  $\mathcal{G}$ . The parameter  $\alpha$  can be interpreted as an external field while the parameter  $\beta$  can be interpreted as the inverse of a temperature adjusting the level of dependency between adjacent sites. In the absence of an external field, when the temperature drops below a fixed threshold, called phase transition, the model exhibits strong dependence between neighbors and values  $x_i$  of a typical realisation of the field are almost all equal. Note that a potential function on nodes can be defined up to an additive constant. To ensure that potential functions on nodes are uniquely determined, one usually imposes the constraint  $\sum_{k=0}^{K-1} \alpha_k = 0$ . The dimension of the parameter space  $\Theta$  is then  $K$ .

In this example we focus on a digital 2-state Potts model defined on a  $16 \times 16$  regular lattice with a first order neighbourhood system, see Figure 2. This example was chosen since it is one for which we can estimate very accurately the underlying posterior distribution of model parameters and therefore we can use this as a pedagogical example to compare our noisy HMC algorithm to the corresponding noisy exchange algorithms. We also note that this is not a particularly challenging example because of the size of the grid and the low number of states. As such, it not one which best exemplifies the performance of the noisy HMC algorithm.

**Ground truth** The R-package GiRaF (Stoehr et al., 2016) allows one to compute exactly the normalising constant  $Z(\boldsymbol{\theta})$  of a Potts model defined on a rectangular  $h \times w$  lattice (Friel and Rue, 2007). The algorithm to evaluate  $Z(\boldsymbol{\theta})$  is exponential in the number of rows  $h$  and linear in the number of columns  $w$ . As such it can handle models defined on a lattice up to  $h = 25$  for  $K = 2$ . For such a lattice, we can then compute ground truth quantities against which we can compare the output from the various algorithms that we consider in this paper. In particular using adaptive cubature algorithms from the R-package cubature (Narasimhan and Johnson, 2017) we



**Figure 2:** (a) First order neighbourhood graphs  $\mathcal{G}$ . (b) a digital 2-states Potts model defined on a  $n = 16 \times 16$  regular lattice with a first order neighbourhood system.

can compute the posterior mean

$$\bar{\boldsymbol{\theta}} = \mathbf{E}_{\pi} \{\boldsymbol{\theta}\} = \left[ \int_{\Theta} \theta_i \pi(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta} \right]_i,$$

and the Kullback-Leibler divergence between the empirical distribution  $Q$  and the stationary distribution  $\pi$ , namely

$$\text{KL}(Q||\pi) = \int_{\Theta} Q(\boldsymbol{\theta}) \log \frac{Q(\boldsymbol{\theta})}{\pi(\boldsymbol{\theta} | \mathbf{x})} d\boldsymbol{\theta}.$$

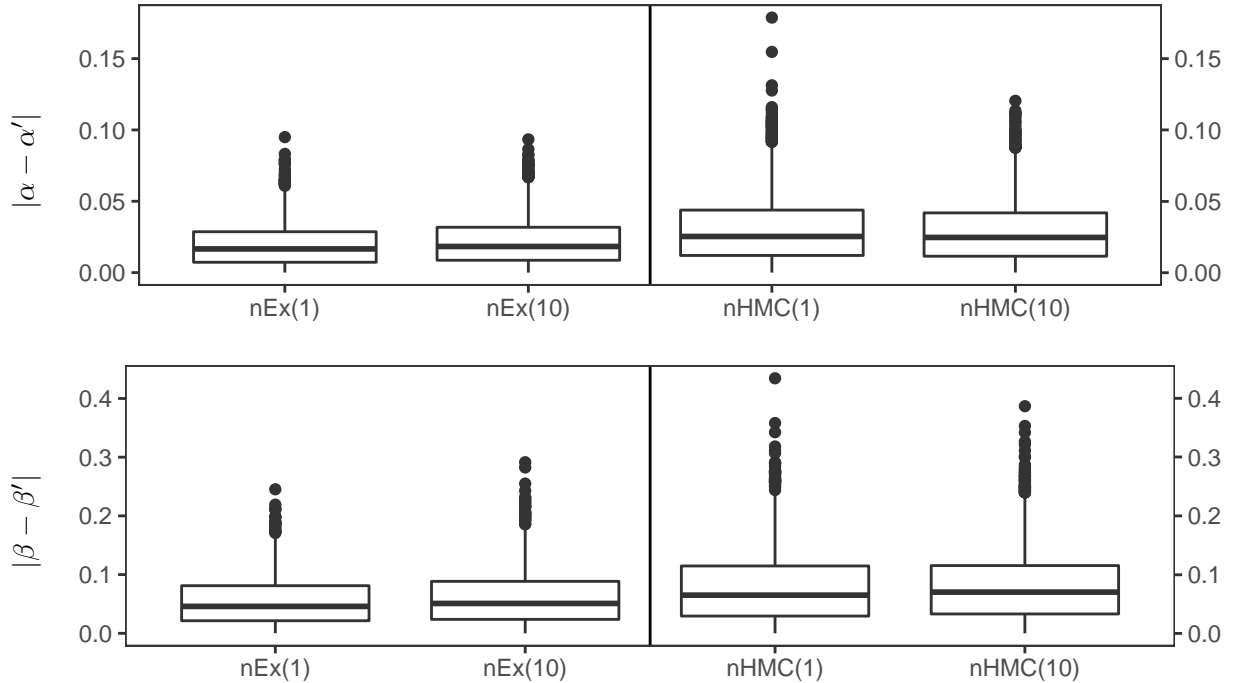
To estimate the latter, we divided the parameter space  $\Theta$  into a set of bins of size  $0.01 \times 0.01$ . Then, we solely computed integrals over the set of non-empty bins, integrals being set to zero otherwise.

**MAP and Hessian estimates** The maximum *a posteriori* (MAP) was estimated using the Ruppert-Polyak averaging method. In this scheme, at each iteration the gradient was estimated using the identity (3.5) with  $N = 10$  draws from the likelihood. The algorithm stops when  $\|\boldsymbol{\theta}_n - \boldsymbol{\theta}_{n+1}\|$  drops below a  $1e^{-3}$  threshold. Once the MAP  $\boldsymbol{\theta}^*$  is estimated, we estimate the Hessian  $\nabla_{\boldsymbol{\theta}}^2 \log \pi(\boldsymbol{\theta}^* | \mathbf{x}) = \Sigma^{-1}$  using identity (3.12) with  $N = 500$  draws from the likelihood.

**First experiment: nHMC(N)** We ran the noisy HMC algorithm for various numbers of draws, namely  $N = 1$  and  $N = 10$ , in order to compute the gradient estimate (3.5) and the leapfrog estimator (3.9). The covariance matrix of the auxiliary variable distribution is set to  $\mathbf{M} = \widehat{\Sigma}^{-1}$ . For each value of  $N$ , the step size  $\epsilon$  and the number of leapfrog steps  $L$  were tuned as described in Section 3.4 with  $\delta = 0.65$ . The burn-in chain was of 500 iterations to tune the parameters for each setting.

**Second experiment: nEx(N)** The gold standard for conducting doubly-intractable Bayesian inference is the exchange algorithm (Murray et al., 2006). We ran the noisy





**Figure 3:** Potts model: boxplot of absolute accepted moves lengths for each component of  $\theta = (\alpha, \beta)$  in the exchange (nEx(1)), the noisy exchange (nEx(10)) and the noisy HMC (nHMC(1), nHMC(10)). The accepted moves of HMC algorithms are greater than the two random walk exchange algorithms across both parameter directions.

exchange algorithm (Alquier et al., 2016) for various number of auxiliary draws, namely  $N = 1$  and  $N = 10$ . The proposal distribution is set to be a 2-dimensional normal distribution with mean  $\theta$  and covariance matrix  $2.38^2 d^{-1} \hat{\Sigma}$  in order to target the optimal acceptance probability for Metropolis-Hastings algorithm, namely  $\delta = 0.234$  (Roberts and Rosenthal, 2001), as already mentioned in Section 3.4.

For the various experiments aforementioned, we considered a pseudo-observation exactly drawn from a Potts model using the R-package GiRaF with parameter  $\theta = (0., 0.5)$  and a uniform prior on  $\Theta = [-0.5; 0.5] \times [0; 1]$ . Once the parameters were tuned, we ran 20 chains of size 4,500 per experiment. Figure 3 shows the typical behaviour of the accepted moves length for both the noisy exchange, nEx( $N$ ), and noisy HMC, nHMC( $N$ ) algorithms, when  $N = 1$  or  $N = 10$  auxiliary draws. First, we can observe that the noisy version of those algorithms still benefits from using a well chosen mass matrix  $\mathbf{M}$ . Indeed, the chain is then able to move more in the less constrained direction, namely  $\beta$ . On the other hand, both algorithms produce a proposal with the same covariance terms but the HMC chain manages to move more across the parameter space despite a sub-optimal choice of  $\epsilon$ . We surmise that the performance could be further improved with a better tuning for  $\epsilon$ .

Table 2 summarises different outputs average obtained over these chains. As a first

step, we note that the dual averaging scheme with a Monte Carlo estimate (3.5) of the gradient still ensures that the average Metropolis acceptance probability  $\rho(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is close to its required value. As regards the value  $\epsilon$  and  $L$ , we observe a trade-off with the quality of the gradient as already mention in Section 3.4. In order to keep proposing transitions which are accepted with high probability, a poor gradient estimate leads to numerous small steps, therefore avoiding to follow an incorrect gradient for too long, and conversely,  $\epsilon$  becomes larger and  $L$  smaller when we put more computation effort to get a more accurate gradient estimate.

The major motivation for using HMC was to have better mixing. Table 2 gives Effective Sample Size (ESS) with respect to the marginal chain in  $\alpha$  (denoted  $\text{ESS}(\alpha)$ ) and the marginal chain in  $\beta$  (denoted  $\text{ESS}(\beta)$ ) so as to stress that difference of behaviour between noisy algorithms. In terms of a per iteration basis, noisy HMC overall mixes better than for various exchange algorithms, as one would expect. However it is significantly more expensive in terms of computation time. Indeed, the cost of the different algorithms is largely determined by the number of auxiliary draws from the model. HMC performs  $N \times (L + 1)$  draws per iteration whereas the noisy exchange does  $N$  draws. On a per unit of computational time basis, the exchange algorithm is then typically superior to the noisy HMC algorithms, at least in terms of ESS. We further remark here that the noisy version proposed by [Alquier et al. \(2016\)](#) comes at a loss in practice as the overall ESS does not increase significantly enough against the extra computational cost.

On a per iteration basis, noisy HMC present better results in terms of the mean squared error of the posterior mean, namely  $\text{MSE}(\bar{\boldsymbol{\theta}})$ , and Kullback-Leibler divergence. One could argue that such results might not hold on a per unit of computational time basis as we would have a chain 10 times longer for the exchange algorithm and therefore expect to have better result in terms of  $\text{MSE}(\bar{\boldsymbol{\theta}})$  and Kullback-Leibler divergence, though it is not completely obvious whether this would be the case.

**Table 2:** Potts model: noisy HMC, exchange and noisy exchange algorithms outputs average (standard deviation) over 20 chains of size 4,500 for various number of auxiliary draws ( $N = 1$  and  $N = 10$ ).

	nHMC(1)	nHMC(10)	nEx(1)	nEx(10)
$\rho(\boldsymbol{\theta}, \boldsymbol{\theta}')$ (%)	62.8 (1e-2)	63.5 (8.2e-3)	22.5 (7.5e-3)	28.9 (7.7e-3)
$\epsilon$	0.13	0.80		NA
L	12	2		NA
Running time (s)	1456 (110)	3268 (34.4)	110.7 (4.4)	1133 (34)
ESS( $\alpha$ )	1604 (103)	1891 (169)	244.4 (34.7)	394.8 (40.8)
ESS( $\alpha$ ) per sec.	1.12 (9.1e-2)	0.57 (5.3e-2)	2.21 (0.34)	0.35 (3.8e-2)
ESS( $\beta$ )	1573 (181)	1904 (186)	241.4 (23.6)	396.1 (41.4)
ESS( $\beta$ ) per sec.	1.09 (0.13)	0.59 (5.7e-2)	2.18 (0.24)	0.35 (3.8e-2)
MSE( $\bar{\boldsymbol{\theta}}$ ) ( $\times 1e-5$ )	0.34	1.68	3.80	1.52
KL-divergence	0.23 (1.8e-2)	0.18 (1.3e-2)	0.43 (5.1e-2)	0.33 (2.5e-2)

## 4.2 Exponential random graph model study

Exponential random graph models are a family of network models that generalise Markov random graphs. A graph is a collection of  $n$  nodes connected by edges. The edges are indicated by a  $n \times n$  adjacency matrix  $\mathbf{x}$  where

$$x_{ij} = \begin{cases} 1, & \text{if node } i \text{ and node } j \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases}$$

These edges may be directed or undirected, in the latter case  $x_{ij} = x_{ji}$  for all pairs  $(i, j)$ . An example of such a network is shown in Figure 4 depicting the friendship connections of 34 individuals (nodes) in a karate club (Zachary, 1977).

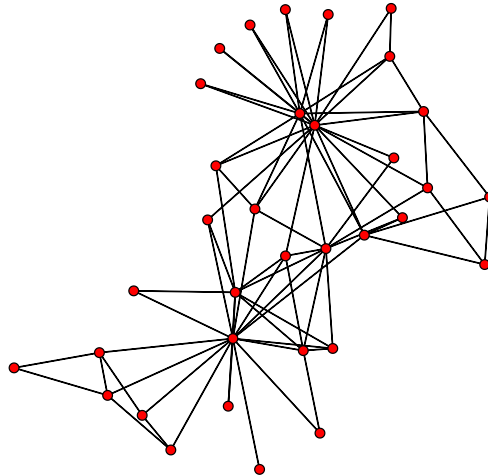
The energy function  $A(\boldsymbol{\theta}, \mathbf{x})$  for an exponential random graph model on a directed or an undirected graph  $\mathbf{x}$  with  $n$  nodes is,

$$A(\boldsymbol{\theta}, \mathbf{x}) = \sum_{i=1}^d \theta_i s_i(\mathbf{x}), \quad (4.1)$$

where the terms  $s_i(\mathbf{x})$  are typically sufficient statistics. These statistics capture local structure of the network  $\mathbf{x}$ , for example, the count of edges or triangles in the network. The normalising constant for the exponential random graph models involves a sum over all  $2^{\binom{n}{2}}$  realisable graphs  $\mathbf{x} \in \mathcal{X}$  in the undirected case and  $2^{n(n-1)}$  in the directed case,

$$Z(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in \mathcal{X}} \exp \left\{ \sum_{i=1}^d \theta_i s_i(\mathbf{x}) \right\}$$

and is intractable for all but small graphs.



**Figure 4:** Zachary karate club: A social network of 34 individuals in a karate club. The network has 78 edges and 528 2-stars.

In this numerical study we consider a two parameter ERGM where the sufficient statistics are  $s_1(\mathbf{x}) = \sum_{i < j} x_{ij}$ , counting the number of undirected edges in  $\mathbf{x}$  and  $s_2(\mathbf{x}) = \sum_i \sum_{j < k} x_{ij} x_{ik}$ , counting the number of so-called 2-stars in the network. The sufficient statistics are assigned parameters  $\boldsymbol{\theta} = (\theta_1, \theta_2)$  with model energy function,

$$A(\boldsymbol{\theta}, \mathbf{x}) = \theta_1 s_1(\mathbf{x}) + \theta_2 s_2(\mathbf{x}). \quad (4.2)$$

The posterior distribution for the parameters of an exponential random graph model as defined with potential (4.1) is doubly-intractable, therefore it is not possible to obtain a ground truth for the posterior distribution of the parameters in order to compare with HMC. This is in contrast to the Potts model study in which a ground truth is available for small lattices using the R-package `GiRaF`. Instead using the R-package `Bergm` (Caimo and Friel, 2014), an exchange algorithm was run for a long period of time in order to sample from the posterior (3.1). This long run consisted of 100,000 iterations for the parameter  $\boldsymbol{\theta}$  with a per iteration auxiliary burn in of 1,000,000 to obtain the estimate the single importance sampling estimate (3.7).

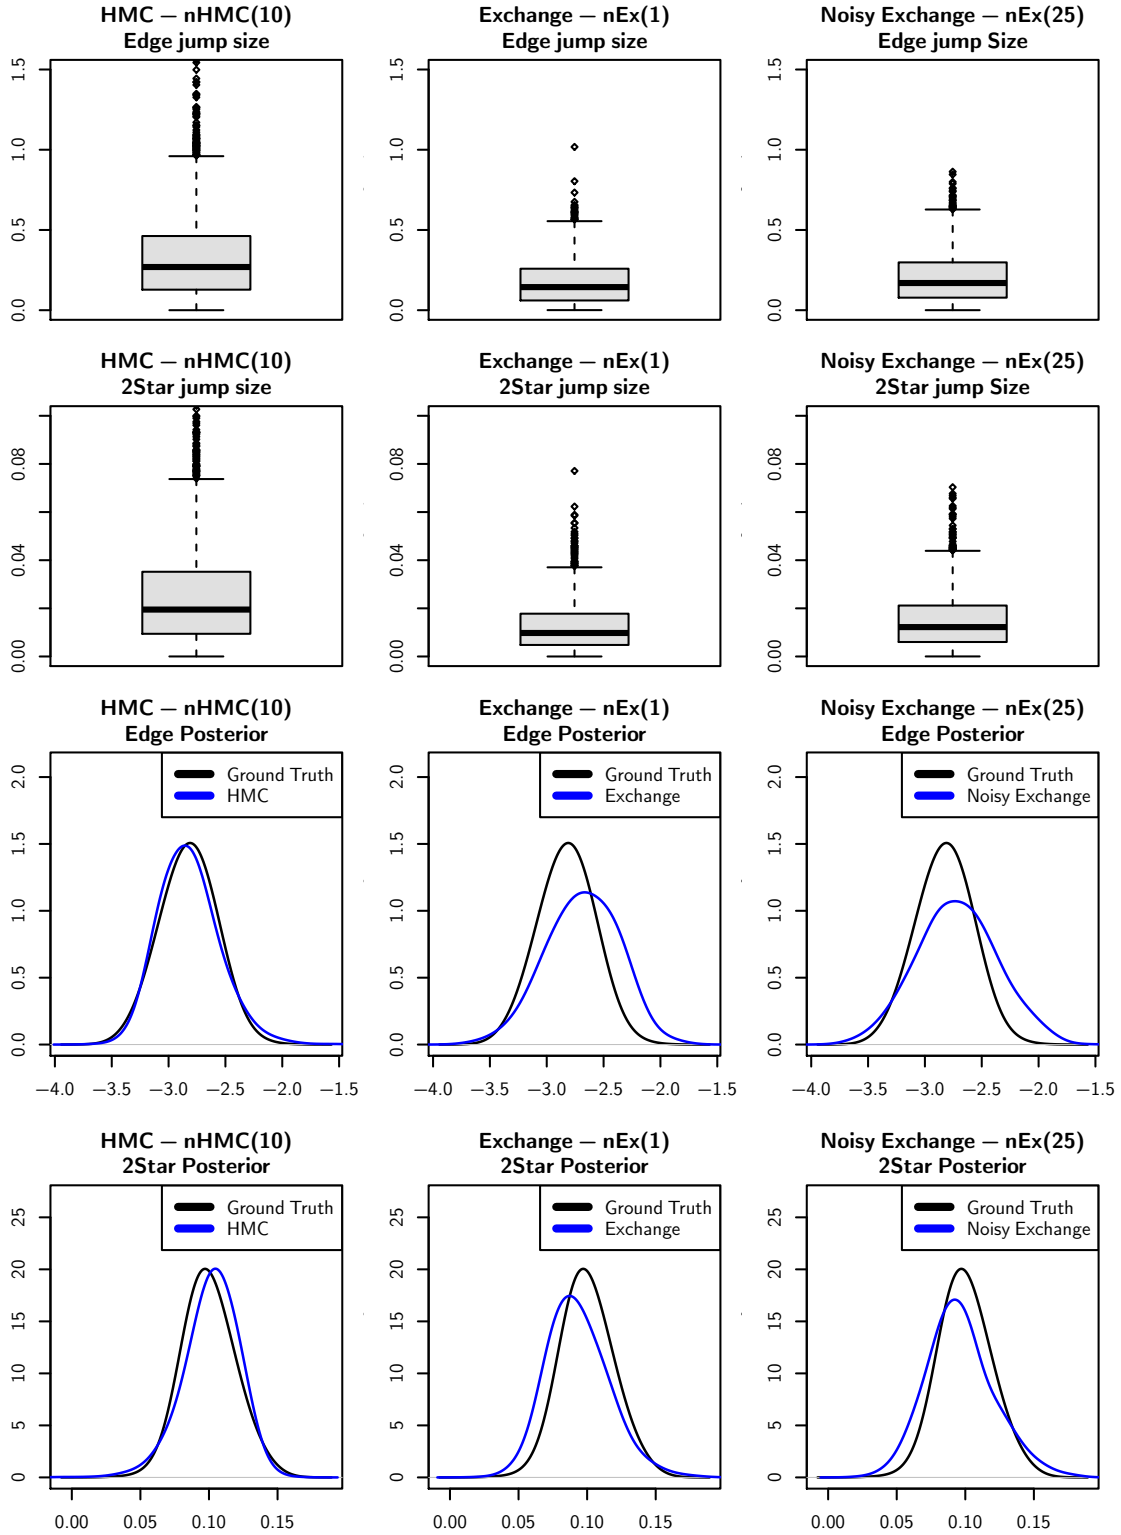
**First experiment: nHMC(N)** We ran the noisy HMC algorithm (Algorithm 2) for  $N = 10$  draws in order to compute the gradient estimate (3.5) and the leapfrog estimator (3.9). To tune the algorithm the  $\mathbf{M}$  matrix was set equal to an approximation of the Hessian of the log posterior (3.12). This approximation was computed to locating the posterior mode  $\boldsymbol{\theta}^*$  by iterating the following Robbins-Munro scheme for  $i = 1, \dots, 200$ ,

$$\boldsymbol{\theta}^{(i+1)} = \boldsymbol{\theta}^{(i)} + \frac{\alpha}{i} \widehat{\nabla}_{\mathbf{u}(\boldsymbol{\theta}, N=10)} \log \pi(\boldsymbol{\theta}^{(i)} | \mathbf{x}),$$

with  $\alpha = 1$ . The approximate location of the mode is taken as  $\boldsymbol{\theta}^* = \boldsymbol{\theta}^{(201)}$ . Then 500 draws are sampled from  $f(\cdot | \boldsymbol{\theta}^*)$  to compute the Hessian estimate  $\widehat{\Sigma}^{-1}$  using (3.12). The mass matrix  $\mathbf{M}$  was then set equal to  $\widehat{\Sigma}^{-1}$ .

**Second experiment: nEx(N)** The exchange algorithm (Murray et al., 2006) was run for 5,000 iterations of the parameter  $\boldsymbol{\theta}$ . The noisy exchange (Alquier et al., 2016) was run for 5,000 iterations of the parameter  $\boldsymbol{\theta}$  using  $N = 25$  auxiliary draws at each parameter value. The proposal distribution for both the exchange and noisy was taken as a normal distribution with mean at the current  $\boldsymbol{\theta}$  and covariance matrix  $2.38^2 d^{-1} \mathbf{M}^{-1}$ , where  $\mathbf{M} = \widehat{\Sigma}^{-1}$  is approximated using (3.12) as with the noisy HMC study.

The different algorithms were replicated 20 times using a different starting point and random seed for each experiment. The results of the various algorithms run on the Zachary karate club for the model defined by (4.2) are shown in Table 3 and Figure 5. The noisy HMC provides a noticeable increase in ESS per unit time with respect to the exchange and noisy exchange algorithms. The mean squared error (MSE) to the ground truth value of  $\theta$  is also significantly lower for HMC and this is visible in the marginal density estimates given in Figure 5, where the HMC algorithm matches closely to each marginal posterior distribution from the ground truth.



**Figure 5:** Zachary karate club: HMC, exchange and noisy exchange algorithm results. The jump sizes or absolute accepted moves lengths of the HMC algorithm are greater than the two random walk exchange algorithms across both parameter directions. The HMC algorithm shows good convergence to the ground truth results.

## 5 Conclusions

The paper yields a noisy version of the HMC algorithm for doubly-intractable Bayesian analysis. This noisy version of the HMC algorithm is based on Monte Carlo estimates for the gradient based leapfrog numerical integrator and a novel estimator for the Metropolis-Hastings proposal compatible with the integration scheme. Throughout the numerical studies described here, we have shown that it improves the statistical efficiency of the gold standard exchange algorithm which is only capable of exploring the parameter space slowly by making local moves. The computational cost of sampling from the model is the main bottleneck of computing gradients and importance sampling estimators required in HMC. Further research is required to assess the trade-off between the gradient quality, in terms of number of samples from the likelihood, and the efficiency of HMC. Noisy HMC opens vistas for further research.

**Theoretical support** On the first hand, the paper motivates an in-depth theoretical study of noisy HMC. Indeed, the transition kernel of the original HMC method has been replaced by an approximate kernel due to Monte Carlo estimates being used within each iteration. This breaks the ergodicity and the noisy HMC may not admit the target posterior distribution  $\pi$  as the invariant distribution. [Alquier et al. \(2016\)](#) have examined the ergodicity for noisy MCMC algorithms and provide bounds on the total variation distance between a Markov chain with the desired target distribution, and the Markov chain of a noisy MCMC algorithm. Their noisy MALA is a special case of our noisy HMC when the number of leapfrog steps is  $L = 1$ . While their results do not hold for larger  $L$ , one might expect similar conclusion for noisy HMC algorithms. Nevertheless the intractable HMC proposal forms to some extent a hurdle for the proof regarding noisy HMC. Another question that has also been opened by [Alquier et al. \(2016\)](#) is the

**Table 3:** Zachary karate club: HMC, exchange and noisy exchange algorithms results for the Zachary karate club network. The HMC was performed with 10 draws from the likelihood to estimate the gradients. Each chain is of size 5,000 and replicated 20 times to produce the results in this table.

	nHMC(10)	nEx(1)	nEx(25)
$\rho(\theta, \theta')$ (%)	61.7 (7e-3)	23.46 (6e-3)	29.44 (8e-3)
$\epsilon$	$3.453 \times 10^{-3}$	NA	NA
L	487	NA	NA
Running time (s)	1622 (92)	706 (111)	821 (149)
ESS ( $\theta_{\text{edge}}$ )	3132 (250)	260 (43)	296 (40)
ESS ( $\theta_{\text{edge}}$ ) per sec.	1.94 (0.22)	0.377 (0.08)	0.368 (0.06)
ESS ( $\theta_{2\text{star}}$ )	3379 (267)	285 (57)	323 (39)
ESS ( $\theta_{2\text{star}}$ ) per sec.	2.09 (0.23)	0.412 (0.10)	0.404 (0.07)
MSE (ground truth)	$1.968 \times 10^{-5}$	0.0249	0.00974

asymptotic variance of estimators from noisy HMC algorithm.

**Optimal tuning** The tuning strategy proposed in this paper is not optimal though the results for the noisy HMC algorithm we presented are already convincing. The overall performances could be improved by a better choice of the integration time and in general by introducing an automatic problem-specific strategy to pick  $\epsilon$  and  $L$ . A path to explore is to adapt the NUTS sampler in the context of doubly-intractable Bayesian inference. A follow-up to the theoretical support of noisy HMC would also be to provide an optimal setting.

**Latent process** Another specific instance of likelihood intractability occurs when the model relies on a latent process  $\mathbf{x}$  whose state space  $\mathcal{X}$  is of high dimension. Consider a likelihood expressed as a multidimensional integral,

$$\ell(\boldsymbol{\theta} \mid \mathbf{y}) = \int_{\mathcal{X}} \ell(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{x}) \mu(d\mathbf{x}),$$

and impossible to evaluate. Such problems arise frequently in applied statistics for example hidden Markov models. A fruitful avenue to explore is the development of noisy HMC methods for this class of statistical model. In the paper, we propose Monte Carlo estimates based on forward-simulations as surrogates for the gradient based integrator. Likewise, the score function of  $\ell(\boldsymbol{\theta} \mid \mathbf{y})$  can be estimated using following identity,

$$\nabla_{\boldsymbol{\theta}} \log \ell(\boldsymbol{\theta} \mid \mathbf{y}) = \mathbf{E}_{\boldsymbol{\theta}, \mathbf{y}} \{ \nabla_{\boldsymbol{\theta}} \log \ell(\boldsymbol{\theta}, \mathbf{X} \mid \mathbf{y}) \},$$

where the expectation is with respect to  $\ell(\mathbf{x} \mid \boldsymbol{\theta}, \mathbf{y})$  as described in [Friel et al. \(2016\)](#). Moreover, for hidden Markov models, the Metropolis-Hastings acceptance probability can still be evaluated using the leapfrog estimator we introduced in this work.

## References

- P. Alquier, N. Friel, R. Everitt, and A. Boland. Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, 26(1): 29–47, 2016.
- C. Andrieu and G. O. Roberts. The Pseudo-Marginal Approach for Efficient Monte Carlo Computations. *The Annals of Statistics*, 37(2):697–725, 2009.
- N. H. Augustin, M. A. Muggleston, and S. T. Buckland. The role of simulation in modelling spatially correlated data. *Environmetrics*, 9(2):175–196, 1998.
- M. A. Beaumont. Estimation of Population Growth or Decline in Genetically Monitored Populations. *Genetics*, 164(3):1139–1160, 2003.

- J. E. Besag. Spatial Interaction and the Statistical Analysis of Lattice Systems (with Discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2): 192–236, 1974.
- A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, and A. Stuart. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534, 2013.
- A. Caimo and N. Friel. Bergm: Bayesian exponential random graphs in R. *Journal of Statistical Software*, 61(2):1–25, 2014. URL <http://www.jstatsoft.org/v61/i02/>.
- B. Carpenter, D. Lee, M. A. Brubaker, A. Riddell, A. Gelman, B. Goodrich, J. Guo, M. Hoffman, M. Betancourt, and P. Li. Stan: A Probabilistic Programming Language, 2016.
- T. Chen, E. Fox, and C. Guestrin. Stochastic gradient hamiltonian monte carlo. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1683–1691, Beijing, China, 22–24 Jun 2014. PMLR.
- P. Clifford. Markov random fields in statistics. *Disorder in physical systems: A volume in honour of John M. Hammersley*, pages 19–32, 1990.
- S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid monte carlo. *Physics Letters B*, 195(2):216 – 222, 1987.
- O. François, S. Ancelet, and G. Guillot. Bayesian Clustering Using Hidden Markov Random Fields in Spatial Population Genetics. *Genetics*, 174(2):805–816, 2006.
- N. Friel. Bayesian Inference for Gibbs Random Fields Using Composite Likelihoods. In *Proceedings of the Winter Simulation Conference*, number 28 in WSC '12, pages 1–8. Winter Simulation Conference, 2012.
- N. Friel and H. Rue. Recursive computing and simulation-free inference for general factorizable models. *Biometrika*, 94(3):661–672, 2007.
- N. Friel, A. Mira, and C. J. Oates. Exploiting Multi-Core Architectures for Reduced-Variance Estimation with Intractable Likelihoods. *Bayesian Analysis*, 11(1):215–245, 2016.
- M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- P. J. Green and S. Richardson. Hidden Markov Models and Disease Mapping. *Journal of the American Statistical Association*, 97(460):1055–1070, 2002.
- G. R. Grimmett. A theorem about random fields. *Bulletin of the London Mathematical Society*, 5(1):81–84, 1973.
- W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.



- M. D. Hoffman and A. Gelman. The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- M. A. Hurn, O. K. Husby, and H. Rue. A Tutorial on Image Analysis. In *Spatial Statistics and Computational Methods*, volume 173 of *Lecture Notes in Statistics*, pages 87–141. Springer New York, 2003.
- T. S. Jaakkola and M. I. Jordan. Bayesian parameter estimation via variational methods. *Statistics and Computing*, 10(1):25–37, 2000.
- M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An Introduction to Variational Methods for Graphical Models. *Machine learning*, 37(2):183–233, 1999.
- B. G. Lindsay. Composite likelihood methods. *Contemporary Mathematics*, 80(1):221–39, 1988.
- A.-M. Lyne, M. Girolami, Y. Atchadé, H. Strathmann, and D. Simpson. On Russian Roulette Estimates for Bayesian Inference with Doubly-Intractable Likelihoods. *Statistical Science*, 30(4):443–467, 2015.
- J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate Bayesian Computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The journal of Chemical Physics*, 21(6):1087–1092, 1953.
- J. Møller, A. N. Pettitt, R. Reeves, and K. K. Berthelsen. An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika*, 93(2):451–458, 2006.
- I. Murray, Z. Ghahramani, and D. J. C. MacKay. MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, pages 359–366. AUAI Press, 2006.
- B. Narasimhan and S. G. Johnson. *cubature: Adaptive Multivariate Integration over Hypercubes*, 2017. URL <https://CRAN.R-project.org/package=cubature>. R package version 1.3-8.
- R. M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- R. M. Neal. *Handbook of Markov Chain Monte Carlo*, chapter 5, pages 113–162. CRC press, 2011.
- B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- R. B. Potts. Some generalized order-disorder transformations. In *Mathematical proceedings of the cambridge philosophical society*, volume 48, pages 106–109. Cambridge Univ Press, 1952.

- 
- H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- C. P. Robert and G. Casella. *Monte Carlo statistical methods*. Springer Texts in Statistics, second edition edition, 2004.
- G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16(4):351–367, 2001.
- G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations". *Bernoulli*, 2(4):341–363, 12 1996.
- G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph ( $p^*$ ) models for social networks. *Social networks*, 29(2):173–191, 2007.
- D. Ruppert. Efficient estimations from a slowly convergent Robbins-Monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.
- J. Stoehr and N. Friel. Calibration of conditional composite likelihood for bayesian inference on gibbs random fields. In *JMLR W&CP: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38, pages 921–929, 2015.
- J. Stoehr, P. Pudlo, and N. Friel. *GiRaF: A Toolbox for Gibbs Random Fields Analysis*, 2016. URL <https://CRAN.R-project.org/package=GiRaF>. R package version 1.0.
- L. Tierney. A note on Metropolis-Hastings kernels for general state spaces. *The Annals of Applied Probability*, 8(1):1–9, 02 1998.
- W. W. Zachary. An Information Flow Model for Conflict and Fission in Small Groups. *Journal of Anthropological Research*, 33(4):452–473, 1977.