

Measuring Sample Quality in Algorithms for Intractable Normalizing Function Problems

Bokgyeong Kang

*Department of Statistical Science
Duke University
Durham, NC 27708, USA*

BOKGYEONG.KANG@DUKE.EDU

John Hughes

*College of Health
Lehigh University
Bethlehem, PA 18015, USA*

JOH621@LEHIGH.EDU

Murali Haran

*Department of Statistics
The Pennsylvania State University
University Park, PA 16802, USA*

MURALI.HARAN@PSU.EDU

Editor: Alexandre Bouchard

Abstract

Models with intractable normalizing functions have numerous applications. Because the normalizing constants are functions of the parameters of interest, standard Markov chain Monte Carlo cannot be used for Bayesian inference for these models. A number of algorithms have been developed for such models. Some have the posterior distribution as their asymptotic distribution. Other “asymptotically inexact” algorithms do not possess this property. There is limited guidance for evaluating approximations based on these algorithms. Hence it is very hard to tune them. We propose two new diagnostics that address these problems for intractable normalizing function models. Our first diagnostic, inspired by the second Bartlett identity, is in principle broadly applicable to Monte Carlo approximations beyond the normalizing function problem. We develop an approximate version of this diagnostic that is applicable to intractable normalizing function problems. Our second diagnostic is a Monte Carlo approximation to a kernel Stein discrepancy-based diagnostic introduced by Gorham and Mackey (2017). We provide theoretical justification for our methods and apply them to several algorithms in challenging simulated and real data examples including an Ising model, an exponential random graph model, and a Conway–Maxwell–Poisson regression model, obtaining interesting insights about the algorithms in these contexts.

Keywords: Bartlett identity, doubly intractable distributions, kernel Stein discrepancy, Markov chain Monte Carlo, sample quality measure

1. Introduction

Models with intractable normalizing functions arise in many contexts, notably the Ising (Lenz, 1920; Ising, 1925) and autologistic models (see Besag, 1974; Hughes et al., 2011, for a review) for binary data on a lattice, exponential random graph models (see Robins et al., 2007; Hunter and Handcock, 2006) and mixed graphical models (see Lauritzen and Wermuth, 1989; Lee and Hastie, 2015; Cheng et al., 2017) for explaining relationships among actors in networks, interaction point pro-

cess models (see Strauss, 1975; Goldstein et al., 2015) for describing spatial patterns among points, and Conway–Maxwell–Poisson regression models (see Conway and Maxwell, 1962; Shmueli et al., 2005; Chaniyalidis et al., 2018) for over- or under-dispersed count data. Interest in models having intractable normalizing functions has increased rapidly during the last two decades. Indeed, a look at the Web of Science database shows that the number of yearly citations of articles on this subject has grown quadratically since 2005. Unfortunately, standard Markov chain Monte Carlo (MCMC) algorithms, the workhorse of Bayesian inference in the past few decades, cannot be applied to these models. Suppose we have data $\mathbf{x} \in \mathcal{X}$ generated from a probability model $f(\mathbf{x} | \boldsymbol{\theta})$ with likelihood function $L(\boldsymbol{\theta} | \mathbf{x}) = h(\mathbf{x} | \boldsymbol{\theta})/c(\boldsymbol{\theta})$, where $c(\boldsymbol{\theta})$ is a normalizing function, and a prior density $p(\boldsymbol{\theta})$. The posterior density of $\boldsymbol{\theta}$ is $\pi(\boldsymbol{\theta} | \mathbf{x}) \propto p(\boldsymbol{\theta})h(\mathbf{x} | \boldsymbol{\theta})/c(\boldsymbol{\theta})$, which brings about so-called doubly intractable posterior distributions. A major problem in constructing a standard MCMC algorithm for such models is that $c(\boldsymbol{\theta})$ cannot be easily evaluated. The Metropolis–Hastings (MH) algorithm (Metropolis et al., 1953; Hastings, 1970) acceptance probability at each step requires evaluating the unknown ratio $c(\boldsymbol{\theta})/c(\boldsymbol{\theta}')$, where $\boldsymbol{\theta}'$ denotes the proposed value.

A wide range of computational methods have been proposed for Bayesian inference for doubly intractable posterior distributions (see Park and Haran, 2018, for a review). There are asymptotically exact algorithms whose Markov chain has a stationary distribution equal to its target distribution (cf. Møller et al., 2006; Murray et al., 2006; Atchadé et al., 2013; Lyne et al., 2015; Liang et al., 2016). Throughout this manuscript we use “target distribution” to refer to the posterior distribution of interest. Some of the asymptotically exact algorithms are available only for a small class of probability models having intractable normalizing functions. The other algorithms are complicated to construct and have to be carefully tuned to provide reliable inference. These algorithms tend to be computationally intensive (Park and Haran, 2018). There are asymptotically inexact algorithms that may be much faster and can be applied to a wider class of problems (cf. Liang, 2010; Alquier et al., 2016; Park and Haran, 2020). An asymptotically inexact algorithm either converges to an approximation of the target or is not known to converge to any distribution. The performance of these algorithms relies on the choice of various tuning parameters. It is also not always easy to judge the tradeoffs between using a faster asymptotically inexact algorithm and a potentially more computationally expensive but asymptotically exact algorithm. Hence, it is crucial to have diagnostics that provide guidance for users to carefully tune their algorithms to provide reliable results.

There is a large literature on convergence diagnostics for MCMC algorithms (see Cowles and Carlin, 1996; Flegal and Jones, 2011; Roy, 2020, for a review). In fact, given the wide usage of MCMC and the importance of MCMC diagnostics, some of the best known MCMC diagnostics papers have thousands of citations (cf. Gelman and Rubin, 1992; Geyer, 1992; Geweke et al., 1991) and at this point perhaps often even get used without citation. However, the literature on assessing the quality of approximations provided by asymptotically inexact algorithms is very limited. There are several approaches that measure the deviation between sample means and target expectations whose values are known (Fan et al., 2006; Gorham and Mackey, 2015, 2017). Lee et al. (2019) and Xing et al. (2019) have provided tools for assessing the coverage of approximate credible intervals. These are laudable innovations, but they are not available for asymptotically inexact algorithms for doubly intractable distributions. This motivates our development of sample quality measures that assist scientists in tuning these algorithms.

In this article we describe a new diagnostic method that uses the well known second Bartlett identity (Bartlett, 1953a,b). Our method is, in principle, applicable in virtually any likelihood-based

context where misspecification is of concern. We develop a Monte Carlo approximation to this new diagnostic that is applicable to intractable normalizing function models. We also develop an approximate version of the kernel Stein discrepancy introduced by Gorham and Mackey (2017), making this available for doubly intractable distributions. This diagnostic asymptotically inherits the same convergence properties as that of Gorham and Mackey (2017) and thus can be used for diagnosing convergence of a sequence of sample distributions to the target distribution. Following Gorham and Mackey (2017), we think of our diagnostics as measuring “sample quality.”

The remainder of this article is organized as follows. In Section 2 we briefly describe computational methods for models with intractable normalizing functions. In Section 3 we discuss the need for diagnostics for tuning asymptotically inexact algorithms. In Section 4 we propose a new diagnostic for asymptotically exact and inexact methods, and we develop an approximation for doubly intractable distributions. In Section 5 we briefly describe the kernel Stein discrepancy introduced by Gorham and Mackey (2017) and propose its Monte Carlo approximation for intractable normalizing function models. We provide theoretical justification for our diagnostics. In Section 7 we describe the application of our diagnostic approaches to several algorithms in the context of three different challenging examples and study the computational complexity and variability of our diagnostics. Finally, in Section 8 we conclude with a brief summary and discussion.

2. Inference for Models with Intractable Normalizing Functions

Several computational methods have been developed for Bayesian inference for models with intractable normalizing functions. Park and Haran (2018) categorized these algorithms into two general, overlapping classes: (1) *auxiliary variable approaches*, which introduce an auxiliary variable and cancel out the normalizing functions in the Metropolis–Hastings acceptance probability (Møller et al., 2006; Murray et al., 2006; Liang, 2010; Liang et al., 2016), and (2) *likelihood approximation approaches*, which directly approximate the normalizing functions and plug the approximations into the Metropolis–Hastings acceptance probability (Atchadé et al., 2013; Lyne et al., 2015; Alquier et al., 2016; Park and Haran, 2020). An important characteristic of these algorithms is whether they are “asymptotically exact” or “asymptotically inexact.” Asymptotically exact algorithms generate a sequence whose asymptotic distribution is exactly equal to the target distribution. Asymptotically inexact algorithms generate a sequence that does not converge to the target distribution (or to any distribution in some cases).

Asymptotically exact algorithms have attractive theoretical properties but can often be computationally burdensome or even infeasible for challenging models. For instance, Møller et al. (2006) and Murray et al. (2006) depend on perfect sampling (cf. Propp and Wilson, 1996), an algorithm that generates an auxiliary variable exactly from the target distribution using bounding Markov chains. Perfect samplers tend to be very computationally expensive and are available only for a small class of probability models. Atchadé et al. (2013) and Liang et al. (2016) propose asymptotically exact algorithms that do not need perfect sampling. Atchadé et al. (2013)’s adaptive MCMC (ALR) algorithm introduces multiple particles $\theta^{(1)}, \dots, \theta^{(d)}$ over the parameter space and approximates $c(\theta)$ in the acceptance probability through importance sampling using the entire sample path of the algorithm. Liang et al. (2016)’s adaptive exchange (AEX) algorithm runs an auxiliary chain and a target chain simultaneously. At each iteration, the auxiliary chain simulates and stores a sample from a set of distributions, $\{h(\mathbf{x} \mid \theta^{(1)})/c(\theta^{(1)}), \dots, h(\mathbf{x} \mid \theta^{(d)})/c(\theta^{(d)})\}$, where $\theta^{(1)}, \dots, \theta^{(d)}$ are predetermined particles over the parameter space. The target chain generates a posterior

sample via the exchange algorithm (Murray et al., 2006), where an auxiliary variable \mathbf{y} is sampled from the cumulative samples in the auxiliary chain instead of exact sampling of \mathbf{y} . These algorithms require storing simulated auxiliary data or their sufficient statistics with each iteration. The computational and memory costs are very expensive for models without low-dimensional sufficient statistics. Pseudo-marginal MCMC algorithms (Beaumont, 2003; Andrieu and Roberts, 2009) are approaches that use an unbiased Monte Carlo approximation of an intractable likelihood. Lyne et al. (2015) constructed an unbiased likelihood estimate for models with intractable normalizing functions. To obtain a single estimate, their approach requires multiple Monte Carlo approximations to the normalizing constant and thus can often be computationally prohibitive.

Several computationally efficient asymptotically inexact algorithms have also been proposed. For instance, the double Metropolis–Hastings (DMH) sampler, proposed by Liang (2010), replaces perfect sampling with a standard Metropolis–Hastings algorithm, an “inner sampler,” at each iteration of the exchange algorithm. The DMH algorithm is easy to implement and is computationally efficient compared to the other algorithms discussed thus far. But the inner sampling becomes more computationally expensive with an increase in the dimension of the data. For large data problems, Park and Haran (2020) proposed a function emulation algorithm that approximates the likelihood normalizing function (or full likelihood function) at several parameter values and interpolates the function using a Gaussian process. This provides significant gains in computational efficiency. More asymptotically exact and inexact algorithms are found in Park and Haran (2018).

Both asymptotically exact and inexact algorithms require careful tuning in order to provide good approximations in a reasonable amount of time. For instance, the ALR and AEX algorithms require users to select an appropriate number d of particles to cover the important region of the parameter space. The DMH algorithm requires users to decide the length m of the inner sampler for generating an auxiliary variable. As m grows large the auxiliary variable becomes approximately a draw from the probability model at the expense of longer computing time. However, currently there is little guidance on how to tune these algorithms, and most of them rely on simulation studies. Also, the behavior of the algorithm varies across models or across datasets for a given model.

3. The Need for Diagnostics for Intractable Normalizing Function Problems

There is a vast literature on MCMC convergence diagnostics (see Cowles and Carlin, 1996; Flegal and Jones, 2011; Roy, 2020, for a review). However, these diagnostics may not be adequate for asymptotically inexact algorithms. Suppose we have a sample generated by an asymptotically inexact method. Such a sample may not have an asymptotic distribution, or said sample may converge but to a mere approximation of the target distribution. Standard MCMC diagnostics assess whether the sample has converged to its asymptotic distribution but do not assess whether the asymptotic distribution is equal to the target distribution. As discussed in Section 2, however, asymptotically exact algorithms for models with intractable normalizing functions are available only for a few special cases, and even for these cases computing tends to be quite burdensome.

Several approaches based on measuring the difference between sample means and target expectations have been proposed for assessing the quality of approximations provided by asymptotically inexact algorithms. Fan et al. (2006) proposed score function (i.e., gradient of the log likelihood) diagnostics for assessing estimates of some quantities, the values of which are known under the target distribution. They suggested plotting the Monte Carlo estimate versus the sample size together with error bounds. Gorham and Mackey (2015) pointed out limitations of the score function

diagnostics caused by considering only a finite class of functions and introduced a new diagnostic method based on a Stein discrepancy. Gorham and Mackey (2015) defined Stein discrepancies that bound the discrepancy between the sample mean and the target expectation over a large collection of functions whose target expectations are zero. The Stein discrepancies are supported by a theory of weak convergence and are attainable by solving a linear program. Combining this idea with the theory of reproducing kernels, Gorham and Mackey (2017) provided a closed-form kernel Stein discrepancy with sound theoretical support analogous to that of Gorham and Mackey (2015) (see Section 5.1 for details). The kernel Stein discrepancy requires evaluating the score $\nabla_{\theta} \log L(\theta | \mathbf{x})$ of a full likelihood. This can be often expressed as a sum over data points, i.e., $\nabla_{\theta} \log L(\theta | \mathbf{x}) = \sum_{i=1}^n \nabla_{\theta} \log L(\theta | x_i)$. Stochastic Stein discrepancies (SSDs) approximate the likelihood’s score using subsampling data points to expedite computation (Liu and Wang, 2016; Ranganath et al., 2016; Aicher et al., 2023). Gorham et al. (2020) provided theoretical justification for SSDs. These approaches are useful for comparing asymptotically approximate samplers and for selecting tuning parameters for such samplers. However, all of them require evaluating the score function of the target density, which is not possible for doubly intractable posterior distributions. In contrast, our approaches apply broadly to asymptotically exact and inexact algorithms even for such challenging problems. To our knowledge, no other diagnostics are currently available for asymptotically inexact algorithms for intractable normalizing function models. We have studied our diagnostics as applied in several challenging real data settings. In addition, we provide theoretical justification for our methods.

4. Curvature Diagnostics

In this section we introduce two new diagnostics that are useful for tuning asymptotically exact and inexact algorithms: a curvature diagnostic (CD) and an approximate curvature diagnostic (ACD). The curvature diagnostic is based on the second Bartlett identity from the classical theory of maximum likelihood. The approximate curvature diagnostic is an approximation of the CD that is suitable for intractable normalizing function problems.

4.1 A General Purpose Curvature Diagnostic

consider a sample $\theta^{(1)}, \dots, \theta^{(n)}$ generated by an algorithm having $\pi(\theta | \mathbf{x})$ as its target distribution. Our aim is to determine whether the sample will produce a good approximation to some quantity of interest, e.g., $E_{\pi}\{z(\theta)\}$, where $z(\theta)$ is a real-valued function. Our diagnostic, the curvature diagnostic, is inspired by the method for obtaining the asymptotic variance of a maximum likelihood estimator under a misspecified model. When the model is misspecified, the second Bartlett identity does not hold, which is to say (see details below) the sensitivity matrix is not equal to the variability matrix. And so the asymptotic variance of the estimator does not simplify to the inverse of the Fisher information. Our curvature diagnostic uses the dissimilarity between the sensitivity matrix and the variability matrix to assess the quality of the sample. We provide details in the following paragraph.

Let $\mathbf{u}(\theta) = \nabla_{\theta} \log \pi(\theta | \mathbf{x})$ be the score function of the posterior density $\pi(\theta | \mathbf{x})$. The posterior density has the sensitivity matrix

$$E_{\pi}\{-\nabla_{\theta} \mathbf{u}(\theta)\} = \int_{\Theta} -\nabla_{\theta} \mathbf{u}(\theta) \pi(\theta | \mathbf{x}) d\theta$$

and the variability matrix $\text{var}_\pi \{\mathbf{u}(\boldsymbol{\theta})\} = \mathbb{E}_\pi \{\mathbf{u}(\boldsymbol{\theta})\mathbf{u}(\boldsymbol{\theta})^\top\}$. Note that the identity for the variability matrix follows from Bartlett's first identity: $\mathbb{E}_\pi \{\mathbf{u}(\boldsymbol{\theta})\} = \mathbf{0}$. Let $H(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \mathbf{u}(\boldsymbol{\theta})$, $J(\boldsymbol{\theta}) = \mathbf{u}(\boldsymbol{\theta})\mathbf{u}(\boldsymbol{\theta})^\top$, and $\mathbf{d}(\boldsymbol{\theta}) = \text{vech}[J(\boldsymbol{\theta}) + H(\boldsymbol{\theta})]$, where $\text{vech}(M)$ denotes the half-vectorization of the matrix M . We have $\mathbb{E}_\pi \{\mathbf{d}(\boldsymbol{\theta})\} = \mathbf{0}$ by Bartlett's second identity: $\mathbb{E}_\pi \{-H(\boldsymbol{\theta})\} = \mathbb{E}_\pi \{J(\boldsymbol{\theta})\}$. Using the sample $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ we construct a Monte Carlo approximation to the half-vectorized difference between the sensitivity and variability matrices as

$$\mathbf{d}_n := \frac{1}{n} \sum_{i=1}^n \mathbf{d}(\boldsymbol{\theta}^{(i)}). \quad (1)$$

Suppose that $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ are independent. By the central limit theorem, we have

$$\sqrt{n} \mathbf{d}_n \xrightarrow{d} \mathbf{N}(\mathbf{0}, V),$$

where $V := \text{cov}_\pi \{\mathbf{d}(\boldsymbol{\theta})\} = \mathbb{E}_\pi \{\mathbf{d}(\boldsymbol{\theta})\mathbf{d}(\boldsymbol{\theta})^\top\}$. Our unbiased and consistent approximation of V is calculated as

$$V_n := \frac{1}{n} \sum_{i=1}^n \mathbf{d}(\boldsymbol{\theta}^{(i)})\mathbf{d}(\boldsymbol{\theta}^{(i)})^\top. \quad (2)$$

If $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ are from a Markov process and thus dependent, then by the Markov chain Monte Carlo central limit theorem we have

$$\sqrt{n} \mathbf{d}_n \xrightarrow{d} \mathbf{N}(\mathbf{0}, \Sigma),$$

where $\Sigma := \text{cov}_\pi \{\mathbf{d}(\boldsymbol{\theta}^{(i)})\} + 2 \sum_{k=1}^{\infty} \text{cov}_\pi \{\boldsymbol{\theta}^{(i)}, \boldsymbol{\theta}^{(i+k)}\}$. The asymptotic covariance matrix Σ can be estimated by the multivariate batch means method (Chen and Seila, 1987; Flegal and Jones, 2011). Let $n = a_n b_n$ where a_n is the number of batches and b_n is the batch size. The batch means estimate of Σ is calculated as

$$\Sigma_n := \frac{b_n}{a_n - 1} \sum_{j=1}^{a_n} (\bar{\mathbf{d}}_{b_n}^j - \mathbf{d}_n)(\bar{\mathbf{d}}_{b_n}^j - \mathbf{d}_n)^\top, \quad (3)$$

where $\bar{\mathbf{d}}_{b_n}^j = \sum_{i=(j-1)b_n+1}^{j b_n} \mathbf{d}(\boldsymbol{\theta}^{(i)})$. The batch means estimator is strongly consistent under some conditions (Damerdjii, 1994; Jones et al., 2006; Vats et al., 2019). Because it is simple to construct and appears to work well in practice under quite a wide range of settings (cf. Flegal et al., 2008), we suggest using batch means. Our curvature diagnostic is then defined as follows.

Definition 1 (Curvature Diagnostic (CD)) *consider a sample $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ generated by an algorithm having $\pi(\boldsymbol{\theta} \mid \mathbf{x})$ as its target. If the sample is independent, our curvature diagnostic is defined as $\mathcal{C}_n(\mathbf{x}) := n \mathbf{d}_n^\top V_n^{-1} \mathbf{d}_n$. If the sample is from a Markov process, our curvature diagnostic is defined as $\mathcal{C}_n^{\text{BM}}(\mathbf{x}) := n \mathbf{d}_n^\top \Sigma_n^{-1} \mathbf{d}_n$.*

The CD has an asymptotic $\chi^2(r)$ distribution, where $r = p(p+1)/2$ and p is the dimension of $\boldsymbol{\theta}$, if the asymptotic distribution of the sample is equal to the target distribution. A simple and effective heuristic for determining n is to plot \hat{V}_n against the posterior sample size n and select n at which the

approximations appear to have stabilized. Given n , an abnormally large value of CD with respect to $\chi^2(r)$ can signal poor sample quality. We use the $1 - \alpha$ quantile of the $\chi^2(r)$ as a threshold for our diagnostic. A sample path for which the CD value is below the threshold is considered to have an asymptotic distribution that is reasonably close to the target distribution. We found that $\alpha = 0.01$ performed well in our simulation experiments, and so we take $\alpha = 0.01$ in the sequel.

We note that our curvature diagnostic is, in principle, applicable not only in intractable normalizing function problems but in virtually any likelihood-based context where misspecification is of concern. Moreover, the framework we present in this article can be extended to incorporate higher moments.

4.2 An Approximate Curvature Diagnostic for Intractable Normalizing Function Problems

If the normalizing function $c(\boldsymbol{\theta})$ of the likelihood is intractable, it is not possible to evaluate the curvature diagnostic since the diagnostic involves $\nabla_{\boldsymbol{\theta}} \log c(\boldsymbol{\theta})$ and $\nabla_{\boldsymbol{\theta}}^2 \log c(\boldsymbol{\theta})$. The intractable terms can be written as

$$\nabla_{\boldsymbol{\theta}} \log c(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta})} \{ \nabla_{\boldsymbol{\theta}} \log h(\mathbf{X} | \boldsymbol{\theta}) \} \quad (4)$$

$$\begin{aligned} \nabla_{\boldsymbol{\theta}}^2 \log c(\boldsymbol{\theta}) &= \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta})} \{ \nabla_{\boldsymbol{\theta}}^2 \log h(\mathbf{X} | \boldsymbol{\theta}) \} + \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta})} \left\{ [\nabla_{\boldsymbol{\theta}} \log h(\mathbf{X} | \boldsymbol{\theta})] [\nabla_{\boldsymbol{\theta}} \log h(\mathbf{X} | \boldsymbol{\theta})]^\top \right\} \\ &\quad - \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta})} \{ \nabla_{\boldsymbol{\theta}} \log h(\mathbf{X} | \boldsymbol{\theta}) \} \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta})} \{ \nabla_{\boldsymbol{\theta}} \log h(\mathbf{X} | \boldsymbol{\theta}) \}^\top. \end{aligned} \quad (5)$$

We can replace the expectations with their Monte Carlo approximations using auxiliary variables $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}$ that are generated exactly from $f(\cdot | \boldsymbol{\theta})$ or generated by a Monte Carlo algorithm having $f(\cdot | \boldsymbol{\theta})$ as its target distribution. We provide details in Appendix A.

To reduce computational cost, we employ self-normalizing importance sampling (SNIS; Tan and Friel, 2020). For $i = 1, \dots, n$ we need to estimate $\mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta}^{(i)})} \{ g(\mathbf{X}) \}$ where g denotes any function inside of the expectations in (4) and (5). For some $\boldsymbol{\psi} \in \Theta$, the expectation can be written as

$$\mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta}^{(i)})} \{ g(\mathbf{X}) \} = \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\psi})} \left\{ g(\mathbf{X}) \frac{h(\mathbf{X} | \boldsymbol{\theta}^{(i)})}{h(\mathbf{X} | \boldsymbol{\psi})} \right\} / \mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\psi})} \left\{ \frac{h(\mathbf{X} | \boldsymbol{\theta}^{(i)})}{h(\mathbf{X} | \boldsymbol{\psi})} \right\}.$$

The SNIS introduces a finite set $S = \{\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_M\}$ of particles over the parameter space Θ and simulates auxiliary variables $\mathbf{y}_r^{(1)}, \dots, \mathbf{y}_r^{(N)}$ from $f(\cdot | \boldsymbol{\psi}_r)$ for each $r = 1, \dots, M$. For $i = 1, \dots, n$, the consistent estimate of $\mathbb{E}_{\mathbf{X} \sim f(\cdot | \boldsymbol{\theta}^{(i)})} \{ g(\mathbf{X}) \}$ can be obtained by $\sum_{j=1}^N w_r^{(j)} g(\mathbf{y}_r^{(j)})$, where

$$w_r^{(j)} = \frac{h(\mathbf{y}_r^{(j)} | \boldsymbol{\theta}^{(i)})}{h(\mathbf{y}_r^{(j)} | \boldsymbol{\psi}_r)} / \sum_{j'=1}^N \frac{h(\mathbf{y}_r^{(j')} | \boldsymbol{\theta}^{(i)})}{h(\mathbf{y}_r^{(j')} | \boldsymbol{\psi}_r)},$$

and $\boldsymbol{\psi}_r$ is the particle nearest to $\boldsymbol{\theta}^{(i)}$. For particle sampling we use a quasi-random number generator (Faure and Lemieux, 2009). The quasi-random number generation can be done easily using R package `qrng`. We use Mahalanobis distance to measure closeness, and we estimate the covariance matrix using $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$. We found that $M = 200p$ and $N = 10,000$ performed well in our simulation experiments.

By plugging these approximations into $\mathbf{d}(\boldsymbol{\theta})$ we obtain its approximation $\hat{\mathbf{d}}_N(\boldsymbol{\theta})$. By replacing $\mathbf{d}(\boldsymbol{\theta})$ with its approximation $\hat{\mathbf{d}}_N(\boldsymbol{\theta})$ in (1), (2), and (3) we obtain $\hat{\mathbf{d}}_{n,N}$, $\hat{V}_{n,N}$, and $\hat{\Sigma}_{n,N}$. We

note that the number $a_{n,N}$ of batches and the length $b_{n,N}$ of each batch depend on n and N for $\hat{\Sigma}_{n,N}$. Then an approximate version of the curvature diagnostic for intractable normalizing function problems can be defined as follows.

Definition 2 (Approximate Curvature Diagnostic (ACD)) *consider a sample $\theta^{(1)}, \dots, \theta^{(n)}$ generated by an algorithm having $\pi(\theta | \mathbf{x})$ as its target. If the sample is independent, our approximate curvature diagnostic is defined as $\hat{C}_{n,N}(\mathbf{x}) := n \hat{\mathbf{d}}_{n,N}^\top \hat{V}_{n,N}^{-1} \hat{\mathbf{d}}_{n,N}$. If the sample is from a Markov process, our approximate curvature diagnostic is defined as $\hat{C}_{n,N}^{\text{BM}}(\mathbf{x}) := n \hat{\mathbf{d}}_{n,N}^\top \hat{\Sigma}_{n,N}^{-1} \hat{\mathbf{d}}_{n,N}$.*

To provide theoretical justification for $\hat{C}_{n,N}(\mathbf{x})$, we make the following assumptions regarding the prior density $p(\theta)$ and the unnormalized probability model density $h(\mathbf{x} | \theta)$. We use $\|\cdot\|_{\max}$ to represent the max norm of vectors or matrices.

Assumption 3 $\|\nabla_{\theta} \log p(\theta)\|_{\max} < \infty$ and $\|\nabla_{\theta}^2 \log p(\theta)\|_{\max} < \infty$.

Assumption 4 $\|\nabla_{\theta} \log h(\mathbf{x} | \theta)\|_{\max} < \infty$ and $\|\nabla_{\theta}^2 \log h(\mathbf{x} | \theta)\|_{\max} < \infty$.

Since the prior density is determined by the user, the prior's score function and Hessian matrix may be assumed to be bounded. In many applications Assumption 4 may be checked easily. In particular, Assumption 4 is satisfied with high probability for exponential families. For a probability model in an exponential family, the model's score is its summary statistics and the first inequality of Assumption 4 is satisfied with high probability (Chazottes et al., 2007). The model's Hessian matrix is zero and the second inequality is satisfied almost surely. Under these assumptions, Theorem 5 quantifies the distance between the asymptotic covariance matrix V and its two-stage approximation $\hat{V}_{n,N}$.

Theorem 5 *Consider an i.i.d. sample $\theta^{(1)}, \dots, \theta^{(n)}$ the asymptotic distribution of which is $\pi(\theta | \mathbf{x})$. If Assumptions 3 and 4 hold, we have $\|\hat{V}_{n,N} - V\|_{\max} \leq \mathcal{O}(n^{-1/2}) + \mathcal{O}(N^{-1/2})$ almost surely.*

A proof of Theorem 5 is provided in Appendix B. Provided the theorem holds, the two-stage approximation $\hat{V}_{n,N}$ will get closer to V as the posterior sample size n and the auxiliary sample size N increase.

To provide theoretical justification for $\hat{C}_{n,N}^{\text{BM}}(\mathbf{x})$, we make the following assumption in addition to Conditions 1 and 2 of Vats et al. (2019).

Assumption 6 *The batch size $b_{n,N}$ is an integer sequence that satisfies $b_{n,N} \rightarrow \infty$ and $(b_{n,N}/N)^{1/2} \rightarrow 0$ as $N \rightarrow \infty$ and $n \rightarrow \infty$.*

Under Assumptions 3, 4, and 6 and Conditions 1 and 2 of Vats et al. (2019), Theorem 7 quantifies the distance between Σ and $\hat{\Sigma}_{n,N}$.

Theorem 7 *consider a sample $\theta^{(1)}, \dots, \theta^{(n)}$ generated from a Markov process that is an π -invariant polynomially ergodic and of order $m > (1 + \epsilon_1)(1 + 2/\delta)$ for some $\epsilon_1 > 0$ and $\delta > 0$. Then Condition 1 of Vats et al. (2019) holds with $\gamma(n) = n^{1/2-\lambda}$ for some $\lambda > 0$. If Assumptions 3, 4, and 6, and Condition 2 of Vats et al. (2019) hold, we have*

$$\|\hat{\Sigma}_{n,N} - \Sigma\|_{\max} \leq \mathcal{O} \left[\{(\log n)/b_{n,N}\}^{1/2} n^{1/2-\lambda} \right] + \mathcal{O} \left[(b_{n,N}/N)^{1/2} \right]$$

almost surely.

A proof of Theorem 7 is provided in Appendix C. Provided the theorem holds, the two-stage approximation $\hat{\Sigma}_{n,N}$ will get closer to Σ as the posterior sample size n and the auxiliary sample size N increase.

A simple heuristic for determining N is to plot the approximation $\hat{d}_N(\boldsymbol{\theta})$ against the auxiliary sample size N and select N at which the approximation appears to stabilize. Similarly we choose n at which the approximation $\hat{V}_{n,N}$ or $\hat{\Sigma}_{n,N}$ appears to stabilize for the pre-determined N . We henceforth take the batch size to be $b_{n,N} = \min\{n^{1/3}, N^{2/5}\}$ which satisfies Assumption 6 and Condition 2 of Vats et al. (2019).

5. A Kernel Stein Discrepancy

In this section we briefly describe an inverse multiquadric kernel Stein discrepancy (IMQ KSD) introduced by Gorham and Mackey (2017) which is a kernel Stein discrepancy-based diagnostic for assessing the convergence of a sequence of sample distributions to its target distribution and has theoretical support. To make this approach available for doubly intractable target distributions, we develop its Monte Carlo approximation, AIKS (approximate inverse multiquadric kernel Stein discrepancy) and show that it asymptotically inherits the same convergence properties as IMQ KSD.

5.1 An Inverse Multiquadric Kernel Stein Discrepancy

Consider a target distribution P under which direct integration is infeasible. Suppose we approximate integration under P using a weighted sample $Q_n = \sum_{i=1}^n q_n(\boldsymbol{\theta}^{(i)}) \delta_{\boldsymbol{\theta}^{(i)}}$ with sample points $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ and a probability mass function q_n . Each Q_n provides an approximation $E_{Q_n} \{z(\boldsymbol{\theta})\} = \sum_{i=1}^n z(\boldsymbol{\theta}^{(i)}) q_n(\boldsymbol{\theta}^{(i)})$ to each intractable expectation $E_P \{z(\boldsymbol{\theta})\}$ for a real-valued function $z(\boldsymbol{\theta})$. To assess the quality of the approximation, one may consider discrepancies quantifying the maximum expectation error over a set of test functions \mathcal{Z} : $d_{\mathcal{Z}}(Q_n, P) = \sup_{z \in \mathcal{Z}} |E_P \{z(\boldsymbol{\theta})\} - E_{Q_n} \{z(\boldsymbol{\theta})\}|$. When \mathcal{Z} is large enough the discrepancy is called an integral probability metric (IPM) (Müller, 1997) and $d_{\mathcal{Z}}(Q_n, P) \rightarrow 0$ only if $Q_n \Rightarrow P$ for any sequence Q_n . We use \Rightarrow to denote the weak convergence of a sequence of probability measures. However, it is not practical to use IPM for assessing a sample since $E_P \{z(\boldsymbol{\theta})\}$ may not be computable for some $z \in \mathcal{Z}$.

According to Stein’s method (Stein, 1972), Gorham and Mackey (2015) defined a Stein discrepancy as $\mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}) = \sup_{g \in \mathcal{G}} |E_{Q_n} \{(\mathcal{T}_P g)(\boldsymbol{\theta})\}|$ for a Langevin Stein operator \mathcal{T}_P and a Stein set \mathcal{G} that satisfy $E_P \{(\mathcal{T}_P g)(\boldsymbol{\theta})\} = 0$ for all $g \in \mathcal{G}$. The Stein discrepancy is the maximum expectation error over the Stein set \mathcal{G} given the Stein operator \mathcal{T}_P . This avoids explicit integration under P by selecting appropriate \mathcal{T}_P and \mathcal{G} that lead the target expectation to zero. Gorham and Mackey (2017) selected kernel Stein set based on the inverse multiquadric kernel and defined the inverse multiquadric kernel Stein discrepancy $\mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|})$ for any norm $\|\cdot\|$, which admits closed-form solution.

Definition 8 (IMQ KSD (Gorham and Mackey, 2017)) *Let $k(\mathbf{x}, \mathbf{y}) = (c^2 + \|\mathbf{x} - \mathbf{y}\|_2^2)^\beta$ for some $c > 0$ and $\beta \in (-1, 0)$. For each $j \in \{1, \dots, p\}$ construct the Stein kernel*

$$k_0^j(\mathbf{x}, \mathbf{y}) = u_j(\mathbf{x})u_j(\mathbf{y})k(\mathbf{x}, \mathbf{y}) + u_j(\mathbf{x})\nabla_{y_j}k(\mathbf{x}, \mathbf{y}) + u_j(\mathbf{y})\nabla_{x_j}k(\mathbf{x}, \mathbf{y}) + \nabla_{x_j}\nabla_{y_j}k(\mathbf{x}, \mathbf{y}),$$

where u_j is the j th entry of the score function of the target density. Then IMQ KSD $\mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|}) = \|\mathbf{w}\|$, where $w_j = \sqrt{\frac{1}{n^2} \sum_{k,l=1}^n k_0^j(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)})}$.

Computation of \mathbf{w} is parallelizable over sample pairs $(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)})$ and coordinates j . Gorham and Mackey (2017) provided theoretical justification for its use for diagnosing convergence of a sequence Q_n to its target distribution P (see Theorem 8 and Proposition 9 of Gorham and Mackey (2017)).

5.2 An Approximate Inverse Multiquadric Kernel Stein Discrepancy for Intractable Normalizing Function Problems

When the target distribution P is doubly intractable, computation of IMQ KSD is not feasible. This is because IMQ KSD requires evaluating the score function of the target density. In this section we introduce an approximate version of IMQ KSD with L^2 norm by replacing the score function with its Monte Carlo approximation. The approximate inverse multiquadric kernel Stein discrepancy (AIKS) is defined as follows.

Definition 9 (AIKS) *Let $k(\mathbf{x}, \mathbf{y}) = (c^2 + \|\mathbf{x} - \mathbf{y}\|_2^2)^\beta$ for some $c > 0$ and $\beta \in (-1, 0)$. Define an approximate Stein kernel as*

$$\hat{k}_0(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^p \left\{ \hat{u}_j(\mathbf{x})\hat{u}_j(\mathbf{y})k(\mathbf{x}, \mathbf{y}) + \hat{u}_j(\mathbf{x})\nabla_{y_j}k(\mathbf{x}, \mathbf{y}) + \hat{u}_j(\mathbf{y})\nabla_{x_j}k(\mathbf{x}, \mathbf{y}) + \nabla_{x_j}\nabla_{y_j}k(\mathbf{x}, \mathbf{y}) \right\},$$

where \hat{u}_j is the j th entry of the approximate score function of the target density. Then the AIKS is defined as $\hat{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) = \sqrt{\frac{1}{n^2} \sum_{k,l=1}^n \hat{k}_0(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)})}$.

For doubly intractable target density $\pi(\boldsymbol{\theta} \mid \mathbf{x})$, the score function can be approximated as described in Section 4.2. Under the assumptions of IMQ KSD on the target distribution, Theorem 10 quantifies the distance between AIKS and IMQ KSD for L^2 norm.

Theorem 10 *For a target distribution P having a bounded score function and a sample distribution Q_n targeting P , let $\mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2})$ and $\hat{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2})$ denote IMQ KSD and AIKS, respectively, for the sample distribution. Then*

$$\left| \hat{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) - \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) \right| \leq \mathcal{O} \left\{ N^{-1/4} \right\}$$

almost surely where N is the auxiliary sample size used to approximate the score function of the target density.

A proof of Theorem 10 is provided in Appendix D. The theorem implies that AIKS will get closer to IMQ KSD as the auxiliary sample size N increases. The following proposition shows that AIKS asymptotically inherits the same convergence properties as IMQ KSD.

Proposition 11 (i) *Detecting non-convergence: for a distantly dissipative target distribution P , if $\hat{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) \rightarrow 0$ as $n, N \rightarrow \infty$, then $Q_n \Rightarrow P$ as $n \rightarrow \infty$.*

(ii) *Detecting convergence: for a target distribution P having Lipschitz score function with $E_P \{ \|\mathbf{u}(\boldsymbol{\theta})\|_2^2 \} < \infty$, if the Wasserstein distance $d_{\mathcal{W}_{\|\cdot\|_2}}(Q_n, P) \rightarrow 0$ as $n \rightarrow \infty$, then $\hat{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) \rightarrow 0$ as $n, N \rightarrow \infty$.*

A proof of Theorem 11 is provided in Appendix E. In this article we use $\beta = -1/2$, and $c = 1$.

The AIKS is a form of two-sample V -statistics. Chwialkowski et al. (2016) obtain the asymptotic distribution of a kernel Stein discrepancy using the theory of V -statistics for dependent samples. Leucht and Neumann (2013) provides the asymptotic distribution of a V -statistic for dependent samples. By a direct application of Theorem 2.1 of Leucht and Neumann (2013), the following proposition characterizes the asymptotic behavior of AIKS.

Proposition 12 *Consider a sample $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ generated from a P -invariant β -mixing process with $\sum_{r=1}^{\infty} r^2 \sqrt{\beta(r)} < \infty$. If \hat{k}_0 is Lipschitz continuous and $E_P[\hat{k}_0(\boldsymbol{\theta}, \boldsymbol{\theta})] < \infty$, then as $n, N \rightarrow \infty$, we have $n\hat{\mathcal{S}}^2(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) \xrightarrow{d} \sum_{k=1}^{\infty} \lambda_k Z_k^2$ almost surely. The $\{\lambda_k\}$ are the eigenvalues of kernel $\hat{k}_0(\boldsymbol{\theta}, \boldsymbol{\theta}')$ under P , i.e., they are the solutions of $\lambda_k \phi_k(\boldsymbol{\theta}) = \int_{\boldsymbol{\theta}_r} \hat{k}_0(\boldsymbol{\theta}, \boldsymbol{\theta}_r) \phi_k(\boldsymbol{\theta}_r) P(\boldsymbol{\theta}_r) d\boldsymbol{\theta}_r$. The $\{Z_k\}$ are centered, jointly Gaussian random variables with $\text{cov}(Z_i, Z_j) = \sum_{r=-\infty}^{\infty} \text{cov}[\phi_i(\boldsymbol{\theta}), \phi_j(\boldsymbol{\theta}_r)]$.*

The β -mixing is a notion of dependence that is weak enough for most practical uses. For instance, iid samples and stationary, geometrically ergodic Markov chains are β -mixing. The proof of Proposition 12 is a simple verification of assumptions and can be found in the Appendix F.

However, the asymptotic distribution of AIKS does not have a closed form in general. To obtain consistent estimates of quantiles of the asymptotic distribution, we employ a bootstrap method proposed by Leucht and Neumann (2013). We compute bootstrap samples as $\hat{\mathcal{S}}^{*2} = \frac{1}{n^2} \sum_{k,l=1}^n (W_k - \bar{W}) \hat{k}_0(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) (W_l - \bar{W})$, where $\bar{W} = \frac{1}{n} \sum_{k=1}^n W_k$ and $\{W_k, k = 1, \dots, n\}$ is an auxiliary random process for accommodating the dependence structure of the original process $\{\boldsymbol{\theta}^{(k)}, k = 1, \dots, n\}$ asymptotically when the kernel \hat{k}_0 is degenerate, i.e., when the sample is drawn from a P -invariant process. By Remark 2 of Leucht and Neumann (2013), we can construct the auxiliary process as $W_k = e^{-1/\xi} W_{k-1} + \sqrt{1 - e^{-2/\xi}} \epsilon_k$ where $W_0, \epsilon_1, \dots, \epsilon_n$ are independent standard normal random variables and ξ is a tuning parameter that performs a similar role to the block length of blockwise bootstrap methods. A simulation experiment by Leucht and Neumann (2013) shows that the choice of ξ does not significantly affect test performance. We use $\xi = 7$ and take the bootstrap sample size to be $b = 1,000$ in the sequel. We calculate the empirical $(1 - \alpha)$ quantile $\hat{\gamma}_{1-\alpha}$ of $n\hat{\mathcal{S}}^{*2}$. We use $\hat{\gamma}_{1-\alpha}$ as a threshold for AIKS. The following proposition shows that $\hat{\gamma}_{1-\alpha}$ consistently approximates the quantile of the limiting distribution of $n\hat{\mathcal{S}}^2(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2})$.

Proposition 13 *Consider a sample $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ generated from a P -invariant β -mixing process with $\sum_{r=1}^{\infty} r^2 \sqrt{\beta(r)} < \infty$. If \hat{k}_0 is Lipschitz continuous and $E_P[\hat{k}_0(\boldsymbol{\theta}, \boldsymbol{\theta})] < \infty$, then as $n, N \rightarrow \infty$, we have $\sup_x |P(n\hat{\mathcal{S}}^{*2} \leq x) - P(n\hat{\mathcal{S}}^2(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2}) \leq x)| \xrightarrow{P} 0$.*

The proof is a straightforward verification of assumptions that can be found in the Appendix G. A sample path for which $n\hat{\mathcal{S}}^2(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_2})$ is below $\hat{\gamma}_{1-\alpha}$ is considered to have an asymptotic distribution that is reasonably close to the target distribution.

A notable special case of doubly intractable posterior distributions arises when a prior distribution has an intractable normalizing constant that is a function of parameter of interest (cf. Rao et al., 2016; Vats et al., 2022). Suppose that we have a prior density $p(\boldsymbol{\theta} | \boldsymbol{\eta}) = g(\boldsymbol{\theta} | \boldsymbol{\eta})/z(\boldsymbol{\eta})$, where $z(\boldsymbol{\eta})$ is intractable. If hyperprior $k(\cdot)$ is assigned to $\boldsymbol{\eta}$, the posterior density is

$$\pi(\boldsymbol{\theta}, \boldsymbol{\eta} | \boldsymbol{x}) \propto \frac{g(\boldsymbol{\theta} | \boldsymbol{\eta})}{z(\boldsymbol{\eta})} k(\boldsymbol{\eta}) L(\boldsymbol{\theta} | \boldsymbol{x}),$$

Table 1: Power and Type I error rates of a multivariate normality test with increasing sample size n and dimension p , averaged over 1,000 simulations.

Method	n	Power						Type I error rate					
		$p=2$	5	10	15	20	25	2	5	10	15	20	25
CD	1000	1	1	1	1	1	1	0.02	0.03	0.04	0.07	0.10	0.14
	2000	1	1	1	1	1	1	0.01	0.03	0.03	0.03	0.05	0.08
	5000	1	1	1	1	1	1	0.01	0.01	0.01	0.01	0.01	0.02
KSD	1000	1	1	1	1	1	1	0.01	0.02	0.00	0.00	0.00	0.00
	2000	1	1	1	1	1	1	0.00	0.03	0.00	0.00	0.00	0.00
	5000	1	1	1	1	1	1	0.01	0.04	0.01	0.00	0.00	0.00

which leads to a doubly intractable posterior distribution. Our diagnostics ACD and AIKS can be applied to this case if auxiliary variables can be generated exactly from the prior distribution or generated by a Monte Carlo algorithm having $p(\boldsymbol{\theta} \mid \boldsymbol{\eta})$ as its target distribution. Using the auxiliary variables, we can approximate intractable derivatives of the prior and obtain our diagnostic quantities.

6. Comparison of Diagnostics

We recreate the experiment in Section 4 of Chwialkowski et al. (2016) and compare the power and Type I error rates of the CD-based and IMQ KSD-based tests. For $i = 1, \dots, n$ we generate $\boldsymbol{\theta}^{(i)} = \mathbf{z}^{(i)} + u_i \mathbf{e}_1$, where $\mathbf{z}^{(i)} \stackrel{\text{iid}}{\sim} \text{MVN}(\mathbf{0}, I_p)$, $u_i \stackrel{\text{iid}}{\sim} \text{Unif}(0, 1)$, and $\mathbf{e}_1 = (1, 0, \dots, 0)^\top$, for various combinations of sample size n and parameter dimension p . We compare the power of the CD test and IMQ KSD test to assess whether $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ were drawn from the target distribution $P = \text{MVN}(\mathbf{0}, I_p)$. We also generate samples from P and compare the Type I error rates of the two tests. The nominal significance level is $\alpha = 0.01$. The results, averaged over 1,000 simulations, are displayed in Table 1. Both tests provide high power for the values of n and p we considered. We see that the Type I error rate of the CD grows as p increases and $n \leq 2000$. The CD requires approximation of the asymptotic covariance matrix of $\mathbf{d}(\boldsymbol{\theta})$, which is a $p(p+1)/2$ by $p(p+1)/2$ matrix. A sufficiently large n is required to accurately estimate this covariance matrix. The Type I error rate of the CD is well controlled when $n = 5000$. We see that the IMQ KSD's Type I error rate is well controlled across the range of n and p values. We recommend AIKS when n is small but p is high.

We note that violation of higher-order identities, corresponding to the third moment or higher, cannot be detected by our curvature diagnostic because our diagnostic does not verify that any moment constraint of order 3 or higher is satisfied. Extending our methodology to include higher-order Bartlett identities is left to future work. The AIKS is supported by a theory of weak convergence and thus can be used for deciding convergence of a sequence of sample points to the target distribution.

A limitation of the ACD is that it requires the first- and second-order gradients of the likelihood and prior while AIKS requires only the first-order gradient. But ACD can be more computationally efficient than AIKS for large n . Let s denote the computational cost of simulating a single auxiliary

variable. Using N auxiliary variables, we approximate $d(\boldsymbol{\theta})$ for ACD and $\mathbf{u}(\boldsymbol{\theta})$ for AIKS at each posterior sample point. The time complexity of the approximation step is $\mathcal{O}(nsN)$. The threshold value for ACD is the $(1 - \alpha)$ quantile of the $\chi^2(r)$ of distribution, which is a known value. The time complexity for ACD is $\mathcal{O}(n)$ given the approximations. In contrast, the asymptotic distribution of AIKS has no closed-form expression. For each test sample path, it requires a bootstrap approximation to the threshold value, which costs $\mathcal{O}(bn^2)$ where b is the bootstrap sample size. Thus the total computational complexity for AIKS is $\mathcal{O}((b + 1)n^2)$ given the score approximations.

7. Applications

Here we apply our methods to both asymptotically exact and asymptotically inexact algorithms in the context of three general classes of models with intractable normalizing functions: (1) the Ising model, (2) a social network model, and (3) a Conway–Maxwell–Poisson regression model. Effective sample size (ESS), which is one of the most widely used MCMC diagnostics, is helpful for asymptotically exact methods since all chains converge to the target distribution. However, for asymptotically inexact methods, ESS is inadequate since an algorithm that mixes better might yield a poorer approximation to the target distribution. To illustrate the usefulness of our approaches for asymptotically inexact methods, we compare the approximate curvature diagnostic (ACD) and the approximate multiquadric kernel Stein discrepancy (AIKS) with ESS. We also study the computational complexity and the variability of our diagnostics. The code for our diagnostics is implemented in R (Ihaka and Gentleman, 1996) and C++, using the Rcpp and RcppArmadillo packages (Edelbuettel and Francois, 2011). The calculation of ESS follows Kass et al. (1998) and Robert and Casella (2004). All code was run on dual 10-core Xeon E5-2680 processors on the Penn State high-performance computing cluster. The source code may be found in the following repository (<https://github.com/bokgyeong/Diagnostics>).

7.1 The Ising Model

The Ising model (Lenz, 1920; Ising, 1925) is one of the most famous and important models from statistical physics and provides an approach for modeling binary images. For an $r \times s$ lattice $\mathbf{x} = \{x_{i,j}\}$ with binary values $x_{i,j} \in \{-1, 1\}$, where i, j denotes the row and column, the Ising model with a parameter $\theta \in (0, 1)$ has likelihood

$$L(\theta | \mathbf{x}) = \frac{1}{c(\theta)} \exp\{\theta S(\mathbf{x})\},$$

where

$$S(\mathbf{x}) = \sum_{i=1}^r \sum_{j=1}^{s-1} x_{i,j} x_{i,j+1} + \sum_{i=1}^{r-1} \sum_{j=1}^s x_{i,j} x_{i+1,j},$$

is the sum over all possible products of neighboring elements and imposes spatial dependence. A larger value for the dependence parameter θ produces stronger interactions among neighboring observations. Calculation of the normalizing function $c(\theta)$ requires summation over all 2^{rs} possible outcomes for the model, which is computationally infeasible even for lattices of moderate size. We carried out our simulation using perfect sampling (Propp and Wilson, 1996) on a 30×30 lattice with the parameter setting $\theta = 0.2$, which represents moderate dependence.

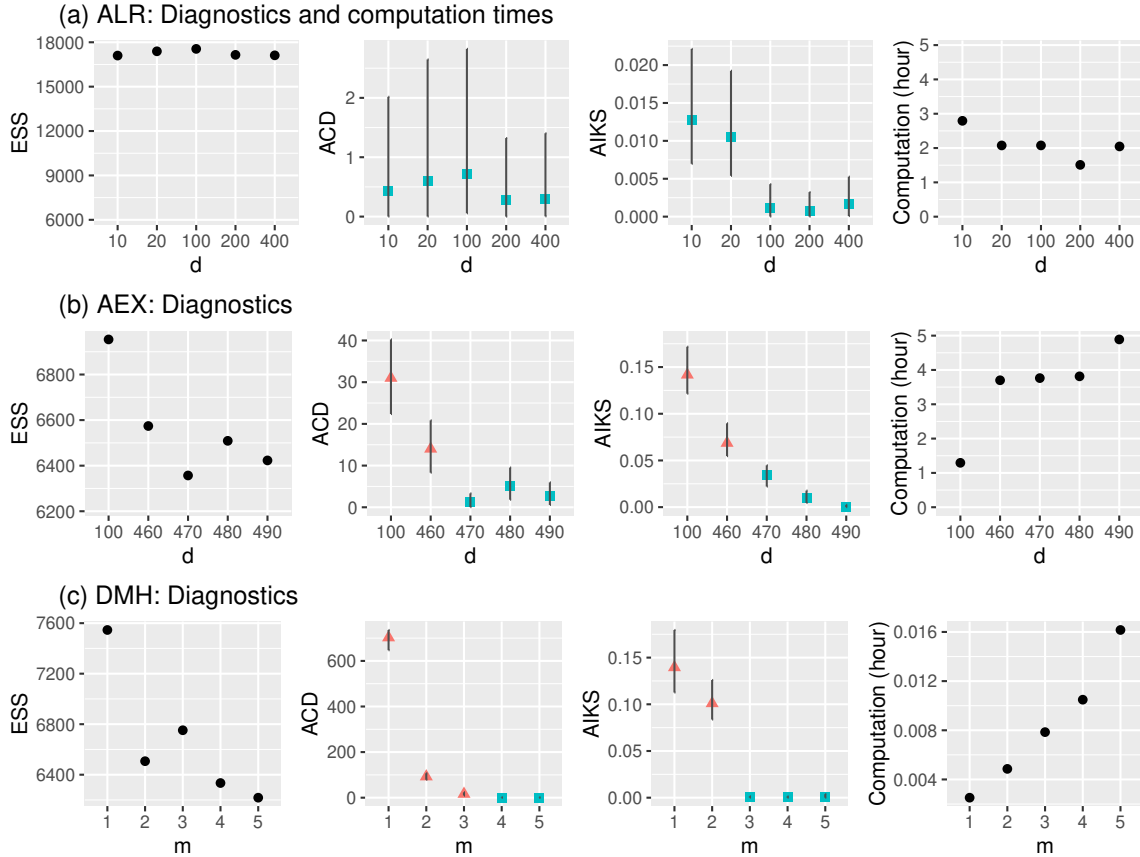


Figure 1: Results for the Ising model. The ESS, ACD, AIKS, and computation time for samples generated from (a) ALR with different numbers d of particles, (b) AEX with different numbers d of particles, and (c) DMH with different numbers m of inner updates. The triangle/square and vertical line show the empirical mean and 95% uncertainty interval, respectively, of 30 replications of the diagnostic. The red triangle and blue square indicate poor sample quality and good sample quality, respectively.

For this example we consider ALR, AEX, and DMH algorithms described in Section 2. The ALR and AEX algorithms are asymptotically exact. The DMH algorithm is asymptotically inexact but was found to be very efficient in terms of effective sample size per unit time and is applicable for doubly intractable problems with higher parameter dimension. To find appropriate values for the tuning parameters of the algorithms, we generate multiple chains from each algorithm with different choices of tuning parameters. The performance of ALR and AEX depends heavily on the set of particles. A sufficiently large number d of particles is required to cover the important region of the parameter space. Selecting a suitable value of d is difficult because said value varies across models or across datasets for a given model. The choices for the other tuning parameters are rather well studied and good heuristics exist. We implement ALR and AEX with different numbers d of particles. We choose values for the other tuning parameters according to Park and Haran (2018). To

Table 2: Summary statistics of posterior samples and a gold standard in the Ising model for a simulated binary lattice.

Algorithm	d or m	Median	LTP	RTP
DMH	1	0.19	0.09	0.09
DMH	4	0.19	0.05	0.05
ALR	10	0.19	0.05	0.05
AEX	470	0.19	0.05	0.04
Gold standard		0.19	0.05	0.05

SD, standard deviation; LTP, left-tail probability; RTP, right-tail probability.

obtain the particles, we used fractional DMH, which is DMH with a larger acceptance probability, and a short run of DMH with a single cycle of (inner) Gibbs updates for AEX and ALR, respectively. We implement DMH with different numbers m of (inner) Gibbs updates. We use a uniform prior with range $[0, 1]$ for θ .

All algorithms were run for $n = 100,000$ iterations, and all 100,000 posterior samples were used for testing. We used parallel computing to obtain 30 replications of each diagnostic for each sample path. We assess the samples using the empirical mean of 30 replications. The threshold value of ACD is the 0.99 quantile of $\chi^2(1)$, which is 6.63. We estimated the 0.99 quantile of the asymptotic distribution of AIKS via the bootstrap method. ACD computation took approximately 3 seconds and AIKS computation took approximately 27 seconds for each test sample path.

Figure 1 shows the diagnostic values and computation time for each sample path. For ALR, Figure 1 (a) shows that all diagnostics provide the same conclusion: ALR is not so sensitive to the number d of particles in this setting. For AEX and DMH, ACD and AIKS provide considerably different conclusions from ESS. In Figure 1 (b) we see that, for AEX, ESS is maximized (the best) at $d = 100$ and generally decreases (worsens) as d increases. On the other hand, ACD and AIKS have their largest (worst) values at $d = 100$ and indicate that d should be at least 470. Likewise, for DMH, Figure 1 (c) shows that ESS is maximized (the best) at $m = 1$ and generally decreases (worsens) as m increases, while ACD and AIKS take their largest (worst) values at $m = 1$ and decrease (improve) as m increases. ACD recommends m of 4 or more, and AIKS recommends m of 3 or more. The computation time is measured using a single run for each case. We observe that DMH with $m = 4$ provides a good approximation of the target distribution while being extremely computationally efficient.

We use the sample generated from the exchange algorithm (Murray et al., 2006), which is asymptotically exact and has no tuning parameter, as the gold standard. Table 2 presents summary statistics for some of the posterior samples and the gold standard. Cutoff values for the left- and right-tail probabilities are the lower 5% and the upper 5% percentiles of the gold standard. Both ACD and AIKS suggest that ALR with $d = 10$, AEX with $d = 470$, and DMH with $m = 4$ provide high quality samples. It is observed that they provide similar values of the summary statistics to the gold standard. On the other hand, the ESS-recommended DMH sample with $m = 1$ does not match the gold standard. It provides higher tail probabilities compared to the gold standard. This shows that ACD and AIKS perform well in assessing how close samples are to the exact posterior. Not




Figure 2: Results for the ERGM. The minimum ESS, ACD, AIKS and computation time for samples generated from DMH with different numbers m of inner updates. The triangle/square and vertical line show the empirical mean and 95% uncertainty interval, respectively, of 30 replications of the diagnostic. The red triangle and blue square indicate poor sample quality and good sample quality, respectively.

surprisingly, ESS is inadequate as a tool for this purpose. In summary, our approaches provide good guidance on how to assess the quality of samples for both the asymptotically exact and asymptotically inexact algorithms. In addition, our diagnostics help one to select the best algorithm and its tuning parameter in terms of statistical efficiency.

7.2 A Social Network Model

An exponential family random graph model (ERGM) (Robins et al., 2007; Hunter et al., 2008) is a statistical model for analysing network data. Consider an undirected ERGM with n vertices. Relationships among the vertices can be represented as an $n \times n$ adjacency matrix \mathbf{x} as follows: for all $i \neq j$, $x_{i,j} = 1$ if the i th and j th vertices are connected, and $x_{i,j} = 0$ otherwise. And $x_{i,i}$ is 0 for all $i \in \{1, \dots, n\}$, i.e., there are no loops. The number of possible network configurations is $2^{n(n-1)/2}$ and summation over those configurations is required to calculate the normalizing function of the model. Thus computing $c(\boldsymbol{\theta})$ is infeasible unless n is quite small.

For an undirected graph with n vertices, the ERGM likelihood is given by

$$L(\boldsymbol{\theta} \mid \mathbf{x}) = \frac{1}{c(\boldsymbol{\theta})} \exp\{\theta_1 S_1(\mathbf{x}) + \theta_2 S_2(\mathbf{x})\},$$

$$S_1(\mathbf{x}) = \sum_{i=1}^n \binom{x_{i+}}{1},$$

$$S_2(\mathbf{x}) = e^\tau \sum_{k=1}^{n-2} \left\{ 1 - (1 - e^{-\tau})^k \right\} \text{EP}_k(\mathbf{x}).$$

The sufficient statistics $S_1(\mathbf{x})$ and $S_2(\mathbf{x})$ are the number of edges in the graph and the geometrically weighted edge-wise shared partnership (GWESP) statistic (Hunter and Handcock, 2006; Hunter, 2007), respectively. x_{i+} denotes the sum of the i th row of the adjacency matrix and $\text{EP}_k(\mathbf{x})$ denotes the number of edges between two vertices that share exactly k neighbors. It is assumed that τ is fixed at a value of 0.25. We simulated a network with 30 actors using 10,000 cycles of Gibbs updates, where the true parameter was $\boldsymbol{\theta} = (\theta_1, \theta_2)^\top = (-0.96, 0.04)^\top$.

For this example we consider the DMH algorithm, which is asymptotically inexact. We explained the DMH algorithm in Section 7.1. We implement DMH with different numbers m of Gibbs updates. We use uniform priors with range $[-5.00, 2.27] \times [-1.57, 2.32]$ for $(\theta_1, \theta_2)^\top$, which are centered around the maximum pseudo-likelihood estimates (MPLE) and have widths of 12 standard deviations.

All algorithms were run for $n = 300,000$ iterations, and all 300,000 posterior samples were used for testing. We used parallel computing to obtain 30 replications of each diagnostic for each sample path. We assess the samples using the empirical mean of 30 replications. The threshold value of ACD is the 0.99 quantile of $\chi^2(3)$, which is 11.34. ACD computation took approximately 54 seconds and AIKS computation took approximately 38 minutes for each test sample path.

Figure 2 shows the minimum effective sample size—this is the conservative way of using ESS when there are multiple ESSs due to the multivariate posterior distribution—, ACD, AIKS, and computation time for the DMH sample, for a sequence of m values. The minimum ESS is maximized (the best) at $m = 4$. ACD and AIKS generally decrease (become better) as m increases. The ACD implies that m should be at least 4, and AIKS suggests that m should be at least 3. We see that $m = 4$ yields a good approximation according to both diagnostics, at a modest computing cost. We treat a run from DMH with 20 cycles of Gibbs updates as the gold standard. We observe that all ESS-, ACD-, and AIKS-selected samples provide almost the same values of the summary statistics as the gold standard. In this case, all of the diagnostics generally agree.

7.3 A Conway–Maxwell–Poisson Regression Model

The Conway–Maxwell–Poisson (COMP) distribution (Conway and Maxwell, 1962) is a two-parameter generalization of the Poisson distribution for modeling count data that are characterized by under-dispersion (variance less than the mean) or over-dispersion (variance greater than the mean). For a $\text{COMP}(\lambda, \nu)$ variable Y , the probability mass function is given by

$$P(Y = y) = \frac{1}{c(\lambda, \nu)} \frac{\lambda^y}{(y!)^\nu},$$

where $\lambda > 0$ is a generalization of the Poisson rate parameter, $\nu \geq 0$ denotes the dispersion parameter, and $c(\lambda, \nu) = \sum_{z=0}^{\infty} \lambda^z / (z!)^\nu$ is a normalizing function. The COMP distribution accommodates under- ($\nu > 1$), equi- ($\nu = 1$), or over-dispersion ($0 \leq \nu < 1$). The COMP distribution forms a continuous bridge between the Poisson ($\nu = 1$), geometric ($\nu = 0$ and $\lambda < 1$), and Bernoulli ($\nu = \infty$ and success probability $\lambda/(1 + \lambda)$) distributions. Guikema and Goffelt (2008) proposed a reparameterization, substituting $\eta = \lambda^{1/\nu}$ to approximate the center of the COMP distribution. For a count variable Y that follows the $\text{COMP}_\eta(\eta, \nu)$ distribution, the probability mass function is

$$P(Y = y) = \frac{1}{c_\eta(\eta, \nu)} \left(\frac{\eta^y}{y!} \right)^\nu,$$

where $c_\eta(\eta, \nu) = \sum_{z=0}^{\infty} (\eta^z / z!)^\nu$ is the normalizing function. Huang (2017) and Ribeiro et al. (2020) have reparameterized the COMP distribution as a function of the mean. Under any parameterization, however, the COMP normalizing function is an infinite sum, making the function intractable.

For count response variables Y_i and corresponding explanatory variables $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p-1})^\top$ for $i = 1, \dots, n$, a COMP regression model with log link function for η is given by $Y_i \sim$




Figure 3: Results for the COMP regression model. The minimum ESS, ACD, AIKS and computation time for samples generated from NormTrunc with different levels k of truncation. The triangle/square and vertical line show the empirical mean and 95% uncertainty interval, respectively, of 30 replications of the diagnostic. The red triangle and blue square indicate poor sample quality and good sample quality, respectively.

COMP $_{\eta}(\eta_i, \nu)$ with $\log(\eta_i) = \beta_0 + x_{i,1}\beta_1 + \dots + x_{i,p-1}\beta_{p-1}$ ($i = 1, \dots, n$), where $\beta = (\beta_0, \dots, \beta_{p-1})^T$ are regression coefficients. We study the takeover bids dataset (Cameron and Johansson, 1997), which comprises the number of bids received by 126 US firms that were targets of tender offers during the period 1978–1985, along with 9 explanatory variables (see Sáez-Castillo and Conde-Sánchez, 2013, for details). The dataset can be obtained from `mpcomp` package (Fung et al., 2019). It is assumed that ν is fixed at a value of 1.754 (Huang, 2017).

For this example we consider a simple and widely used approach, namely, the NormTrunc algorithm. Specifically, we approximate the normalizing function by truncating the infinite sum to a truncation level k such that $c_{\eta}(\eta, \nu) \approx \sum_{z=0}^k (\eta^z / z!)^{\nu}$ and use the approximation for each Metropolis–Hastings accept-reject ratio. The NormTrunc algorithm is asymptotically inexact. In order to determine a suitable value of k , we implement NormTrunc with different levels k of truncation.

All algorithms were run for $n = 300,000$ iterations, and all 300,000 posterior samples were used for testing. We used parallel computing to obtain 30 replications of each diagnostic for each sample path. We assess the samples using the empirical mean of 30 replications. The threshold value of ACD is the 0.99 quantile of $\chi^2(55)$, which is 82.29. ACD computation took approximately 5.8 hours and AIKS computation took approximately 6.3 hours for each test sample path.

Table 3: Summary statistics of posterior samples and a gold standard for β_0 in the COMP regression model.

Algorithm	k	Median	LTP	RTP
NormTrunc	3	0.45	0.01	0.47
NormTrunc	10	0.31	0.05	0.05
Gold standard		0.31	0.05	0.05

SD, standard deviation; LTP, left-tail probability; RTP, right-tail probability.

Figure 3 shows minimum effective sample size, ACD, AIKS, and computing time for the NormTrunc sample for a range of k values. The minimum ESS is maximized (the best) at $k = 3$. On the other hand, both ACD and AIKS recommend k of 10 or more. We see that the uncertainty of our diagnostics is very small since we generate auxiliary variables exactly from the model using the rejection sampler. It appears that $k = 10$ is sufficient.

We treat as the gold standard a run from the exchange algorithm (Murray et al., 2006). Table 3 presents summary statistics for some of the posterior samples and a gold standard. The cutoff values for left- and right-tail probabilities were chosen as in the previous example. Both ACD and AIKS suggest that NormTrunc with $k = 10$ provides high quality samples. It is observed that the NormTrunc sample with $k = 10$ provides the same values of the summary statistics as the gold standard. On the other hand, the ESS-recommended NormTrunc sample with $k = 3$ does not match the gold standard. It provides a large median and a high right-tail probability compared to the gold standard. In summary, our approaches help tune algorithms. In particular, for asymptotically inexact algorithms, our methods can guide users to appropriately choose their tuning parameters and help provide reliable inference.

8. Discussion

In this article we proposed new methods that provide guidance for tuning algorithms and give some measure of sample quality for a wide range of Monte Carlo algorithms, including a particularly challenging class of algorithms: asymptotically inexact algorithms for distributions with intractable normalizing functions. We describe three methods. CD applies broadly to most any likelihood-based context where misspecification is of concern while ACD and AIKS are specifically targeted at likelihoods with intractable normalizing functions. Our study mainly focuses on intractable normalizing function problems and shows that our methods can assess the quality of samples and provide good guidance for tuning algorithms. We have studied applications of ACD and AIKS to several asymptotically exact and inexact algorithms in the context of challenging simulated and real data examples. This shows that our methods provide useful results not only for asymptotically exact algorithms, for which some other methods may be useful, but also for asymptotically inexact algorithms, for which we are not aware of other methods.

There are of course simple and reasonable heuristics one can use for diagnostic purposes, such as increasing m until DMH approximations stabilize. However, our diagnostics go beyond simple approaches, allowing one to compare sample quality across different algorithms, including comparing asymptotically exact algorithms with asymptotically inexact algorithms.

We note that ACD and AIKS could be slightly different in their conclusions, as our examples illustrate, but the difference is quite small and conclusions are nearly the same overall. The difference between ACD and AIKS can stem from the fact that they consider different sets of test functions to quantify the difference between a sample mean and a target expectation. ACD considers two functions whose target expectations are identical and measures the difference between sample means of the two functions. AIKS considers a set of functions whose target expectations are zero and finds the maximum error between the sample mean and the target expectation. Unless the score function of the target density is the origin, their sets of test functions are mutually exclusive, which is the case of all the examples in this article.

ACD and AIKS are obtained by replacing some functions that have to be evaluated for CD and IMQ KSD with their Monte Carlo approximations. This allows for the diagnostics to be avail-

able for doubly intractable posterior distributions. The computing time of ACD and AIKS mainly accounts for the Monte Carlo approximation, which might be computationally expensive for high-dimensional datasets. However, the approximation step is embarrassingly parallel and the computational cost can be reduced drastically via the self-normalized importance sampling. We want to note that the power of our diagnostics heavily depends on the accuracy of the gradient estimations. The approximation can be less reliable in cases where the data distribution has high variance. This can lead to high variance in gradient estimates and make the estimates unstable. In order to alleviate the impact of the approximation uncertainty to the power of the diagnostics, we compute 30 independent replications of each diagnostic and use its mean for testing. Large disparities between the posterior sample point $\theta^{(i)}$ and the chosen particles $\{\psi_1, \dots, \psi_M\}$ might yield poor SNIS estimates. When SNIS performs poorly, we suggest sampling auxiliary variables from $f(\cdot | \theta^{(i)})$ and obtaining Monte Carlo approximations to the gradients. An important caveat of ACD is that it can be misled if the first two moments match the target distribution but higher order moments do not. AIKS cannot be misled in such fashion. Extending ACD to higher order moments may provide interesting avenues for future research.

Acknowledgments

The authors are grateful to Jaewoo Park for providing useful code and to Galin Jones and Geoff K. Nicholls for illuminating discussions. JH and MH were partially supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award Number R01GM123007.

Appendix A. Monte Carlo Approximations to Intractable Terms

A.1 Derivative of Log Normalizing Function

For a p -dimensional parameter vector θ , consider a posterior density $\pi(\theta | \mathbf{x})$ whose likelihood function is $L(\theta | \mathbf{x}) = h(\mathbf{x} | \theta)/c(\theta)$ and prior density is $p(\theta)$. The k th entry of $\nabla_{\theta} \log c(\theta | \mathbf{x})$, the score function of the normalizing function, can be written as

$$\begin{aligned}
 \frac{\partial \log c(\theta)}{\partial \theta_k} &= \frac{1}{c(\theta)} \frac{\partial c(\theta)}{\partial \theta_k} \\
 &= \frac{1}{c(\theta)} \frac{\partial}{\partial \theta_k} \int_{\mathcal{X}} h(\mathbf{x} | \theta) d\mathbf{x} \\
 &= \frac{1}{c(\theta)} \int_{\mathcal{X}} \frac{\partial h(\mathbf{x} | \theta)}{\partial \theta_k} d\mathbf{x} \tag{6} \\
 &= \frac{1}{c(\theta)} \int_{\mathcal{X}} h(\mathbf{x} | \theta) \frac{\partial \log h(\mathbf{x} | \theta)}{\partial \theta_k} d\mathbf{x} \\
 &= \int_{\mathcal{X}} \frac{\partial \log h(\mathbf{x} | \theta)}{\partial \theta_k} f(\mathbf{x} | \theta) d\mathbf{x} \\
 &= E_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_k} \right\}. \tag{7}
 \end{aligned}$$

Equation (6) follows from the dominated convergence theorem. Under the assumptions that the score function exists and the score function and the normalizing function are bounded, we have exchanged the derivative with the integral.

A.2 Second Derivative of Log Normalizing Function

The (k, l) th entry of $\nabla_{\theta}^2 \log c(\theta | \mathbf{x})$, the Hessian matrix of the normalizing function, is given by

$$\begin{aligned} \frac{\partial^2 \log c(\theta)}{\partial \theta_k \partial \theta_l} &= \frac{\partial}{\partial \theta_k} \left\{ \frac{1}{c(\theta)} \frac{\partial c(\theta)}{\partial \theta_l} \right\} \\ &= \frac{1}{c(\theta)} \frac{\partial^2 c(\theta)}{\partial \theta_k \partial \theta_l} - \left\{ \frac{1}{c(\theta)} \frac{\partial c(\theta)}{\partial \theta_k} \right\} \left\{ \frac{1}{c(\theta)} \frac{\partial c(\theta)}{\partial \theta_l} \right\} \\ &= \frac{1}{c(\theta)} \frac{\partial^2 c(\theta)}{\partial \theta_k \partial \theta_l} - \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_k} \right\} \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_l} \right\}. \end{aligned} \quad (8)$$

Now consider the first term on the right hand side of equation (8). It can be written as

$$\begin{aligned} \frac{1}{c(\theta)} \frac{\partial^2 c(\theta)}{\partial \theta_k \partial \theta_l} &= \frac{1}{c(\theta)} \frac{\partial^2}{\partial \theta_k \partial \theta_l} \int_{\mathbf{x}} h(\mathbf{x} | \theta) d\mathbf{x} \\ &= \frac{1}{c(\theta)} \int_{\mathbf{x}} \frac{\partial}{\partial \theta_k} \left\{ \frac{\partial \log h(\mathbf{x} | \theta)}{\partial \theta_l} h(\mathbf{x} | \theta) \right\} d\mathbf{x} \end{aligned} \quad (9)$$

$$\begin{aligned} &= \int_{\mathbf{x}} \frac{\partial^2 \log h(\mathbf{x} | \theta)}{\partial \theta_k \partial \theta_l} f(\mathbf{x} | \theta) d\mathbf{x} + \int_{\mathbf{x}} \frac{\partial \log h(\mathbf{x} | \theta)}{\partial \theta_k} \frac{\partial \log h(\mathbf{x} | \theta)}{\partial \theta_l} f(\mathbf{x} | \theta) d\mathbf{x} \\ &= \mathbb{E}_f \left\{ \frac{\partial^2 \log h(\mathbf{X} | \theta)}{\partial \theta_k \partial \theta_l} \right\} + \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_k} \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_l} \right\}. \end{aligned} \quad (10)$$

In the equation (9) we have exchanged the derivative with the integral, owing to the dominated convergence theorem. Combining (8) and (10), the entry of $\nabla_{\theta}^2 \log c(\theta | \mathbf{x})$ is

$$\begin{aligned} \frac{\partial^2 \log c(\theta)}{\partial \theta_k \partial \theta_l} &= \mathbb{E}_f \left\{ \frac{\partial^2 \log h(\mathbf{X} | \theta)}{\partial \theta_k \partial \theta_l} \right\} + \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_k} \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_l} \right\} \\ &\quad - \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_k} \right\} \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X} | \theta)}{\partial \theta_l} \right\}. \end{aligned} \quad (11)$$

A.3 Monte Carlo Approximations

We approximate the k th entry of the vector $\nabla_{\theta} \log c(\theta)$ using a sample generated from the model distribution:

$$\frac{\partial \log c(\theta)}{\partial \theta_k} \approx \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)} | \theta)}{\partial \theta_k},$$

where $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}$ are auxiliary variables generated exactly from $f(\cdot | \theta)$ or generated by a Monte Carlo algorithm having $f(\cdot | \theta)$ as its target distribution. In a similar fashion, the approximation of the (k, l) th entry of $\{\nabla_{\theta} \log c(\theta)\} \{\nabla_{\theta} \log c(\theta)\}^{\top}$ is calculated as

$$\frac{\partial \log c(\theta)}{\partial \theta_k} \frac{\partial \log c(\theta)}{\partial \theta_l} \approx \left\{ \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)} | \theta)}{\partial \theta_k} \right\} \left\{ \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)} | \theta)}{\partial \theta_l} \right\}, \quad k, l = 1, \dots, p.$$

Lastly we approximate (k, l) th entry of $\nabla_{\boldsymbol{\theta}}^2 \log c(\boldsymbol{\theta})$ by

$$\begin{aligned} \frac{\partial^2 \log c(\boldsymbol{\theta})}{\partial \theta_k \partial \theta_l} &\approx \frac{1}{N} \sum_{j=1}^N \frac{\partial^2 \log h(\mathbf{y}^{(j)} | \boldsymbol{\theta})}{\partial \theta_k \partial \theta_l} + \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)} | \boldsymbol{\theta})}{\partial \theta_k} \frac{\partial \log h(\mathbf{y}^{(j)} | \boldsymbol{\theta})}{\partial \theta_l} \\ &\quad - \left\{ \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)} | \boldsymbol{\theta})}{\partial \theta_k} \right\} \left\{ \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)} | \boldsymbol{\theta})}{\partial \theta_l} \right\}. \end{aligned}$$

Appendix B. Proof of Theorem 5

Let $D(\boldsymbol{\theta}) = J(\boldsymbol{\theta}) + H(\boldsymbol{\theta})$ and $D_{k,l}(\boldsymbol{\theta})$ be the (k, l) entry of $D(\boldsymbol{\theta})$. Let $\hat{D}_{k,l}(\boldsymbol{\theta})$ be the approximation, computed using auxiliary variables $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}$ generated from the model distribution, to $D_{k,l}(\boldsymbol{\theta})$. Suppose we have an i.i.d. sample $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ generated from $\pi(\cdot | \boldsymbol{x})$. For all $k, l, m, s \in \{1, \dots, p\}$, the difference between an entry of $V = \mathbb{E}_{\pi} \{ \mathbf{d}(\boldsymbol{\theta}) \mathbf{d}(\boldsymbol{\theta})^\top \}$ and its two-state approximation is

$$\begin{aligned} &\left| \mathbb{E}_{\pi} \{ D_{k,l}(\boldsymbol{\theta}) D_{m,s}(\boldsymbol{\theta}) \} - \frac{1}{n} \sum_{i=1}^n \hat{D}_{k,l}(\boldsymbol{\theta}^{(i)}) \hat{D}_{m,s}(\boldsymbol{\theta}^{(i)}) \right| \\ &\leq \left| \frac{1}{n} \sum_{i=1}^n D_{k,l}(\boldsymbol{\theta}^{(i)}) D_{m,s}(\boldsymbol{\theta}^{(i)}) - \mathbb{E}_{\pi} \{ D_{k,l}(\boldsymbol{\theta}) D_{m,s}(\boldsymbol{\theta}) \} \right| \\ &\quad + \left| \frac{1}{n} \sum_{i=1}^n \left\{ \hat{D}_{k,l}(\boldsymbol{\theta}^{(i)}) \hat{D}_{m,s}(\boldsymbol{\theta}^{(i)}) - D_{k,l}(\boldsymbol{\theta}^{(i)}) D_{m,s}(\boldsymbol{\theta}^{(i)}) \right\} \right| \\ &\leq \delta(n) + \frac{1}{n} \sum_{i=1}^n \left| \hat{D}_{k,l}(\boldsymbol{\theta}^{(i)}) \hat{D}_{m,s}(\boldsymbol{\theta}^{(i)}) - D_{k,l}(\boldsymbol{\theta}^{(i)}) D_{m,s}(\boldsymbol{\theta}^{(i)}) \right|, \end{aligned} \quad (12)$$

where $\delta(n) = \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{d}(\boldsymbol{\theta}^{(i)}) \mathbf{d}(\boldsymbol{\theta}^{(i)})^\top - \mathbb{E}_{\pi} \{ \mathbf{d}(\boldsymbol{\theta}) \mathbf{d}(\boldsymbol{\theta})^\top \} \right\|_{\max} = \mathcal{O}(n^{-1/2})$ from the ergodic theorem. Now we consider the second term on the right hand side of inequality (12). Let $u_k(\boldsymbol{\theta})$ be the k th entry of the score function $\mathbf{u}(\boldsymbol{\theta})$ of the posterior distribution and $\hat{u}_k(\boldsymbol{\theta})$ be its approximation computed using the auxiliary variables. Let $H_{k,l}(\boldsymbol{\theta})$ be the (k, l) th entry of the hessian matrix $H(\boldsymbol{\theta})$ of the posterior distribution and $\hat{H}_{k,l}(\boldsymbol{\theta})$ be its approximation computed using the auxiliary variables. We can bound the difference between $\hat{D}_{k,l}(\boldsymbol{\theta}) \hat{D}_{m,s}(\boldsymbol{\theta})$ and $D_{k,l}(\boldsymbol{\theta}) D_{m,s}(\boldsymbol{\theta})$ as follows. For all $k, l, m, s \in \{1, \dots, p\}$, we have

$$\begin{aligned} \left| \hat{D}_{k,l}(\boldsymbol{\theta}) \hat{D}_{m,s}(\boldsymbol{\theta}) - D_{k,l}(\boldsymbol{\theta}) D_{m,s}(\boldsymbol{\theta}) \right| &\leq \left| \hat{u}_k(\boldsymbol{\theta}) \hat{u}_l(\boldsymbol{\theta}) \hat{u}_m(\boldsymbol{\theta}) \hat{u}_s(\boldsymbol{\theta}) - u_k(\boldsymbol{\theta}) u_l(\boldsymbol{\theta}) u_m(\boldsymbol{\theta}) u_s(\boldsymbol{\theta}) \right| \\ &\quad + \left| \hat{u}_k(\boldsymbol{\theta}) \hat{u}_l(\boldsymbol{\theta}) \hat{H}_{m,s}(\boldsymbol{\theta}) - u_k(\boldsymbol{\theta}) u_l(\boldsymbol{\theta}) H_{m,s}(\boldsymbol{\theta}) \right| \\ &\quad + \left| \hat{u}_m(\boldsymbol{\theta}) \hat{u}_s(\boldsymbol{\theta}) \hat{H}_{k,l}(\boldsymbol{\theta}) - u_m(\boldsymbol{\theta}) u_s(\boldsymbol{\theta}) H_{k,l}(\boldsymbol{\theta}) \right| \\ &\quad + \left| \hat{H}_{k,l}(\boldsymbol{\theta}) \hat{H}_{m,s}(\boldsymbol{\theta}) - H_{k,l}(\boldsymbol{\theta}) H_{m,s}(\boldsymbol{\theta}) \right|. \end{aligned} \quad (13)$$

Let

$$\begin{aligned}
 \epsilon_1(N) &= \max_k \left| \frac{1}{N} \sum_{j=1}^N \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_k} - \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_k} \right\} \right|, \\
 \epsilon_2(N) &= \max_{k,l} \left| \frac{1}{N} \sum_{j=1}^N \frac{\partial^2 \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_k \partial \theta_l} - \mathbb{E}_f \left\{ \frac{\partial^2 \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_k \partial \theta_l} \right\} \right|, \\
 \epsilon_3(N) &= \max_l \left| \frac{1}{N} \sum_{j=1}^N \frac{\partial^3 \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_k \partial \theta_l \partial \theta_m} - \mathbb{E}_f \left\{ \frac{\partial^3 \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_k \partial \theta_l \partial \theta_m} \right\} \right|, \\
 \epsilon_4(N) &= \max_{k,l} \left| \frac{1}{N} \sum_{j=1}^N \left\{ \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_k} \right\} \left\{ \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_l} \right\} - \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_k} \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_l} \right\} \right|, \\
 \epsilon_5(N) &= \max_{k,l,m} \left| \frac{1}{N} \sum_{j=1}^N \left\{ \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_k} \right\} \left\{ \frac{\partial^2 \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_l \partial \theta_m} \right\} - \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_k} \frac{\partial^2 \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_l \partial \theta_m} \right\} \right|, \\
 \epsilon_6(N) &= \max_{k,l,m} \left| \frac{1}{N} \sum_{j=1}^N \left\{ \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_k} \right\} \left\{ \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_l} \right\} \left\{ \frac{\partial \log h(\mathbf{y}^{(j)}|\boldsymbol{\theta})}{\partial \theta_m} \right\} \right. \\
 &\quad \left. - \mathbb{E}_f \left\{ \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_k} \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_l} \frac{\partial \log h(\mathbf{X}|\boldsymbol{\theta})}{\partial \theta_m} \right\} \right|,
 \end{aligned}$$

and $\epsilon(N) = \max \{\epsilon_1(N), \epsilon_2(N), \epsilon_3(N), \epsilon_4(N), \epsilon_5(N), \epsilon_6(N)\} = \mathcal{O}(N^{-1/2})$ from the ergodic theorem. Now consider the first term on the right hand side of inequality (13). For all $k, l, m, s \in \{1, \dots, p\}$, we have

$$|\hat{u}_k(\boldsymbol{\theta})\hat{u}_l(\boldsymbol{\theta})\hat{u}_m(\boldsymbol{\theta})\hat{u}_s(\boldsymbol{\theta}) - u_k(\boldsymbol{\theta})u_l(\boldsymbol{\theta})u_m(\boldsymbol{\theta})u_s(\boldsymbol{\theta})| \leq B_1\epsilon(N),$$

where B_1 is some constant. The above inequality follows from Assumptions 1 and 2, and the following elementary inequalities:

$$\begin{aligned}
 |ab - \hat{a}\hat{b}| &= |ab - \hat{a}b + \hat{a}b - \hat{a}\hat{b}| = |(a - \hat{a})b + (b - \hat{b})\hat{a}| \\
 &\leq |a - \hat{a}||b| + |b - \hat{b}||\hat{a}|,
 \end{aligned} \tag{14}$$

$$\begin{aligned}
 |abc - \hat{a}\hat{b}\hat{c}| &= |abc - \hat{a}bc + \hat{a}bc - \hat{a}\hat{b}\hat{c}| = |(a - \hat{a})bc + (bc - \hat{b}\hat{c})\hat{a}| \\
 &\leq |a - \hat{a}||b||c| + |bc - \hat{b}\hat{c}||\hat{a}|,
 \end{aligned} \tag{15}$$

$$\begin{aligned}
 |abcd - \hat{a}\hat{b}\hat{c}\hat{d}| &= |abcd - \hat{a}bcd + \hat{a}bcd - \hat{a}\hat{b}\hat{c}\hat{d}| = |(a - \hat{a})bcd + (bcd - \hat{b}\hat{c}\hat{d})\hat{a}| \\
 &\leq |a - \hat{a}||b||c||d| + |bcd - \hat{b}\hat{c}\hat{d}||\hat{a}|.
 \end{aligned} \tag{16}$$

Now consider the remaining terms on the right hand side of inequality (13). In a similar fashion, for all $k, l, m, s \in \{1, \dots, p\}$, we have

$$\begin{aligned} \left| \hat{u}_k(\boldsymbol{\theta}) \hat{u}_l(\boldsymbol{\theta}) \hat{H}_{m,s}(\boldsymbol{\theta}) - u_k(\boldsymbol{\theta}) u_l(\boldsymbol{\theta}) H_{m,s}(\boldsymbol{\theta}) \right| &\leq B_2 \epsilon(N), \\ \left| \hat{u}_m(\boldsymbol{\theta}) \hat{u}_s(\boldsymbol{\theta}) \hat{H}_{k,l}(\boldsymbol{\theta}) - u_m(\boldsymbol{\theta}) u_s(\boldsymbol{\theta}) H_{k,l}(\boldsymbol{\theta}) \right| &\leq B_3 \epsilon(N), \\ \left| \hat{H}_{k,l}(\boldsymbol{\theta}) \hat{H}_{m,s}(\boldsymbol{\theta}) - H_{k,l}(\boldsymbol{\theta}) H_{m,s}(\boldsymbol{\theta}) \right| &\leq B_4 \epsilon(N), \end{aligned}$$

where B_2, B_3 , and B_4 are some constants. These inequalities follow from Assumptions 1 and 2 in the manuscript and inequalities (14) to (16). Therefore, the approximation error of the two-stage approximation to the asymptotic covariance matrix is

$$\|\hat{V}_{n,N} - V\|_{\max} \leq \mathcal{O}\left(n^{-1/2}\right) + \mathcal{O}\left(N^{-1/2}\right),$$

almost surely. We use $\|\cdot\|_{\max}$ to represent the max norm of vectors or matrices.

Appendix C. Proof of Theorem 7

Suppose we have a sample $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ generated from a polynomially ergodic Markov chain with $\pi(\boldsymbol{\theta} \mid \boldsymbol{x})$ as its stationary distribution. Let $D(\boldsymbol{\theta}) = J(\boldsymbol{\theta}) + H(\boldsymbol{\theta})$ and $\hat{D}(\boldsymbol{\theta})$ be the approximation, computed using auxiliary variables $\boldsymbol{y}^{(1)}, \dots, \boldsymbol{y}^{(N)}$ generated from the model distribution, to $D(\boldsymbol{\theta})$. Let $\bar{D} = \frac{1}{n} \sum_{i=1}^n D(\boldsymbol{\theta}^{(i)})$, $\hat{\bar{D}} = \frac{1}{n} \sum_{i=1}^n \hat{D}(\boldsymbol{\theta}^{(i)})$, $\bar{D}^j = \frac{1}{b_n} \sum_{i=(j-1)b_n+1}^{jb_n} D(\boldsymbol{\theta}^{(i)})$, and $\hat{\bar{D}}^j = \frac{1}{b_n} \sum_{i=(j-1)b_n+1}^{jb_n} \hat{D}(\boldsymbol{\theta}^{(i)})$ for $j = 1, \dots, a_n$, where a_n is the number of batches and b_n is the batch size. For all $k, l, m, s \in \{1, \dots, p\}$, the difference between an entry σ_r of Σ and its two-state batch means approximation is

$$\begin{aligned} &\left| \sigma_r - \frac{b_n}{a_n - 1} \sum_{j=1}^{a_n} (\hat{\bar{D}}_{k,l}^j - \hat{\bar{D}}_{k,l})(\hat{\bar{D}}_{m,s}^j - \hat{\bar{D}}_{m,s}) \right| \\ &\leq \left| \frac{b_n}{a_n - 1} \sum_{j=1}^{a_n} (\bar{D}_{k,l}^j - \bar{D}_{k,l})(\bar{D}_{m,s}^j - \bar{D}_{m,s}) - \sigma_r \right| \\ &\quad + \left| \frac{b_n}{a_n - 1} \sum_{j=1}^{a_n} \left\{ (\hat{\bar{D}}_{k,l}^j - \hat{\bar{D}}_{k,l})(\hat{\bar{D}}_{m,s}^j - \hat{\bar{D}}_{m,s}) - (\bar{D}_{k,l}^j - \bar{D}_{k,l})(\bar{D}_{m,s}^j - \bar{D}_{m,s}) \right\} \right| \\ &\leq \gamma(n) + \frac{b_n}{a_n - 1} \sum_{j=1}^{a_n} \left| (\hat{\bar{D}}_{k,l}^j - \hat{\bar{D}}_{k,l})(\hat{\bar{D}}_{m,s}^j - \hat{\bar{D}}_{m,s}) - (\bar{D}_{k,l}^j - \bar{D}_{k,l})(\bar{D}_{m,s}^j - \bar{D}_{m,s}) \right|, \quad (17) \end{aligned}$$

where $\gamma(n) = \|\hat{\Sigma}_n - \Sigma\|_{\max} = \mathcal{O}\left((\log n/b_n)^{1/2} n^{1/2-\lambda}\right) \rightarrow 0$ as $n \rightarrow \infty$ for some $\lambda > 0$ under some conditions (Theorem 2 in Vats et al., 2019). Now we consider the second term on the right hand side of inequality (17). For all $k, l, m, s \in \{1, \dots, p\}$, we have

$$\begin{aligned}
 & b_n \left| (\hat{D}_{k,l}^j - \hat{D}_{k,l}) (\hat{D}_{m,s}^j - \hat{D}_{m,s}) - (\bar{D}_{k,l}^j - \bar{D}_{k,l}) (\bar{D}_{m,s}^j - \bar{D}_{m,s}) \right| \\
 &= \left| b_n^{1/2} (\hat{D}_{k,l}^j - \hat{D}_{k,l}) b_n^{1/2} (\hat{D}_{m,s}^j - \hat{D}_{m,s}) - b_n^{1/2} (\bar{D}_{k,l}^j - \bar{D}_{k,l}) b_n^{1/2} (\bar{D}_{m,s}^j - \bar{D}_{m,s}) \right| \\
 &\leq \left| b_n^{1/2} (\bar{D}_{k,l}^j - \bar{D}_{k,l}) - b_n^{1/2} (\hat{D}_{k,l}^j - \hat{D}_{k,l}) \right| \left| b_n^{1/2} (\bar{D}_{m,s}^j - \bar{D}_{m,s}) \right| \\
 &\quad + \left| b_n^{1/2} (\bar{D}_{m,s}^j - \bar{D}_{m,s}) - b_n^{1/2} (\hat{D}_{m,s}^j - \hat{D}_{m,s}) \right| \left| b_n^{1/2} (\hat{D}_{k,l}^j - \hat{D}_{k,l}) \right| \tag{18} \\
 &\leq \left| b_n^{1/2} (\bar{D}_{k,l}^j - \hat{D}_{k,l}^j) - b_n^{1/2} (\bar{D}_{k,l} - \hat{D}_{k,l}) \right| \\
 &\quad \times \left| b_n^{1/2} \{ \bar{D}_{m,s}^j - \mathbb{E}_\pi [D_{m,s}(\boldsymbol{\theta})] \} - b_n^{1/2} \{ \bar{D}_{m,s} - \mathbb{E}_\pi [D_{m,s}(\boldsymbol{\theta})] \} \right| \\
 &\quad + \left| b_n^{1/2} (\bar{D}_{m,s}^j - \hat{D}_{m,s}^j) - b_n^{1/2} (\bar{D}_{m,s} - \hat{D}_{m,s}) \right| \\
 &\quad \times \left| b_n^{1/2} (\hat{D}_{k,l}^j - \bar{D}_{k,l}^j) - b_n^{1/2} (\hat{D}_{k,l} - \bar{D}_{k,l}) + b_n^{1/2} \{ \bar{D}_{k,l}^j - \mathbb{E}_\pi [D_{k,l}(\boldsymbol{\theta})] \} - b_n^{1/2} \{ \bar{D}_{k,l} - \mathbb{E}_\pi [D_{k,l}(\boldsymbol{\theta})] \} \right| \\
 &\leq \left\{ b_n^{1/2} |\bar{D}_{k,l}^j - \hat{D}_{k,l}^j| + b_n^{1/2} |\bar{D}_{k,l} - \hat{D}_{k,l}| \right\} \\
 &\quad \times \left\{ b_n^{1/2} |\bar{D}_{m,s}^j - \mathbb{E}_\pi [D_{m,s}(\boldsymbol{\theta})]| + b_n^{1/2} |\bar{D}_{m,s} - \mathbb{E}_\pi [D_{m,s}(\boldsymbol{\theta})]| \right\} \\
 &\quad + \left\{ b_n^{1/2} |\bar{D}_{m,s}^j - \hat{D}_{m,s}^j| + b_n^{1/2} |\bar{D}_{m,s} - \hat{D}_{m,s}| \right\} \\
 &\quad \times \left\{ b_n^{1/2} |\bar{D}_{k,l}^j - \hat{D}_{k,l}^j| + b_n^{1/2} |\bar{D}_{k,l} - \hat{D}_{k,l}| + b_n^{1/2} \left| \bar{D}_{k,l}^j - \mathbb{E}_\pi [D_{k,l}(\boldsymbol{\theta})] \right| + b_n^{1/2} \left| \bar{D}_{k,l} - \mathbb{E}_\pi [D_{k,l}(\boldsymbol{\theta})] \right| \right\} \\
 &\leq \left\{ b_n^{1/2} \mathcal{O}(N^{-1/2}) + b_n^{1/2} \mathcal{O}(N^{-1/2}) \right\} \left\{ b_n^{1/2} \mathcal{O}(b_n^{-1/2}) + b_n^{1/2} \mathcal{O}(n^{-1/2}) \right\} \\
 &\quad + \left\{ b_n^{1/2} \mathcal{O}(N^{-1/2}) + b_n^{1/2} \mathcal{O}(N^{-1/2}) \right\} \\
 &\quad \times \left\{ b_n^{1/2} \mathcal{O}(N^{-1/2}) + b_n^{1/2} \mathcal{O}(N^{-1/2}) + b_n^{1/2} \mathcal{O}(b_n^{-1/2}) + b_n^{1/2} \mathcal{O}(n^{-1/2}) \right\} \\
 &\leq \mathcal{O}\left((b_n/N)^{1/2}\right) + \mathcal{O}\left((b_n/n)^{1/2}(b_n/N)^{1/2}\right) + \mathcal{O}(b_n/N) = \mathcal{O}\left((b_n/N)^{1/2}\right),
 \end{aligned}$$

where inequality (18) follows from inequality (14). Therefore, the approximation error of the two-stage batch means approximation to the asymptotic covariance matrix is

$$\|\hat{\Sigma}_{n,N} - \Sigma\|_{\max} \leq \mathcal{O}\left((\log n/b_n)^{1/2} n^{1/2-\lambda}\right) + \mathcal{O}\left((b_n/N)^{1/2}\right),$$

almost surely for some $\lambda > 0$.

Appendix D. Proof of Theorem 10

Consider a target distribution P and a weighted sample $Q_n = \sum_{i=1}^n q_n(\boldsymbol{\theta}^{(i)}) \delta_{\boldsymbol{\theta}^{(i)}}$ targeting P , where $\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(n)}$ are sample points and q_n is a probability mass function. By Minkowski's inequality, the difference between IMQ KSD and AIKS of Q_n is

$$\left| \hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k,\|\cdot\|_p}) - \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k,\|\cdot\|_p}) \right| \leq \|\hat{\boldsymbol{w}} - \boldsymbol{w}\|_p.$$

We now consider the term on the right hand side of this inequality.

$$\begin{aligned}
 \|\hat{\mathbf{w}} - \mathbf{w}\|_p^p &= \sum_{j=1}^p |\hat{w}_j - w_j|^p \\
 &= \sum_{j=1}^p \frac{|\hat{w}_j - w_j|^{p-1}}{\hat{w}_j + w_j} |\hat{w}_j^2 - w_j^2| \\
 &= \sum_{j=1}^p \frac{|\hat{w}_j - w_j|^{p-1}}{\hat{w}_j + w_j} \sum_{k,l=1}^n q_n(\boldsymbol{\theta}^{(k)}) q_n(\boldsymbol{\theta}^{(l)}) \left| \hat{k}_0^j(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) - k_0^j(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) \right|.
 \end{aligned}$$

Let $\epsilon(N) = \max_k \|\hat{\mathbf{u}}(\boldsymbol{\theta}^{(k)}) - \mathbf{u}(\boldsymbol{\theta}^{(k)})\|_{\max} = \mathcal{O}(1/\sqrt{N})$ from ergodic theorem. The approximate error for the Stein kernel can be derived as follows.

$$\begin{aligned}
 &\left| \hat{k}_0^j(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) - k_0^j(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) \right| \\
 &\leq k(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) \left| \hat{u}_j(\boldsymbol{\theta}^{(k)}) \hat{u}_j(\boldsymbol{\theta}^{(l)}) - u_j(\boldsymbol{\theta}^{(k)}) u_j(\boldsymbol{\theta}^{(l)}) \right| \\
 &\quad + \left| \nabla_{\boldsymbol{\theta}_j^{(l)}} k(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) \right| \left| \hat{u}_j(\boldsymbol{\theta}^{(k)}) - u_j(\boldsymbol{\theta}^{(k)}) \right| \\
 &\quad + \left| \nabla_{\boldsymbol{\theta}_j^{(k)}} k(\boldsymbol{\theta}^{(l)}, \boldsymbol{\theta}^{(k)}) \right| \left| \hat{u}_j(\boldsymbol{\theta}^{(l)}) - u_j(\boldsymbol{\theta}^{(l)}) \right| \\
 &\leq \left[k(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) \left\{ |\hat{u}_j(\boldsymbol{\theta}^{(k)})| + |u_j(\boldsymbol{\theta}^{(l)})| \right\} + 2 \left| \nabla_{\boldsymbol{\theta}_j^{(l)}} k(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)}) \right| \right] \epsilon(N) \quad (19) \\
 &\leq c_1 \epsilon(N)
 \end{aligned}$$

for bounded constant c_1 . The inequality in (19) follows from (14) and the fact that $\nabla_{\boldsymbol{\theta}_j^{(l)}} k(\boldsymbol{\theta}^{(k)}, \boldsymbol{\theta}^{(l)})$ and $\nabla_{\boldsymbol{\theta}_j^{(k)}} k(\boldsymbol{\theta}^{(l)}, \boldsymbol{\theta}^{(k)})$ are symmetric around zero. Now,

$$\begin{aligned}
 \|\hat{\mathbf{w}} - \mathbf{w}\|_p^p &\leq c_1 \epsilon(N) \sum_{j=1}^p \frac{|\hat{w}_j - w_j|^{p-1}}{\hat{w}_j + w_j} \sum_{k,l=1}^n q_n(\boldsymbol{\theta}^{(k)}) q_n(\boldsymbol{\theta}^{(l)}) \\
 &\leq n^2 c_1 c_2 \epsilon(N)
 \end{aligned}$$

for bounded constant c_2 . Therefore, the approximate error for AIKS is

$$\left| \hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) - \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) \right| \leq \mathcal{O}\left(N^{-1/(2p)}\right)$$

almost surely.

Appendix E. Proof of Proposition 11

From theorem 10, we have

$$\begin{aligned}
 \hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) &\leq \left| \hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) - \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) + \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) \right| \\
 &\leq \mathcal{O}\left\{N^{-1/(2p)}\right\} + \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}).
 \end{aligned}$$

The AIKS $\hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p})$ goes to zero when $Q_n \rightarrow P$ and $N \rightarrow \infty$. This implies that AIKS can detect convergence. Detecting non-convergence means that if AIKS goes to zero then $Q_n \rightarrow P$, equivalently, if $Q_n \not\rightarrow P$ then AIKS does not go to zero. From theorem 10 we have

$$\hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) \geq \mathcal{S}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p}) - \mathcal{O}\left\{N^{-1/(2p)}\right\}$$

The AIKS $\hat{\mathcal{S}}(Q_n, \mathcal{T}_P, \mathcal{G}_{k, \|\cdot\|_p})$ does not go to zero if $Q_n \not\rightarrow P$ and $N \rightarrow \infty$. This indicates that AIKS can detect non-convergence.

Appendix F. Proof of Proposition 12

We check the following assumptions from Theorem 2.1 of Leucht and Neumann (2013).

A1 The process $\{\boldsymbol{\theta}^{(i)}\}_{i \in \mathbb{Z}}$ is β -mixing with $\sum_{r=1}^{\infty} \sqrt{\beta(r)} < \infty$.

A2 Let $k_0 = \lim_{N \rightarrow \infty} \hat{k}_0$.

- (i) $k_0 : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is symmetric and degenerate.
- (ii) k_0 is positive semidefinite.
- (iii) $E_P\{k_0(\boldsymbol{\theta}, \boldsymbol{\theta})\} < \infty$.
- (iv) k_0 is Lipschitz continuous.

The conditions A1, A2 (iii), and A2 (iv) are assumed in Proposition 12.

The k is symmetric, so is k_0 . Let Q be the stationary distribution of the process $\{\boldsymbol{\theta}^{(i)}\}_{i \in \mathbb{Z}}$. We will show that k_0 is degenerate, i.e., $E_Q\{k_0(\boldsymbol{\theta}, \boldsymbol{\theta})\} = 0$ if and only if $Q = P$. If $Q = P$, then $E_Q\{k_0(\boldsymbol{\theta}, \boldsymbol{\theta})\} = 0$ by Proposition 1 of Gorham and Mackey (2017). Suppose that $Q \neq P$ but $E_Q\{k_0(\boldsymbol{\theta}, \boldsymbol{\theta})\} = 0$. Since $E_Q\{k_0(\boldsymbol{\theta}, \boldsymbol{\theta})\} = 0$, we have $E_Q\left\{u_j(\boldsymbol{\theta})k(\boldsymbol{\theta}, \cdot) + \frac{\partial k(\boldsymbol{\theta}, \cdot)}{\partial \theta_j}\right\} = 0$ for $j = 1, \dots, p$ by the proof of Proposition 2 in Gorham and Mackey (2017). Let π and q be the probability density functions of P and Q , respectively. The $u_j(\boldsymbol{\theta}) = \frac{\partial \log \pi(\boldsymbol{\theta})}{\partial \theta_j} = \frac{\partial \log q(\boldsymbol{\theta})}{\partial \theta_j} + \frac{\partial [\log \pi(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})]}{\partial \theta_j}$. We have

$$\begin{aligned} \mathbf{0} &= E_Q\left\{\nabla_{\boldsymbol{\theta}} \log q(\boldsymbol{\theta})k(\boldsymbol{\theta}, \cdot) + \nabla_{\boldsymbol{\theta}}[\log \pi(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})]k(\boldsymbol{\theta}, \cdot) + \nabla_{\boldsymbol{\theta}}k(\boldsymbol{\theta}, \cdot)\right\} \\ &= E_Q\left\{\nabla_{\boldsymbol{\theta}} \log q(\boldsymbol{\theta})k(\boldsymbol{\theta}, \cdot) + \nabla_{\boldsymbol{\theta}}k(\boldsymbol{\theta}, \cdot)\right\} + E_Q\left\{\nabla_{\boldsymbol{\theta}}[\log \pi(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})]k(\boldsymbol{\theta}, \cdot)\right\} \\ &= E_Q\left\{\nabla_{\boldsymbol{\theta}}[\log \pi(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})]k(\boldsymbol{\theta}, \cdot)\right\} \end{aligned} \quad (20)$$

We note that the expectation of (20) is the kernel embedding of a function $g(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \log \left(\frac{\pi(\boldsymbol{\theta})}{q(\boldsymbol{\theta})}\right)$ with respect to Q . By Theorem 4.4 c of Carmeli et al. (2010), the embedding is zero if and only if $g = 0$. Therefore, $\nabla_{\boldsymbol{\theta}} \log \left(\frac{\pi(\boldsymbol{\theta})}{q(\boldsymbol{\theta})}\right) = \mathbf{0}$ and thus $\log \left(\frac{\pi(\boldsymbol{\theta})}{q(\boldsymbol{\theta})}\right) = \mathbf{c}$ where \mathbf{c} is a constant vector. This implies that $\pi(\boldsymbol{\theta}) = e^{\mathbf{c}} q(\boldsymbol{\theta})$. Since π and q integrate to one, $\mathbf{c} = \mathbf{0}$, which contradicts with $Q \neq P$.

The $k_0(\boldsymbol{\theta}, \boldsymbol{\theta}') = \langle \mathcal{T}_P k(\boldsymbol{\theta}, \cdot), \mathcal{T}_P k(\boldsymbol{\theta}', \cdot) \rangle_{\mathcal{K}_k}$ is an inner product and hence positive definite where \mathcal{K}_k is a Hilbert space of functions such that, for all $\boldsymbol{\theta} \in \mathbb{R}^p$, $k(\boldsymbol{\theta}, \cdot) \in \mathcal{K}_k$ and $f(\boldsymbol{\theta}) = \langle f, k(\boldsymbol{\theta}, \cdot) \rangle_{\mathcal{K}_k}$ whenever $f \in \mathcal{K}_k$.

Appendix G. Proof of Proposition 13

We check assumptions from Theorem 3.1 of Leucht and Neumann (2013). The condition B1 is assumed in Proposition 13. The condition B2 is satisfied by Remark 2 (i) of Leucht and Neumann (2013). The condition A2 follows from the proof of Proposition 12.

References

- Christopher Aicher, Srshti Putcha, Christopher Nemeth, Paul Fearnhead, and Emily Fox. Stochastic gradient MCMC for nonlinear state space models. *Bayesian Analysis*, pages 1–23, 2023. doi: 10.1214/23-ba1395.
- P. Alquier, N. Friel, R. Everitt, and A. Boland. Noisy Monte Carlo: Convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, 26(1-2):29–47, 2016.
- Christophe Andrieu and Gareth O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics*, 37:697–725, 2009.
- Yves F. Atchadé, Nicolas Lartillot, and Christian Robert. Bayesian computation for statistical models with intractable normalizing constants. *Brazilian Journal of Probability and Statistics*, 27(4): 416–436, 2013.
- M S Bartlett. Approximate confidence intervals. *Biometrika*, 40(1/2):12–19, 1953a.
- M S Bartlett. Approximate confidence intervals. II. More than one unknown parameter. *Biometrika*, 40(3/4):306–317, 1953b.
- Mark A. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164:1139–1160, 2003.
- Julian Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B*, 36(2):192–225, 1974.
- A Colin Cameron and Per Johansson. Count data regression using series expansions: With applications. *Journal of Applied Econometrics*, 12:203–223, 1997.
- Claudio Carmeli, Ernesto De Vito, Alessandro Toigo, and Veronica Umanitá. Vector valued reproducing kernel Hilbert spaces and universality. *Analysis and Applications*, 8(01):19–61, 2010.
- Charalampos Chaniavidis, Ludger Evers, Tereza Neocleous, and Agostino Nobile. Efficient Bayesian inference for COM-Poisson regression models. *Statistics and Computing*, 28:595–608, 2018.
- J. R. Chazottes, P. Collet, C. Külske, and F. Redig. Concentration inequalities for random fields via coupling. *Probability Theory and Related Fields*, 137:201–225, 2007.
- Der Fa R. Chen and Andrew F. Seila. Multivariate inference in stationary simulation using batch means. In *Proceedings of the 19th conference on Winter Simulation*, pages 302–304, 1987.
- Jie Cheng, Tianxi Li, Elizaveta Levina, and Ji Zhu. High-dimensional mixed graphical models. *Journal of Computational and Graphical Statistics*, 26:367–378, 2017.

- Kacper Chwialkowski, Heiko Strathmann, and Arthur Gretton. A kernel test of goodness of fit. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48, pages 2606–2615, 2016.
- R. W. Conway and W. L. Maxwell. Network dispatching by the shortest-operation discipline. *Operations Research*, 10:51–73, 1962.
- Mary Kathryn Cowles and Bradley P. Carlin. Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904, 1996.
- Halim Damerджи. Strong consistency of the variance estimator in steady-state simulation output analysis. *Mathematics of Operations Research*, 19:494–512, 1994.
- Dirk Eddelbuettel and Romain Francois. Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40:1–18, 2011.
- Y Fan, S P Brooks, and A Gelman. Output assessment for Monte Carlo simulations via the score statistic. *Journal of Computational and Graphical Statistics*, 15(1):178–206, 2006.
- Henri Faure and Christiane Lemieux. Generalized Halton sequences in 2008: a comparative study. *ACM Transactions on Modeling and Computer Simulation*, 19:15:1–15:31, 2009.
- James M. Flegal and Galin L. Jones. Implementing MCMC: Estimating with confidence. In *Handbook of Markov Chain Monte Carlo*, pages 175–197. Chapman and Hall/CRC, 2011.
- James M Flegal, Murali Haran, and Galin L Jones. Markov chain monte carlo: Can we trust the third significant figure? *Statistical Science*, 23(2):250–260, 2008.
- Ho-Ting Fung, Aya Alwan, Justin Wishart, and Alan Huang. mpcmp: Mean-parametrized Conway–Maxwell–Poisson regression, 2019.
- Andrew Gelman and Donald B Rubin. Inference from iterative simulation using multiple sequences. *Statistical science*, pages 457–472, 1992.
- John F Geweke et al. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. Technical report, Federal Reserve Bank of Minneapolis, 1991.
- Charles J Geyer. Practical markov chain monte carlo. *Statistical science*, pages 473–483, 1992.
- Joshua Goldstein, Murali Haran, Ivan Simeonov, John Fricks, and Francesca Chiaromonte. An attraction-repulsion point process model for respiratory syncytial virus infections. *Biometrics*, 71(2):376–385, 2015.
- Jackson Gorham and Lester Mackey. Measuring sample quality with Stein’s method. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- Jackson Gorham and Lester Mackey. Measuring sample quality with kernels. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1292–1301, 2017.
- Jackson Gorham, Anant Raj, and Lester Mackey. Stochastic stein discrepancies. *Advances in Neural Information Processing Systems*, 33:17931–17942, 2020.

- Seth D. Guikema and Jeremy P. Goffelt. A flexible count data regression model for risk analysis. *Risk Analysis*, 28:213–223, 2008.
- W Keith Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- Alan Huang. Mean-parametrized Conway–Maxwell–Poisson regression models for dispersed counts. *Statistical Modelling*, 17:359–380, 2017.
- John Hughes, Murali Haran, and Petruța C Caragea. Autologistic models for binary data on a lattice. *Environmetrics*, 22(7):857–871, 2011.
- David R. Hunter. Curved exponential family models for social networks. *Social Networks*, 29(2): 216–230, 2007.
- David R. Hunter and Mark S. Handcock. Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, 15(3):565–583, 2006.
- David R. Hunter, Mark S. Handcock, Carter T. Butts, Steven M. Goodreau, and Martina Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3):1–29, 2008.
- Ross Ihaka and Robert Gentleman. R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, 5(3):299–314, 1996.
- Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925.
- Galin L Jones, Murali Haran, Brian S Caffo, and Ronald Neath. Fixed-width output analysis for Markov chain Monte Carlo. *Journal of the American Statistical Association*, 101:1537–1547, 2006.
- Robert E Kass, Bradley P Carlin, Andrew Gelman, and Radford M Neal. Markov Chain Monte Carlo in Practice: A Roundtable Discussion. *The American Statistician*, 52(2):93–100, 1998.
- S L Lauritzen and N Wermuth. Graphical models for associations between variables, some of which are qualitative and some quantitative. *The Annals of Statistics*, 17:31–57, 1989.
- Jason D. Lee and Trevor J. Hastie. Learning the structure of mixed graphical models. *Journal of Computational and Graphical Statistics*, 24:230–253, 2015.
- Jeong Eun Lee, Geoff K. Nicholls, and Robin J. Ryder. Calibration Procedures for Approximate Bayesian Credible Sets. *Bayesian Analysis*, 14(4):1245–1269, 12 2019.
- Wilhelm Lenz. Beitrag zum Verständnis der magnetischen Erscheinungen in festen Körpern. *Physikalische Zeitschrift*, 21:613–615, 1920.
- Anne Leucht and Michael H. Neumann. Dependent wild bootstrap for degenerate U- and V-statistics. *Journal of Multivariate Analysis*, 117:257–280, 2013.

- Faming Liang. A double Metropolis-Hastings sampler for spatial models with intractable normalizing constants. *Journal of Statistical Computation and Simulation*, 80(9):1007–1022, 2010.
- Faming Liang, Ick Hoon Jin, Qifan Song, and Jun S. Liu. An adaptive exchange algorithm for sampling from distributions with intractable normalizing constants. *Journal of the American Statistical Association*, 111(513):377–393, 2016.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: a general purpose Bayesian inference algorithm. *Advances in Neural Information Processing Systems*, 29, 2016.
- Anne Marie Lyne, Mark Girolami, Yves Atchadé, Heiko Strathmann, and Daniel Simpson. On Russian Roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statistical Science*, 30(4):443–467, 2015.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward 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.
- Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.
- Iain Murray, Zoubin Ghahramani, and David J C Mackay. MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*, pages 359–366, 2006.
- Jaewoo Park and Murali Haran. Bayesian inference in the presence of intractable normalizing functions. *Journal of the American Statistical Association*, 113(523):1372–1390, 2018.
- Jaewoo Park and Murali Haran. A function emulation approach for doubly intractable distributions. *Journal of Computational and Graphical Statistics*, 29(1):66–77, 2020.
- James Gary Propp and David Bruce Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures & Algorithms*, 9(1-2):223–252, 1996.
- Rajesh Ranganath, Jaan Altosaar, Dustin Tran, and David M Blei. Operator variational inference. *Advances in Neural Information Processing Systems*, 29, 2016.
- Vinayak Rao, Lizhen Lin, and David B. Dunson. Data augmentation for models based on rejection sampling. *Biometrika*, 103:319–335, 2016.
- Eduardo E. Ribeiro, Walmes M. Zeviani, Wagner H. Bonat, Clarice G.B. Demetrio, and John Hinde. Reparametrization of COM–Poisson regression models with applications in the analysis of experimental data. *Statistical Modelling*, 20:443–466, 2020.
- Christian P Robert and George Casella. *Monte Carlo statistical methods*. Springer, 2004.

- Garry Robins, Pip Pattison, Yuval Kalish, and Dean Lusher. An introduction to exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):173–191, 2007.
- Vivekananda Roy. Convergence diagnostics for Markov chain Monte Carlo. *Annual Review of Statistics and Its Application*, 7:387–412, 2020.
- A. J. Sáez-Castillo and A. Conde-Sánchez. A hyper-Poisson regression model for overdispersed and underdispersed count data. *Computational Statistics and Data Analysis*, 61:148–157, 2013.
- Galit Shmueli, Thomas P. Minka, Joseph B. Kadane, Sharad Borle, and Peter Boatwright. A useful distribution for fitting discrete data: Revival of the Conway–Maxwell–Poisson distribution. *Journal of the Royal Statistical Society: Series C*, 54:127–142, 2005.
- Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability*, volume 2, pages 583–602, 1972.
- David J Strauss. A model for clustering. *Biometrika*, 62(2):467–475, 1975.
- Linda S.L. Tan and Nial Friel. Bayesian variational inference for exponential random graph models. *Journal of Computational and Graphical Statistics*, 29:910–928, 2020.
- D Vats, FB Gonçalves, K Łatuszyński, and GO Roberts. Efficient Bernoulli factory Markov chain Monte Carlo for intractable posteriors. *Biometrika*, 109:369–385, 2022.
- Dootika Vats, James M. Flegal, and Galin L. Jones. Multivariate output analysis for Markov chain Monte Carlo. *Biometrika*, 106:321–337, 2019.
- Hanwen Xing, Geoff K. Nicholls, and Jeong Eun Lee. Calibrated approximate Bayesian inference. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6912–6920, 2019.