Published in final edited form as:

Technometrics. 2016 ; 58(3): 604–318. doi:10.1080/00401706.2016.1142905.

A Bootstrap Metropolis–Hastings Algorithm for Bayesian Analysis of Big Data

Faming Liang,

Professor, Department of Biostatistics, University of Florida, Gainesville, FL 32611

Jinsu Kim, and

Enterprise Solution Specialist, Big Data Analysis Consulting Team, LG CNS, Seoul, Republic of Korea

Qifan Song

Assistant Professor, Department of Statistics, Purdue University, West Lafayette, IN 47906

Abstract

Markov chain Monte Carlo (MCMC) methods have proven to be a very powerful tool for analyzing data of complex structures. However, their computer-intensive nature, which typically require a large number of iterations and a complete scan of the full dataset for each iteration, precludes their use for big data analysis. In this paper, we propose the so-called bootstrap Metropolis-Hastings (BMH) algorithm, which provides a general framework for how to tame powerful MCMC methods to be used for big data analysis; that is to replace the full data log-likelihood by a Monte Carlo average of the log-likelihoods that are calculated in parallel from multiple bootstrap samples. The BMH algorithm possesses an embarrassingly parallel structure and avoids repeated scans of the full dataset in iterations, and is thus feasible for big data problems. Compared to the popular divide-and-combine method, BMH can be generally more efficient as it can asymptotically integrate the whole data information into a single simulation run. The BMH algorithm is very flexible. Like the Metropolis-Hastings algorithm, it can serve as a basic building block for developing advanced MCMC algorithms that are feasible for big data problems. This is illustrated in the paper by the tempering BMH algorithm, which can be viewed as a combination of parallel tempering and the BMH algorithm. BMH can also be used for model selection and optimization by combining with reversible jump MCMC and simulated annealing, respectively.

Keywords

Big Data; Bootstrap; Markov Chain Monte Carlo; Metropolis-Hastings; Parallel Computing

1 Introduction

The integration of computer technology into science and daily life has enabled the collection of massive volumes of data, such as climate data, high-throughput biological assay data,

website transaction logs, and credit card records. However, such massive data cannot be practically analyzed on a single commodity computer because their sizes are too large to fit in memory or it is too time consuming to process when the current statistical methods are used. To circumvent this obstacle, one may have to resort to parallel and distributed architectures, with multicore and cloud computing platforms providing access to hundreds to thousands of processors. While the parallel and distributed architectures present new capabilities for storage and manipulation of data, from an inferential point of view, it is unclear how the current statistical methodology can be transported to the paradigm of big data. Also, with growing size typically comes a growing complexity of data structures and of the models needed to account for the structures. Big data have put a great challenge on the current statistical methodology.

During the past few decades, Markov chain Monte Carlo (MCMC) methods have been widely used in statistical data analysis, and they have proven to be a very powerful and typically unique computational tool for analyzing data of complex structures. However, MCMC methods are compute intensive, which typically require a large number of iterations and a complete scan of the full dataset for each iteration. This characteristic precludes their use for big data analysis.

Motivated by the successes of MCMC methods in analyzing data of complex structures, we propose in this paper a bootstrap Metropolis-Hastings (BMH) algorithm which is feasible for big data and workable on parallel and distributed architectures. The BMH works by replacing the full data log-likelihood by a Monte Carlo average of the log-likelihoods that are calculated in parallel from multiple bootstrap samples, where the bootstrap sample refers to a small set of observations drawn from the full dataset at random and with/without replacement. By this way, BMH avoids repeated scans of the full dataset in iterations, while is still able to produce sensible solutions, such as parameter estimates or posterior samples, to the problem under consideration. The convergence of the BMH algorithm is also studied in this paper. It is interesting to point out that BMH can be generally more efficient than the divide-and-combine method, as BMH is able to integrate the whole data information into a single simulation run. The BMH algorithm is very flexible. Like the Metropolis-Hastings (MH) algorithm, it can serve as a basic building block for developing advanced MCMC algorithms that are feasible for big data problems. This has been illustrated in the paper through a combined use of BMH with parallel tempering (Geyer, 1991; Hukushima and Nemoto, 1996) for a neural network example. BMH can also be used for model selection and optimization by combining with reversible jump MCMC (Green, 1995) and simulated annealing (Kirkpatrick et al., 1983), respectively. Finally we note that the new algorithm is named with the word “bootstrap”, because it employs a resampling-based Monte Carlo method to infer the posterior distribution of the whole data rather than to assess the quality of some estimators.

In the literature, there have been a few methods proposed for big data analysis, such as the aggregated estimating equation method (Lin and Xi, 2011), the resampling-based stochastic approximation method (Liang et al., 2013), the bag of little Bootstraps (Kleiner et al., 2012), the approximate Metropolis-Hastings test (AMHT) method (Korattikara *et al.*, 2014, Bardenet *et al.*, 2014), and the stochastic gradient-based Monte Carlo methods (Welling and

Teh, 2011; Ahn *et al.*, 2012; Chen *et al.*, 2014). The aggregated estimating equation method employs a divide-and-combine strategy. It is first to compress the raw data of each partition of the full dataset into some low dimensional statistics, and then to obtain an approximation to the estimating equation estimator, the aggregated estimating equation estimator, by solving an equation aggregated from the saved low dimensional statistics in all partitions. Liang et al. (2013) proposed a new parameter estimator, *maximum mean log-likelihood estimator*, for big data problems, and a resampling-based stochastic approximation method for calculating such an estimator. The resampling-based stochastic approximation method successfully avoids some difficulties involved in big data computation, such as high-dimensional matrix inversion. The bag of little bootstraps provides an efficient way of bootstrapping for big data estimators, which functions by combining the results of bootstrapping multiple small subsets of the big original dataset. The AMHT method will be described in detail in the Supplementary Material. It provides an approximate way to draw samples from the big data posterior, however, there is no clear theory on how to make Bayesian inference based on its samples. The stochastic gradient-based methods are to estimate the gradient, which are required in Langevin dynamics and Hamiltonian Monte Carlo, using a subset of the data. To counteract the effect of the noisy gradient, the moving step size needs to be annealed to zero to reach the fixed point and, consequently, the mixing rate of these methods can be very slow. Although these methods can be improved in certain ways, e.g., see Chen *et al.* (2014) and Ahn *et al.* (2012), their moving step size still needs to be kept at an extremely small value to counterbalance the factor n/m (full dataset size/subset size) appeared in the gradient estimator. Hence, these methods are hard to be used when n is extremely large. The BMH is different from the AMHT and stochastic gradient methods. Instead of directly drawing samples from the full data posterior, BMH essentially aims to draw samples from a tempered posterior distribution, and then makes inference based on the technique of importance sampling. Sampling from the tempered posterior improves the mixing rate of the algorithm. BMH provides a simple yet effective way of uncertainty quantification for big data problems.

The remainder of this paper is organized as follows. Section 2 describes the BMH algorithm. Section 3 studies the convergence of the BMH algorithm and establishes the asymptotic normality of the approximated posterior distribution. Section 4 illustrates the flexibility of the BMH algorithm through a combined use with parallel tempering for a neural network example. Section 5 describes two efficient schemes for drawing bootstrap samples from distributed architectures. Section 6 concludes the paper with a brief discussion.

2 The Bootstrap Metropolis-Hastings Algorithm

Let $\mathcal{D}$ denote a big dataset of n observations, where n can be very large. Let $f(\mathcal{D}|\boldsymbol{\theta})$ denote the likelihood function of the data, where $\boldsymbol{\theta} \in \Theta$ is the parameter vector and Θ is the parameter space. Suppose that we are interested in conducting Bayesian inference for the big data through drawing samples from the posterior distribution

$$\pi(\boldsymbol{\theta}|\mathcal{D}) \propto \pi(\boldsymbol{\theta})f(\mathcal{D}|\boldsymbol{\theta}), \quad \boldsymbol{\theta} \in \Theta, \quad (1)$$

where $\pi(\boldsymbol{\theta})$ denotes the prior distribution of $\boldsymbol{\theta}$. Since n is large, a direct simulation of $\pi(\boldsymbol{\theta}|\mathcal{D})$ is infeasible. A natural idea is divide-and-combine (D&C), which is to simulate a sequence of posterior distributions $\pi(\boldsymbol{\theta}|\mathcal{D}_1), \dots, \pi(\boldsymbol{\theta}|\mathcal{D}_k)$, where $\{\mathcal{D}_1, \dots, \mathcal{D}_k\}$ forms a partition of $\mathcal{D}$, and then to infer $\pi(\boldsymbol{\theta}|\mathcal{D})$ based on the combined samples simulated from $\pi(\boldsymbol{\theta}|\mathcal{D}_1), \dots, \pi(\boldsymbol{\theta}|\mathcal{D}_k)$. However, this method is potentially inefficient, as it still requires to scan the full dataset for a large number of times as the simulation goes on. To circumvent this obstacle, we propose the BMH algorithm, which avoids repeated scans of the full dataset in iterations while being able to asymptotically integrate the whole data information into a single simulation run. This makes the BMH algorithm potentially much more efficient than the D&C method. The BMH algorithm can be described as follows.

Let D_i denote a bootstrap sample of $\mathcal{D}$, which is drawn from the full dataset at random and with/without replacement. Let m denote the size of D_i . If sampling is done without

replacement, D_i is called a subsample or $\binom{n}{m}$ -bootstrap sample (Politis and Romano, 1994). Otherwise, D_i is called an m -out-of- n bootstrap sample or m/n -bootstrap sample (Bickel et al., 1997). Let $\tilde{l}(D_i|\boldsymbol{\theta})$ denote a likelihood-like function of D_i , and define

$$l_{m,n,k}(\boldsymbol{D}_s|\boldsymbol{\theta}) = \frac{1}{k} \sum_{i=1}^k \log \tilde{l}(D_i|\boldsymbol{\theta}), \quad (2)$$

where k denotes the number of bootstrap samples drawn from $\mathcal{D}$, and $\boldsymbol{D}_s = \{D_1, \dots, D_k\}$ is the collection of the bootstrap samples. In this paper, we assume that all observations in $\mathcal{D}$

are independently and identically distributed (*i.i.d.*). Then, regardless $\binom{n}{m}$ - or m/n -bootstrap sample, we define

$$\tilde{f}(D_i|\boldsymbol{\theta}) = \tilde{f}(X_1^{(i)}, \dots, X_m^{(i)}|\boldsymbol{\theta}) = \prod_{j=1}^m f(X_j^{(i)}|\boldsymbol{\theta}), \quad (3)$$

where $X_j^{(i)}$ denotes the j th element of D_i . The BMH algorithm works by iterating between the following steps:

Algorithm 1

1. Draw $\boldsymbol{\vartheta}$ from a proposal distribution $Q(\boldsymbol{\theta}, \boldsymbol{\vartheta})$.
2. Draw k bootstrap samples $D_1, \dots, D_k$ via $\binom{n}{m}$ - or m/n -bootstrapping. Let $\boldsymbol{D}_s = \{D_1, \dots, D_k\}$.
3. Calculate the BMH ratio:

$$r(\boldsymbol{\theta}_t, \mathbf{D}_s, \boldsymbol{\vartheta}) = \exp\{l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\vartheta}) - l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}_t)\} \frac{\pi_m(\boldsymbol{\vartheta})}{\pi_m(\boldsymbol{\theta}_t)} \frac{Q(\boldsymbol{\vartheta}, \boldsymbol{\theta}_t)}{Q(\boldsymbol{\theta}_t, \boldsymbol{\vartheta})},$$

where $\pi_m(\boldsymbol{\theta})$ denotes the prior density/mass function of $\boldsymbol{\theta}$ and it can depend on the value of m .

4. Set $\boldsymbol{\theta}_{t+1} = \boldsymbol{\vartheta}$ with probability $\alpha(\boldsymbol{\theta}_t, \mathbf{D}_s, \boldsymbol{\vartheta}) = \min\{1, r(\boldsymbol{\theta}_t, \mathbf{D}_s, \boldsymbol{\vartheta})\}$, and set $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t$ with the remaining probability.

In BMH, $\{\boldsymbol{\theta}_t\}$ form a Markov chain with the transition kernel given by

$$P_{m,n,k}(\boldsymbol{\theta}, d\boldsymbol{\vartheta}) = \sum_{\mathbf{D}_s \in \mathbb{D}} \alpha(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}) Q(\boldsymbol{\theta}, \boldsymbol{\vartheta}) \psi(\mathbf{D}_s) + \delta_{\boldsymbol{\theta}}(d\boldsymbol{\vartheta}) \left[1 - \sum_{\mathbf{D}'_s \in \mathbb{D}} \int_{\Theta} \alpha(\boldsymbol{\theta}, \mathbf{D}'_s, \boldsymbol{\vartheta}') Q(\boldsymbol{\theta}, d\boldsymbol{\vartheta}') \psi(\mathbf{D}'_s) \right], \quad (4)$$

where $\mathbb{D}$ denote the space of $\mathbf{D}_s$, $\psi(\mathbf{D}_s)$ denotes the probability of drawing $\mathbf{D}_s$, and $\delta_{\boldsymbol{\theta}}(\cdot)$ is an

indicator function. For $\binom{n}{m}$ -bootstrapping, $\psi(\mathbf{D}_s) = \binom{n}{m}^{-k}$; and for m/n -bootstrapping, $\psi(\mathbf{D}_s) = 1/n^{mk}$. Let

$$P_{m,n,k,\mathbf{D}_s}(\boldsymbol{\theta}, d\boldsymbol{\vartheta}) = \alpha(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}) Q(\boldsymbol{\theta}, \boldsymbol{\vartheta}) + \delta_{\boldsymbol{\theta}}(d\boldsymbol{\vartheta}) \left[1 - \int_{\Theta} \alpha(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}') Q(\boldsymbol{\theta}, d\boldsymbol{\vartheta}') \right],$$

denote the transitional kernel corresponding to a particular subset data $\mathbf{D}_s$. Then $P_{m,n,k}(\boldsymbol{\theta}, d\boldsymbol{\vartheta})$ can be written as a mixture of $P_{m,n,k,\mathbf{D}_s}(\boldsymbol{\theta}, d\boldsymbol{\vartheta})$'s; that is,

$$P_{m,n,k}(\boldsymbol{\theta}, d\boldsymbol{\vartheta}) = \sum_{\mathbf{D}_s \in \mathbb{D}} P_{m,n,k,\mathbf{D}_s}(\boldsymbol{\theta}, d\boldsymbol{\vartheta}) \psi(\mathbf{D}_s). \quad (5)$$

To make $P_{m,n,k}(\boldsymbol{\theta}, d\boldsymbol{\vartheta})$ have an interpretable stationary distribution, one may assume that each of the kernels $P_{m,n,k,\mathbf{D}_s}(\boldsymbol{\theta}, d\boldsymbol{\vartheta})$'s has about the same stationary distribution; that is, $P_{m,n,k,\mathbf{D}_s}$ is approximately independent of $\mathbf{D}_s$. If n is fixed, one may choose a small value of k and a value of m close to n , then the algorithm will have a stationary distribution approximately equal to $\pi(\boldsymbol{\theta})$. However, this setting does not benefit much to big data problems. In this paper, we study the performance of the BMH algorithm under the setting $n \rightarrow \infty$. As the sample size n increases, one may wish to model the data by more complex models. Hence, a reasonable assumption is that m will increase with n . However, the dependence of the stability of $P_{m,n,k,\mathbf{D}_s}$ on the orders of m and k is tricky. In this paper, we study the convergence of the algorithm under the setting $m/k \rightarrow 0$, while both m and k

increase with n . Obviously, for a fixed value of k , increasing the value of m can improve the stability of P_{m,n,k,D_s} ; which implies that the algorithm is also possibly valid under the setting $m/k \rightarrow \infty$. This will be left for a future study. We note that BMH introduces a new type of Markov chain, where the target distribution $\pi_m(\boldsymbol{\theta}) \exp(l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}))$ is changing from iteration to iteration while the detailed balance condition is reserved. As long as k or m is reasonably large, the target distribution will become very stable, and the resulting samples of $\boldsymbol{\theta}$ will not deviate much from the equilibrium of the Markov chain.

Regarding the BMH algorithm, we have a few additional remarks:

- As shown in Section 3, the two sampling schemes, $\binom{n}{m}$ -bootstrapping and m/n -bootstrapping, result in the same stationary distribution of the Markov chain.
- If we define

$$\mathbb{H}(\boldsymbol{\theta}|\mathbf{D}_s) = -l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - \log \pi_m(\boldsymbol{\theta}),$$

which is the so-called energy function of the posterior distribution, then the BMH ratio can be written as

$$r(\boldsymbol{\theta}_t, \mathbf{D}_s, \boldsymbol{\vartheta}) = \exp\{\mathbb{H}(\boldsymbol{\theta}_t|\mathbf{D}_s) - \mathbb{H}(\boldsymbol{\vartheta}|\mathbf{D}_s)\} \frac{Q(\boldsymbol{\vartheta}, \boldsymbol{\theta}_t)}{Q(\boldsymbol{\theta}_t, \boldsymbol{\vartheta})}.$$

This formula will be used in Section 4 for the Bayesian neural network example.

- In BMH, the proposal distribution can be chosen as in the MH algorithm for a fixed target distribution; that is, choosing an appropriate proposal distribution such that the resulting BMH chain is irreducible and aperiodic, and the BMH moves have a reasonable acceptance rate, e.g., between 0.2 and 0.4 as suggested by Gelman *et al.* (1996) for conventional MH algorithms. A formal statement for the requirement of the proposal distribution will be given in condition (B) of Section 3.

3 Convergence of the Bootstrap Metropolis-Hastings Algorithm

In this section, we first prove the ergodicity of the BMH algorithm and then discuss how to make Bayesian inference for the full dataset based on the BMH samples. The ergodicity of

BMH will be studied in two scenarios, namely, $\binom{n}{m}$ -bootstrapping and m/n -bootstrapping.

3.1 $\binom{n}{m}$ -bootstrapping

To study the ergodicity of BMH, we first assume the following condition holds:

$$(A) \quad \sup_{\theta \in \Theta} E |\log f(X|\theta)|^2 < \infty.$$

Conditioned on the dataset $\mathcal{D}$, $\{\log \tilde{f}(D_1|\theta), \dots, \log \tilde{f}(D_k|\theta)\}$ forms a simple random sample without replacement from a finite population. Motivated by this observation, we define the U -statistic

$$U_{m,n}(\mathcal{D}|\theta) = \binom{n}{m}^{-1} \sum_{D_i \in \mathfrak{D}} h(D_i) = \binom{n}{m}^{-1} \sum_{D_i \in \mathfrak{D}} \log \tilde{f}(D_i|\theta), \quad (6)$$

where $\mathfrak{D}$ is the space of D_i and it contains all the possible $\binom{n}{m}$ subsamples of size m , and $h(D_i) = \log \tilde{f}(D_i|\theta)$ is called the kernel of the U -statistic. Thus, $U_{m,n}(\mathcal{D}|\theta)$ is the conditional mean of $\log \tilde{f}(D_1|\theta)$ on the dataset $\mathcal{D}$. U -statistics were introduced by Hoeffding (1948), which represent a class of statistics that is especially important in estimation theory. Many well known test statistics and estimators, such as mean and variance, are in fact members of this class. The simple structure of U -statistics has made them widely used for studying general estimation processes such as bootstrapping and jackknifing, and for generalizing those parts of asymptotic theory concerned with sample means. Refer to Lee (1990) for an overview of theory and practice of U -statistics.

By the law of iterated expectations, it is easy to show that

$$E(l_{m,n,k}(D_s|\theta) - U_{m,n}(\mathcal{D}|\theta))^2 = E[E((l_{m,n,k}(D_s|\theta) - U_{m,n}(\mathcal{D}|\theta))^2 | \mathcal{D})] \leq \frac{m}{k} \text{Var}(\log f(X|\theta)), \quad (7)$$

which, by condition (A) and Markov inequality, implies that as $k \rightarrow \infty$, $n \rightarrow \infty$ and $m/k \rightarrow 0$,

$$l_{m,n,k}(D_s|\theta) - U_{m,n}(\mathcal{D}|\theta) \xrightarrow{p} 0, \quad (8)$$

where $\xrightarrow{p}$ denotes the convergence in probability. The inequality in (7) is due to the factor $[\binom{n}{m} - k] / [\binom{n}{m} - 1]$ of finite population correction for variance of sample mean. Let

$g_m(\mathcal{D}|\theta) = \exp\{E[\log \tilde{f}(D_1|\theta)]\}$. In the scenario of $\binom{n}{m}$ -bootstrapping for *i.i.d* observations, it follows from (3) that

$$g_m(\mathcal{D}|\boldsymbol{\theta}) = \exp\{mE[\log f(X_1|\boldsymbol{\theta})]\}. \quad (9)$$

The variance of a U -statistic based on i.i.d random variables can be expressed in terms of certain conditional expectations. Define for $c = 1, 2, \dots, m$ the conditional expectation

$$h_c(x_1, \dots, x_c) = E\left\{\log \tilde{f}(X_1, \dots, X_m | X_1 = x_1, \dots, X_c = x_c, \boldsymbol{\theta})\right\},$$

and their variances

$$\sigma_c^2 = \text{Var}(h_c(X_1, \dots, X_c)).$$

Then, according to Hoeffding's theorem (see e.g., Lee, 1990, p.12),

$$\text{Var}(U_{m,n}) = \binom{n}{m}^{-1} \sum_{c=1}^m \binom{m}{c} \binom{n-m}{m-c} \sigma_c^2,$$

provided condition (A) holds. Since $\sigma_c^2 = c \text{Var}(\log f(X|\boldsymbol{\theta}))$ for the U -statistic defined in (6), we have

$$\text{Var}(U_{m,n}) = \frac{m^2}{n} \text{Var}(\log f(X|\boldsymbol{\theta})), \quad (10)$$

which implies the following theorem holds:

Theorem 1: Assume that the condition (A) holds and $m = O(n^\gamma)$. If $0 < \gamma < 1/2$, then

$$U_{m,n}(\mathcal{D}|\boldsymbol{\theta}) - \log(g_m(\mathcal{D}|\boldsymbol{\theta})) \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty. \quad (11)$$

Proof: It follows from (10) and Markov inequality.

This theorem implies that the U -statistic $U_{m,n}(\mathcal{D}|\boldsymbol{\theta})$, the mean of $l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta})$ conditioned on the observed data set, provides a consistent estimator for the population mean $\log(g_m(\mathcal{D}|\boldsymbol{\theta}))$. Note that $U_{m,n}(\mathcal{D}|\boldsymbol{\theta})$ depends on both m and n , while $\log(g_m(\mathcal{D}|\boldsymbol{\theta}))$ depends on m only. Combining (8) and (11), we have for any $\boldsymbol{\theta} \in \Theta$,

$$l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - \log(g_m(\mathcal{D}|\boldsymbol{\theta})) \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty, \quad (12)$$

where, as implied by (8) and Theorem 1,

$$m = O(n^\gamma) \quad \text{and} \quad k = O(n^{\gamma+e_0}), \quad (13)$$

for some $e_0 > 0$ and $0 < \gamma < 1/2$. Define

$$\Gamma_{m,n,k}(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}) = \frac{\exp\{l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\vartheta}) - l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta})\}}{g_m(\mathcal{D}|\boldsymbol{\vartheta})/g_m(\mathcal{D}|\boldsymbol{\theta})},$$

and

$$\lambda_{m,n,k}(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}) = |\log(\Gamma_{m,n,k}(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}))| = |[l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\vartheta}) - \log(g_m(\mathcal{D}|\boldsymbol{\vartheta}))] - [l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - \log(g_m(\mathcal{D}|\boldsymbol{\theta}))]|.$$

It follows from (12) that

$$\lambda_{m,n,k}(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}) \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty. \quad (14)$$

Define

$$\rho(\boldsymbol{\theta}) = 1 - \sum_{\mathbf{D}_s \in \mathbb{D}} \int_{\Theta} \alpha(\boldsymbol{\theta}, \mathbf{D}_s, \boldsymbol{\vartheta}) Q(\boldsymbol{\theta}, d\boldsymbol{\vartheta}) \psi(\mathbf{D}_s).$$

which represents the mean rejection probability of a BMH move starting from $\boldsymbol{\theta}$.

To establish the convergence of BMH, we also consider the transition kernel

$$P_m(\boldsymbol{\theta}, \boldsymbol{\vartheta}) = \alpha(\boldsymbol{\theta}, \boldsymbol{\vartheta}) Q(\boldsymbol{\theta}, \boldsymbol{\vartheta}) + \delta_{\boldsymbol{\theta}}(d\boldsymbol{\vartheta}) \left[1 - \int_{\Theta} \alpha(\boldsymbol{\theta}, \boldsymbol{\vartheta}') Q(\boldsymbol{\theta}, \boldsymbol{\vartheta}') d\boldsymbol{\vartheta}' \right], \quad (15)$$

where $\alpha(\boldsymbol{\theta}, \boldsymbol{\vartheta}') = \min\{1, r(\boldsymbol{\theta}, \boldsymbol{\vartheta}')\}$, and $r(\boldsymbol{\theta}, \boldsymbol{\vartheta}')$ is the MH ratio for a move with the proposal $Q(\boldsymbol{\theta}, \boldsymbol{\vartheta}')$ and the invariant distribution

$$\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D}) \propto \pi_m(\boldsymbol{\theta}) g_m(\mathcal{D}|\boldsymbol{\theta}). \quad (16)$$

Further, we assume the following conditions hold:

- (B) Assume that P_m defines an irreducible and aperiodic Markov chain such that $\tilde{\pi}_m(\cdot)P_m = \tilde{\pi}_m(\cdot)$. Therefore, for any starting point $\theta_0 \in \Theta$,
 $\lim_{t \rightarrow \infty} \|P_m^t(\theta_0, \cdot) - \tilde{\pi}_m(\cdot)\| = 0$, where $\|\cdot\|$ denotes the total variation norm.
- (C) For any $(\theta, \vartheta) \in \Theta \times \Theta$,

$$0 < \Gamma_{m,n,k}(\theta, D_s, \vartheta) < \infty, \quad \psi(\cdot) - a.s.,$$

where $\psi(D_s)$ is the sampling probability of D_s from $\mathcal{D}$.

Following from the standard theory of the MH algorithm (see e.g. Tierney, 1994), condition (B) can be simply satisfied by choosing an appropriate proposal distribution $Q(\cdot, \cdot)$. Condition (C) is equivalent to assuming $0 < \exp \{I_{m,n,k}(D_s|\vartheta) - I_{m,n,k}(D_s|\theta)\} < \infty$, which ensures the BMH ratio to be well defined in simulations.

Lemma 1 states that the kernel $P_{m,n,k}$ defined in (4), has a stationary distribution. Its proof follows the proof of Lemma 1 (except for some notational changes) of Liang and Jin (2013) for the Monte Carlo MH algorithm, where the MH ratio includes a random quantity calculated using Monte Carlo samples. A similar theorem has also been proved in Andrieu and Roberts (2009) for the pseudo-marginal approach, where the likelihood function is approximated using a Monte Carlo approach such as the importance sampling method.

Lemma 1: Assume conditions (B) and (C) hold. Then for any $m, n, k \in \mathbb{N}$ such that $\rho(\theta) > 0$ for all $\theta \in \Theta$, $P_{m,n,k}$ is also irreducible and aperiodic, and hence there exists a stationary distribution $\hat{\pi}_{m,n,k}(\theta|\mathcal{D})$ such that for any $\theta_0 \in \Theta$,

$$\lim_{t \rightarrow \infty} \|P_{m,n,k}^t(\theta_0, \cdot) - \hat{\pi}_{m,n,k}(\cdot)\| = 0.$$

Lemma 2 concerns the distance between the kernel $P_{m,n,k}$ and the kernel P_m . It states that the two kernels can be arbitrarily close to each other, provided that k and n are large enough. By assuming (13) holds, both m and k are treated as a given function of n . Therefore, we consider only the growth of n in the following lemmas and theorems.

Lemma 2: Assume the conditions (A), (B) and (C) hold. If (13) holds, then for any $\varepsilon \in (0, 1]$ and any $\theta \in \Theta$, there exists $N(\theta) \in \mathbb{N}$ such that for any $\phi: \Theta \rightarrow [-1, 1]$ and any $n > N(\theta)$,

$$|P_{m,n,k}\phi(\theta) - P_m\phi(\theta)| \leq 4\varepsilon.$$

The proof of this lemma is given in the Supplementary Material. Theorem 2 concerns the ergodicity of BMH, which states that the kernel $P_{m,n,k}$ shares the same stationary distribution with the kernel P_m when both n and k become large. The proof of this theorem follows directly from the proof of Theorem 1 of Liang and Jin (2013).

Theorem 2: Assume the conditions (A), (B) and (C) hold and the observations are i.i.d. If (13) holds, then for any $\varepsilon \in (0, 1]$ and any $\theta_0 \in \Theta$, there exist $N(\varepsilon, \theta_0) \in \mathbb{N}$ and $T(\varepsilon, \theta_0, n) \in \mathbb{N}$ such that for any $n > N(\varepsilon, \theta_0)$ and $t > T(\varepsilon, \theta_0, n, k)$,

$$\|P_{m,n,k}^t(\theta_0, \cdot) - \tilde{\pi}_m(\cdot)\| \leq \varepsilon,$$

where $\tilde{\pi}_m(\cdot)$ is the stationary distribution of P_m as defined in (16).

Theorem 2 establishes the convergence of BMH under the setting that the bootstrap samples D_s are updated at each iteration. In practice, to avoid frequent updating of bootstrap samples, one may repeatedly use them for a small number of iterations. This may accelerate BMH, especially when m is large. Let κ_0 denote the number of repeated iterations. Following from (5), the transitional kernel of BMH for this repeated bootstrap version can be written as

$$\tilde{P}_{m,n,k}(\theta, d\vartheta) = \sum_{D_s \in \mathbb{D}} P_{m,n,k,D_s}^{\kappa_0}(\theta, d\vartheta) \psi(D_s). \quad (17)$$

Under the same conditions of Lemma 2, it is shown in the Supplementary Material that

$$\|\tilde{P}_{m,n,k}\phi(\theta) - P_m^{\kappa_0}\phi(\theta)\| \leq 4\kappa_0\varepsilon. \quad (18)$$

Hence, the convergence established in Theorem 2 still follows for the BMH algorithm with repeated use of bootstrap samples.

In Section 3.3, we will consider the asymptotics of $\tilde{\pi}_m(\theta^{\mathcal{D}})$. In particular, we will discuss how $\tilde{\pi}_m(\theta^{\mathcal{D}})$ is related to the whole data posterior $\pi(\theta^{\mathcal{D}})$ as m becomes large, and how to make Bayesian inference for $\pi(\theta^{\mathcal{D}})$ based on the samples simulated by the BMH algorithm.

3.2 m/n -bootstrapping

Under this scenario, the ergodicity of the BMH algorithm can be studied in a similar way to $\binom{n}{m}$ -bootstrapping. In what follows, we show with appropriate conditions that BMH has the same stationary distribution for the two bootstrapping schemes.

First, we define

$$V_{m,n}(\mathcal{D}|\theta) = n^{-m} \sum_{1 \leq j_1, \dots, j_m \leq n} \log \tilde{f}(X_{j_1}, \dots, X_{j_m}|\theta),$$

which is the conditional mean of $l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta})$ on the dataset $\mathcal{D}$ and is called a von Mises statistic or V -statistic. A straightforward calculation (see the Supplementary Material for the details) shows that

$$E(l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - V_{m,n}(\mathcal{D}|\boldsymbol{\theta}))^2 = \frac{m}{k} \left(1 - \frac{1}{n}\right) \text{Var}(\log f(X|\boldsymbol{\theta})), \quad (19)$$

which, by condition (A) and Chebyshev's inequality, implies

$$l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - V_{m,n}(\mathcal{D}|\boldsymbol{\theta}) \xrightarrow{p} 0, \quad (20)$$

as $k \rightarrow \infty$, $n \rightarrow \infty$ and $m/k \rightarrow 0$. Further, we have the following theorem.

Theorem 3: Assume that the condition (A) holds. Let $m = O(n^\gamma)$. If $0 < \gamma < 1/2$, then

$$V_{m,n}(\mathcal{D}|\boldsymbol{\theta}) - \log(g_m(\mathcal{D}|\boldsymbol{\theta})) \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty. \quad (21)$$

Proof: It follows from (3) that $V_{m,n}(\mathcal{D}|\boldsymbol{\theta}) = \frac{m}{n} \sum_{X_i \in \mathcal{D}} \log f(X_i|\boldsymbol{\theta})$. Therefore, $\text{Var}(V_{m,n}(\mathcal{D}|\boldsymbol{\theta})) = \frac{m^2}{n} \text{Var}(\log f(X|\boldsymbol{\theta}))$, which, by Markov inequality, concludes the theorem.

Combining (20) and (21), we have

$$l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - \log(g_m(\mathcal{D}|\boldsymbol{\theta})) \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty,$$

under the setting (13). Then, by the same reasoning as Theorem 2, we have the following theorem:

Theorem 4: Assume that the observations are i.i.d., and the conditions (A), (B) and (C) hold. If (13) is satisfied, then BMH with m/n -bootstrapping has asymptotically the same

stationary distribution as with $\binom{n}{m}$ -bootstrapping.

3.3 Bayesian Inference

This subsection is organized as follows. In Section 3.3.1, we establish the asymptotic normality of $\tilde{\boldsymbol{\pi}}_m(\boldsymbol{\theta}|\mathcal{D})$. In Section 3.3.2 and Section 3.3.3, we discuss how to estimate the mean and asymptotic covariance matrix of $\boldsymbol{\pi}_n(\boldsymbol{\theta}|\mathcal{D})$, respectively.

3.3.1 Asymptotic Normality of $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$ —For convenience, we rewrite the full data posterior $\pi(\boldsymbol{\theta}|\mathcal{D})$ by $\pi_n(\boldsymbol{\theta}|\mathcal{D})$; i.e.,

$$\pi_n(\boldsymbol{\theta}|\mathcal{D}) \propto \pi_n(\boldsymbol{\theta})f(\mathcal{D}|\boldsymbol{\theta}),$$

where $f(\mathcal{D}|\boldsymbol{\theta}) = \prod_{i=1}^n f(x_i|\boldsymbol{\theta})$ denotes the likelihood function of $\mathcal{D}$, and $\pi_n(\boldsymbol{\theta})$ denotes the prior of $\boldsymbol{\theta}$ which may depend on the value of n .

The asymptotic normality of posterior distributions has long been studied in the literature. Walker (1969) presented a straightforward approach to the problem for i.i.d. observations. Later, this result was generalized to different statistical models and the conditions were also weakened, see e.g., Dawid (1970), Heyde and Johnstone (1979), and Chen (1985). Among those work, the conditions given in Chen (1985) are very general and flexible. For convenience, we shall present Chen's result as

$$\sqrt{n} \sum_0^{-1/2} (\boldsymbol{\theta}^{(n)} - \boldsymbol{\mu}_n) \xrightarrow{d} N(0, I), \quad (22)$$

where $\boldsymbol{\theta}^{(n)}$ denotes a generic sample of the full data posterior $\pi_n(\boldsymbol{\theta}|\mathcal{D})$, $\boldsymbol{\mu}_n$ denotes a local mode of $\pi_n(\boldsymbol{\theta}|\mathcal{D})$, $\xrightarrow{d}$ denotes convergence in distribution, I denotes an identity matrix, and

$$\sum_0 = n \left\{ -\partial^2 \log \pi_n(\boldsymbol{\theta}|\mathcal{D}) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T \right\}^{-1} \Big|_{\boldsymbol{\theta}=\boldsymbol{\mu}_n} = \left\{ -\partial^2 \left[\frac{1}{n} \log \pi_n(\boldsymbol{\theta}) + \frac{1}{n} \sum_{i=1}^n \log f(x_i|\boldsymbol{\theta}) \right] / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T \right\}^{-1} \Big|_{\boldsymbol{\theta}=\boldsymbol{\mu}_n}.$$

Here, we have assumed that $\pi_n(\boldsymbol{\theta}|\mathcal{D})$ satisfies appropriate conditions, to be specific, the conditions (E₁)–(E₅) given in Lemma A of the Supplemental Material. These conditions require that for each n , $\boldsymbol{\mu}_n$ is a local maximum of $\pi_n(\boldsymbol{\theta}|\mathcal{D})$ such that $\log \pi_n(\boldsymbol{\theta}|\mathcal{D}) / \boldsymbol{\theta} = \boldsymbol{\mu}_n = 0$ and Σ_0 is positive definite, and that for large n , $\pi_n(\boldsymbol{\theta}|\mathcal{D})$ becomes highly peaked and behaves like a normal kernel inside a small neighborhood of $\boldsymbol{\mu}_n$ and the probability outside the neighborhood is negligible. As pointed out by Chen (1985), the posterior $\pi_n(\boldsymbol{\theta}|\mathcal{D})$ can be multimodal, and $\boldsymbol{\mu}_n$ does not need to be the global maximum. However, the concentration condition must be satisfied at $\boldsymbol{\mu}_n$; that is, the probability outside a small neighborhood of $\boldsymbol{\mu}_n$ is negligible.

For BMH, $l_{m,n,k}(\boldsymbol{D}_s|\boldsymbol{\theta})$ works as the log-likelihood function, which, as shown in Sections 3.1 and 3.2, converges to $g_m(\mathcal{D}|\boldsymbol{\theta})$ in probability under appropriate conditions. Let $\pi_m(\boldsymbol{\theta})$ denote the prior distribution of $\boldsymbol{\theta}$. Then the posterior density function $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$ is approximately given by

$$\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D}) \dot{\propto} \pi_m(\boldsymbol{\theta})g_m(\mathcal{D}|\boldsymbol{\theta}) = \exp \left\{ m \left[\frac{1}{m} \log \pi_m(\boldsymbol{\theta}) + E \log f(X|\boldsymbol{\theta}) \right] \right\},$$

where $\dot{\propto}$ denotes “approximately proportional to”. Let $\tilde{l}_m(\boldsymbol{\theta}) = m \left[\frac{1}{m} \log \pi_m(\boldsymbol{\theta}) + E \log f(X|\boldsymbol{\theta}) \right]$. It is assumed that for each m ,

- (D₁) $\tilde{l}_m(\boldsymbol{\theta})$ is uniformly continuous on the parameter space Θ , and it has a unique global maximum and a finite number of local maxima.
- (D₂) At the global maximum of $\tilde{l}_m(\boldsymbol{\theta})$, denoted by $\boldsymbol{\mu}_m$, the following conditions are satisfied:
- i. $\dot{\tilde{l}}_m(\boldsymbol{\mu}_m) = \partial \tilde{l}_m(\boldsymbol{\theta}) / \partial \boldsymbol{\theta} |_{\boldsymbol{\theta}=\boldsymbol{\mu}_m} = 0$;
  - ii. $\partial^2 \tilde{l}_m(\boldsymbol{\theta}) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T$ is continuous on Θ , and $\ddot{\tilde{l}}_m(\boldsymbol{\mu}_m) = \partial^2 \tilde{l}_m(\boldsymbol{\theta}) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T |_{\boldsymbol{\theta}=\boldsymbol{\mu}_m}$ is negative definite.

The uniform continuity condition can be satisfied by restricting Θ to a large compact set, say, $\Theta = [10^{-100}, 10^{100}]^d$, where d is dimension of $\boldsymbol{\theta}$. As a practical matter, this is equivalent to set $\Theta = \mathbb{R}^d$. The maxima $\boldsymbol{\mu}_m$'s may be different for different values of m . As $m \rightarrow \infty$, $\frac{1}{m} \log \pi_m(\boldsymbol{\theta})$ tends to 0, and thus $\boldsymbol{\mu}_m$ converges to the maximum of $E \log f(X|\boldsymbol{\theta})$. It follows from Jensen's inequality that $E_{\boldsymbol{\theta}_*} \log(f(X|\boldsymbol{\theta})/f(X|\boldsymbol{\theta}_*)) \leq 0$ for any $\boldsymbol{\theta} \in \Theta$ and $\boldsymbol{\theta}_* \in \Theta_*$, where Θ_* denotes the set of true values of $\boldsymbol{\theta}$ and $E_{\boldsymbol{\theta}_*}$ denotes expectation with respect to $f(X|\boldsymbol{\theta}_*)$. That is, for any $\boldsymbol{\theta} \in \Theta$,

$$E_{\boldsymbol{\theta}_*} \log f(X|\boldsymbol{\theta}) \leq E_{\boldsymbol{\theta}_*} \log f(X|\boldsymbol{\theta}_*).$$

Moreover, this inequality is strict unless $P(f(X|\boldsymbol{\theta}) = f(X|\boldsymbol{\theta}_*)) = 1$. Hence, as $m \rightarrow \infty$, $\boldsymbol{\mu}_m$ will converge to $\boldsymbol{\theta}_*$. The uniqueness condition for the global maximum requires Θ_* is a singleton, i.e., $\boldsymbol{\theta}_*$ is unique. In the case that Θ_* is not a singleton, e.g., in mixture models, the BMH samples can still be used for model inference after applying a label switching procedure (see e.g., Stephens, 2000). Alternatively, one may impose some constraints on $\boldsymbol{\theta}$ such that Θ_* is a singleton.

Theorem 5 shows that under conditions (D₁) and (D₂), $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$ will converge to a normal density function. Its proof can be found in the Supplementary Material.

Theorem 5: Assume that conditions (D₁) and (D₂) hold for each $m > 0$. Then, as $m \rightarrow \infty$, we have

$$\sqrt{m} \sum_{i=0}^{\infty} (\boldsymbol{\theta}^{(i)} - \boldsymbol{\mu}_m)^{-1/2} \xrightarrow{d} N(0, I), \quad \text{as } m \rightarrow \infty, \quad (23)$$

where $\boldsymbol{\theta}^{(m)}$ denotes a generic sample of $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$, and $\hat{\Sigma}_0 = m[l_m''(\boldsymbol{\mu}_m)]^{-1}$.

Let $b(\boldsymbol{\theta})$ be a function of $\boldsymbol{\theta}$. Suppose that $b(\boldsymbol{\theta})/\boldsymbol{\theta}$ exists and is not 0. Then, by Delta method, we have

$$\sqrt{m}(b(\boldsymbol{\theta}^{(m)}) - b(\boldsymbol{\mu}_m)) \xrightarrow{d} N\left(0, \left(\frac{\partial b(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right)^T \hat{\Sigma}_0 \left(\frac{\partial b(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right)\right), \quad \text{as } m \rightarrow \infty. \quad (24)$$

Further, by the convergence of the averaged observed information

$-\frac{1}{n} \sum_{i=1}^n \partial^2 \log f(x_i|\boldsymbol{\theta}) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T$ to the Fisher information $-E(-\partial^2 \log f(x|\boldsymbol{\theta}) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T)$ under regularity conditions, we have

$$\|\hat{\Sigma}_0 - \Sigma_0\| \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty. \quad (25)$$

That is, Σ_0 can be estimated based on the BMH samples simulated from $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$.

As implied by (23) and (25), BMH has the capability to incorporate the whole data information into a single simulation run. Hence, it can have quite different performance from the D&C method. For the latter, suppose that the dataset has been divided into K subsets $\mathcal{D}_1, \dots, \mathcal{D}_K$, and each is of size m , i.e., $n = m \times K$. If m is reasonably large, for each subset the corresponding posterior distribution is approximately normal; that is,

$$\sqrt{m} \sum_{m,i}^{-1/2} (\boldsymbol{\theta}_i^{(m)} - \boldsymbol{\mu}_{m,i}) \xrightarrow{d} N(0, I),$$

where i indexes the i th subset, and $\boldsymbol{\mu}_{m,i}$ and $\Sigma_{m,i}$ denote, respectively, the mean and covariance matrix of the posterior distribution based on the subset $\mathcal{D}_i$. If m is small, $\boldsymbol{\mu}_{m,i}$'s and $\Sigma_{m,i}$'s can be quite different from others. It is true that $\boldsymbol{\mu}_{m,i}$ and $\Sigma_{m,i}$ will asymptotically lose their dependence on i when m becomes large, but this comes at a price of increasing computational cost. As shown by a simulated example in the Supplementary Material, this dependence can lose surprisingly slowly.

We note that the D&C method can be implemented in many different ways. For example, an interesting way is to dynamically incorporate the split samples into MH simulations. This leads to a BMH-like algorithm: Instead of bootstrap samples, we simply use consecutive chunks of the data as each subsample. For convenience, we call this algorithm the Divide-and-Combine-MH (DCMH) algorithm. Suppose that k chunks of the data are used at each iteration, where $k \leq K = n/m$. Then, as $n \rightarrow \infty$, $k \rightarrow \infty$ and $m/k \rightarrow 0$, we have

$$\frac{1}{k} \sum_{i=1}^k \log \tilde{f}(D_i | \boldsymbol{\theta}) - N_{m,n}(\mathcal{D} | \boldsymbol{\theta}) \xrightarrow{p} 0,$$

where $N_{m,n}(\mathcal{D} | \boldsymbol{\theta}) = \frac{1}{K} \sum_{j=1}^K h(D_j) = \frac{1}{K} \sum_{j=1}^K \log \tilde{f}(D_j | \boldsymbol{\theta})$. It follows from (3) that

$$N_{m,n}(\mathcal{D} | \boldsymbol{\theta}) = U_{m,n}(\mathcal{D} | \boldsymbol{\theta}) = \frac{m}{n} \sum_{i=1}^n \log f(X_i | \boldsymbol{\theta}),$$

which implies that DCMH can be as efficient as BMH. However, for the kernels that cannot be expressed as an independent sum, say, $h(D_j) = \prod_{i=1}^m f(X_i^{(j)} | \boldsymbol{\theta})$, BMH should be asymptotically more efficient than DCMH, as the inequality $\text{Var}(N_{m,n}) \leq \text{Var}(U_{m,n})$ generally holds.

3.3.2 Estimation of the Mean of $\pi_n(\boldsymbol{\theta} | \mathcal{D})$ —First, we explore the relationship between $b(\boldsymbol{\mu}_n)$ and $b(\boldsymbol{\mu}_m)$. Consider the standard Laplace approximation for posterior means, see e.g., Lindley (1961, 1980), Kass et al. (1990) and Miyata (2004) for the details. Given a prior π , a log-likelihood $\log p(\mathbf{x} | \boldsymbol{\theta})$, a positive function ξ , and a real function h , we define h_n and ρ by $h_n(\boldsymbol{\theta}) = -\log p(\mathbf{x} | \boldsymbol{\theta})/n - \log \xi/n$ and $\rho = \pi/\xi$. Suppose that we are interested in estimating the posterior mean $E_{\pi_n}[b(\boldsymbol{\theta})]$ for an integrable function $b(\boldsymbol{\theta})$, where $E_{\pi_n}[\cdot]$ denotes expectation with respect to the posterior $\pi_n(\boldsymbol{\theta} | \mathcal{D})$. Under regularity conditions, $E_{\pi_n}[b(\boldsymbol{\theta})]$ can be approximated as follows:

$$\begin{aligned} E_{\pi_n}[b(\boldsymbol{\theta})] &= \frac{\int_{\Theta} b(\boldsymbol{\theta}) \rho(\boldsymbol{\theta}) \exp\{-nh_n(\boldsymbol{\theta})\} d\boldsymbol{\theta}}{\int_{\Theta} \rho(\boldsymbol{\theta}) \exp\{-nh_n(\boldsymbol{\theta})\} d\boldsymbol{\theta}} \\ &= b(\hat{\boldsymbol{\theta}}) + \frac{1}{n} \sum_{ij} b_i h^{ij} \left\{ \frac{\rho_i(\hat{\boldsymbol{\theta}})}{\rho(\hat{\boldsymbol{\theta}})} - \frac{1}{2} \sum_{rs} h^{rs} h_{rsj} \right\} + \frac{1}{2n} \sum_{ij} h^{ij} b_{ij} + O(n^{-2}), \end{aligned} \quad (26)$$

where $\rho_i(\hat{\boldsymbol{\theta}}) = \partial \rho(\hat{\boldsymbol{\theta}}) / \partial \theta_i$, $b_i = \partial b(\hat{\boldsymbol{\theta}}) / \partial \theta_i$, $b_{ij} = \partial^2 b(\hat{\boldsymbol{\theta}}) / \partial \theta_i \partial \theta_j$, $h_{rsj} = \partial^3 h_n(\hat{\boldsymbol{\theta}}) / \partial \theta_r \partial \theta_s \partial \theta_j$, h^{ij} is the component of the matrix $[\partial^2 h_n(\hat{\boldsymbol{\theta}}) / \partial \theta_i \partial \theta_j]^{-1}$, and θ_j denotes the j th component of $\boldsymbol{\theta}$. There are two special cases for the choices of ξ and ρ : If $\xi = 1$ and $\rho = \pi$, then $\hat{\boldsymbol{\theta}}$ becomes the MLE; and if $\xi = \pi$, then $\rho = 1$ and $\hat{\boldsymbol{\theta}}$ becomes the posterior mode.

Suppose that $\boldsymbol{\theta}$ is subject to the following prior:

$$\pi_m(\boldsymbol{\theta}) \propto [\pi_n(\boldsymbol{\theta})]^{m/n}. \quad (27)$$

Note that this prior setting facilitates the following theoretical analysis, but is not an essential requirement. By condition (A), the central limit theorem implies that ρ

$$-h_n(\theta) = E_h(\theta) + \sigma_\theta Z / \sqrt{n},$$

where $E_h(\theta) = -E \log f(X|\theta) - \log(\xi)/n$, σ_θ denotes a function of θ (independent of n), and Z is a standard Gaussian random variable. Then it follows from (26) that the posterior mean $E_{\tilde{\pi}_m}[b(\theta)]$, which is defined with respect to the posterior $\tilde{\pi}_m(\theta|\mathcal{D})$, can be approximated by

$$\begin{aligned} E_{\tilde{\pi}_m}[b(\theta)] &= \frac{\int_{\Theta} b(\theta) \rho(\theta) \exp\{-m E_h(\theta)\} d\theta}{\int_{\Theta} \rho(\theta) \exp\{-m E_h(\theta)\} d\theta} \\ &= \frac{\int_{\Theta} b(\theta) \rho(\theta) \exp\{-m[h_n(\theta) + \sigma_\theta Z / \sqrt{n}]\} d\theta}{\int_{\Theta} \rho(\theta) \exp\{-m[h_n(\theta) + \sigma_\theta Z / \sqrt{n}]\} d\theta} \\ &= b(\hat{\theta}) + \frac{1}{m} \sum_{ij} b_i (h^{ij} + O(1/\sqrt{n})) \left\{ \frac{\rho_j(\hat{\theta})}{\rho(\hat{\theta})} - \frac{1}{2} \sum_{rs} (h^{rs} + O(1/\sqrt{n})) (h_{rsj} + O(1/\sqrt{n})) \right\} + \frac{1}{2m} \sum_{ij} (h^{ij} + O(1/\sqrt{n})) b_{ij} + O(m^{-2}), \\ &= b(\hat{\theta}) + \frac{1}{m} \sum_{ij} b_i h^{ij} \left\{ \frac{\rho_j(\hat{\theta})}{\rho(\hat{\theta})} - \frac{1}{2} \sum_{rs} h^{rs} h_{rsj} \right\} + \frac{1}{2m} \sum_{ij} h^{ij} b_{ij} + O\left(\frac{1}{m^2}\right) + O\left(\frac{1}{m\sqrt{n}}\right), \end{aligned}$$

where ρ_j , b_j , b_{ij} , h_{rsj} and h^{ij} are as defined in (26). Under the setting $m = O(n^\gamma)$ for $0 < \gamma < 1/2$, given in Theorem 3, we have

$$E_{\tilde{\pi}_m}[b(\theta)] = b(\hat{\theta}) + \frac{1}{m} \sum_{ij} b_i h^{ij} \left\{ \frac{\rho_j(\hat{\theta})}{\rho(\hat{\theta})} - \frac{1}{2} \sum_{rs} h^{rs} h_{rsj} \right\} + \frac{1}{2m} \sum_{ij} h^{ij} b_{ij} + O(m^{-2}). \quad (28)$$

Hence, $E_{\tilde{\pi}_m}[b(\theta)] \rightarrow b(\hat{\theta})$ as $m \rightarrow \infty$. As n goes to infinity, m will also go to infinity, then it follows from (26) and (28) that

$$\|E_{\tilde{\pi}_m}[b(\theta)] - E_{\pi_n}[b(\theta)]\| \rightarrow 0, \quad \text{as } m, n \rightarrow \infty. \quad (29)$$

and, by setting $\xi = \pi$ and $\rho = 1$ in (26) and (28),

$$\|b(\mu_m) - b(\mu_n)\| \rightarrow 0, \quad \text{as } m, n \rightarrow \infty. \quad (30)$$

Equation (29) suggests that we can use the sample average

$$E_{\tilde{\pi}_m}[\widehat{b(\theta)}] = \frac{1}{T} \sum_{t=1}^T b(\theta_t), \quad (31)$$

to estimate $E_{\pi_n}[b(\theta)]$, where $\theta_1, \dots, \theta_T$ denote T samples simulated by BMH from the approximate posterior $\tilde{\pi}_m(\theta|\mathcal{D})$. It follows from (29) and the property of MCMC that

$E_{\hat{\pi}_m}[\widehat{b(\theta)}]$ provides a consistent estimator for $E_{\pi_n}[b(\theta)]$. Further, it follows from (26) and the consistency of MLE that $E_{\hat{\pi}_m}[\widehat{b(\theta)}]$ is consistent for $b(\theta_*)$; that is,

$$E_{\hat{\pi}_m}[\widehat{b(\theta)}] \xrightarrow{p} b(\theta_*), \quad \text{as } m \rightarrow \infty. \quad (32)$$

In practice, m cannot be very large for the reason of computational efficiency. As implied by (28), we can improve the accuracy of the estimator of $b(\theta_*)$ using an extrapolation method by fitting the linear regression

$$E_{\hat{\pi}_m}[\widehat{b(\theta)}] = \beta_0 + \beta_1/m + \varepsilon,$$

for a small set of m , where $1/m$ works as the explanatory variable, and ε is the normal random error. Then $\hat{\beta}_0$, the least square estimator of β_0 , will serve as an estimator of $b(\theta_*)$, which corresponds to the limit $m \rightarrow \infty$.

3.3.3 Estimation of the Covariance Matrix of $\pi_n(\theta|\mathcal{D})$ —In addition to the mean of the posterior $\pi_n(\theta|\mathcal{D})$, the asymptotic covariance matrix of $\pi_n(\theta|\mathcal{D})$ can also be simply estimated based on the BMH samples. It follows from (23) and (25) that $m\hat{\Sigma}_m$ provides a consistent estimator of Σ_0 , where $\hat{\Sigma}_m$ denotes the covariance matrix of $\theta^{(m)}$ calculated based on the BMH samples. In summary, BMH provides a simple way to asymptotically integrate the whole data information into a single simulation run and thus a convenient way for Bayesian analysis of big data.

4 Tempering BMH for Bayesian Neural Network Classifiers

Like the MH algorithm for small data Bayesian computation, BMH can also be used as a basic building block for constructing MCMC algorithms for big data Bayesian computation. In this section, we propose a tempering BMH algorithm, which can be viewed as a BMH-version of the parallel tempering algorithm (Geyer, 1991; Hukushima and Nemoto, 1996), with applications to Bayesian neural network (BNN) classifiers.

This section is organized as follows. Section 4.1 describes the Bayesian neural network classifier. Section 4.2 describes the tempering BMH algorithm and its parallel implementation. Section 4.3 reports the numerical results for a real data example.

4.1 Bayesian Neural Network Classifier

Suppose we have a group of observations $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, where $x_i \in \mathbb{R}^p$ are explanatory variables, y_i indicates the class membership, and n denotes the number of observations. For the problems of binary classes, y_i is represented by a single digit 0 or 1. For the problems of multiple classes, y_i is represented by a binary vector with the class membership indicated by the position of the only element of 1. If there is not any structural information assumed for the classifier, a feed-forward neural network, which is also known as the multiple-layer

perceptron (MLP), can be used due to its universal approximation ability (Cybenko, 1989; Hornik, 1991). A one hidden layer MLP model (illustrated by Figure 1) with weight eliminations approximates y_i by a group of q functions (corresponding to q output units) of the form

$$g_j(\mathbf{x}_i) = \phi_o \left[\alpha_{j0} I_{\alpha_{j0}} + \sum_{s=1}^p x_{is} \alpha_{js} I_{\alpha_{js}} + \sum_{l=1}^M \beta_{jl} I_{\beta_{jl}} \phi_h \left(\gamma_{l0} I_{\gamma_{l0}} + \sum_{s=1}^p x_{is} \gamma_{ls} I_{\gamma_{ls}} \right) \right], \quad j=1, \dots, q, \quad (33)$$

where M is the number of hidden units, p is the number of input units, $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$, $\boldsymbol{\theta} = \{\alpha_{js}, \beta_{jl}, \gamma_{ls}; j=1, \dots, q; l=1, \dots, M; s=0, 1, \dots, p\}$ is the collection of the MLP connection weights, I_{ζ} is an indicator function which indicates the effectiveness of the connection ζ , and $\phi_h(\cdot)$ and $\phi_o(\cdot)$ are activation functions of the hidden units and the output units, respectively. The weights α_{js} , β_{jl} and γ_{ls} correspond to the connections from the s th input unit to the j th output unit, from the l th hidden unit to the j th output unit, and from the s th input unit to the l th hidden unit, respectively. In (33), the bias unit is treated as a special input unit, indexed by 0 and having a constant input of 1. The connections from input units to the output units are called shortcut connections. Sigmoid and hyperbolic tangent functions are two popular choices for the activation functions. In this paper, we set $\phi_h(\cdot) = \tanh(\cdot)$, which ensures that the output of a hidden unit is 0 if all connections to the hidden unit from the input units have been eliminated, and thus the hidden unit can also be eliminated from the network without any effect on the outputs. The function $\phi_o(\cdot)$ can be set according to the problem under consideration. For the problems of binary classes, $\phi_o(\cdot)$ can be set to the sigmoid function $\phi_o(z) = 1/(1 + \exp(-z))$. For the problems of multiple classes, $\phi_o(\cdot)$ can be set to the softmax function $\exp(z_j) / \sum_{i=1}^q \exp(z_i)$ which ensures that the outputs sum to one.

Let $\Lambda = \{\alpha_{js}, \beta_{jl}, \gamma_{ls}; j=1, \dots, q; l=1, \dots, M; s=0, 1, \dots, p\}$ denote the structure of a network, and let $\boldsymbol{\theta}_{\Lambda}$ denote the subset of $\boldsymbol{\theta}$ corresponding to the nonzero elements of Λ . Thus, the model (33) can be specified by the tuple $(\Lambda, \boldsymbol{\theta}_{\Lambda})$. For simplicity, in what follows we will denote the model by $\boldsymbol{\theta}_{\Lambda}$ only. Then the likelihood function of the BNN classifier is given by

$$f(\mathcal{D}|\boldsymbol{\theta}_{\Lambda}) = \prod_{i=1}^n g_1(\mathbf{x}_i, \boldsymbol{\theta}_{\Lambda})^{y_i} (1 - g_1(\mathbf{x}_i, \boldsymbol{\theta}_{\Lambda}))^{1-y_i} \quad (34)$$

for the problems of two classes, and

$$f(\mathcal{D}|\boldsymbol{\theta}_\Lambda) = \prod_{i=1}^n \prod_{j=1}^q g_j(\mathbf{x}_i, \boldsymbol{\theta}_\Lambda)^{y_{ij}} \quad (35)$$

for the problems of multiple classes.

To conduct Bayesian analysis for the BNN model (33), we follow Liang (2003) to specify an independent Gaussian prior for each element of $\boldsymbol{\theta}_\Lambda$, i.e., $\boldsymbol{\theta}_\Lambda \sim N(0, \sigma^2 I_{|\Lambda|})$, where $|\Lambda|$ denotes the number of elements of $\boldsymbol{\theta}_\Lambda$, and $I_{|\Lambda|}$ denotes an identity matrix of order $|\Lambda|$. Let Λ be subject to a prior probability that is proportional to the mass put on $|\Lambda|$ by a truncated Poisson with rate λ , i.e.,

$$P(\Lambda) = \begin{cases} \frac{1}{Z} \frac{\lambda^{|\Lambda|}}{|\Lambda|!}, & |\Lambda| = 3, 4, \dots, U, \\ 0, & \text{otherwise,} \end{cases}$$

where $U = (p+1)(M+q) + MQ$ is the total number of connections of the full model, 3 is the minimum number of connections allowed in BNN, and Z is the normalizing constant. Here σ^2 and λ are hyperparameters to be specified by the user. In this paper, we set $\sigma^2 = 5$, which favors small connection weights and thus enhances the generalization ability of the BNN. The value of λ can be set according to the complexity of the problem under consideration. In general, the value of λ should increase with the complexity of the problem.

Motivated by LASSO (Tibshirani, 1996), a Laplace prior can also be specified for $\boldsymbol{\theta}_\Lambda$. Under the Laplace prior, the resulting neural networks can be more sparse. However, as the main purpose of this example is to demonstrate the use of tempering BMH, the Gaussian prior is still used here.

Let $(\Lambda^*, \boldsymbol{\theta}_\Lambda^*)$ denote the true value of $(\Lambda, \boldsymbol{\theta}^*)$ for a problem under study. Since the hidden units can be relabeled without affecting the output of the neural network, $(\Lambda^*, \boldsymbol{\theta}_\Lambda^*)$ is not unique. To make it unique, we suggest to impose some constraints to fix the order of the hidden units. For example, we can impose on β_{ji} 's the following constraints $\beta_{11} \ \beta_{12} \ \cdots \ \beta_{1M} \ \beta_{21} \ \beta_{22} \ \cdots \ \beta_{2M} \ \dots, \beta_{q1} \ \beta_{q2} \ \cdots \ \beta_{qM}$ (in an importance decreasing order). Since the connection selection will be done in our BNN, certain coefficients β_{ji} 's will be set to zero. Hence, multiple constraints may be needed to fix the order of the hidden units. We note that if our aim is at prediction instead of BNN structure inference, such constraints are not necessary.

4.2 Tempering BMH for Learning BNN with Big Data

BNN learning suffers from two difficulties for big data problems. First, a complete scan of all observations for each iteration is prohibitive. When the network is large, the likelihood evaluation can be time consuming even for a single observation, as which involves many long-operations such as $\exp(\cdot)$. BMH is an obvious solution to alleviate this difficulty. Second, the energy landscape of BNN can be rugged, which can contain a multitude of local

energy minima separated by high energy barriers. Here the energy refers to the negative of the log-posterior distribution in terms of physics. A straightforward application of BMH to this problem can lead to the so-called local trap problem; that is, the simulation will get trapped into a local energy minimum indefinitely, rendering an ineffective inference of the model. To overcome this difficulty, we propose the tempering BMH algorithm, which is a combination of parallel tempering and the BMH algorithm. The tempering BMH algorithm can be described as follows.

Suppose we would like to sample from a posterior distribution $f(\boldsymbol{\theta}|\mathcal{D}) \propto \exp(-\mathbb{H}(\boldsymbol{\theta}|\mathcal{D}))$, where $\mathbb{H}(\cdot)$ denotes the energy function. In tempering BMH, a sequence of distributions $f_1(\boldsymbol{\theta}), \dots, f_\pi(\boldsymbol{\theta})$ are first constructed: $f_i(\boldsymbol{\theta}) \propto \exp\{-\mathbb{H}(\boldsymbol{\theta}|\mathcal{D})/\tau_i\}$, $i = 1, \dots, \pi$, where τ_i is called the temperature of $f_i(\cdot)$. The temperature sequence $\boldsymbol{\tau} = (\tau_1, \dots, \tau_\pi)$ forms a ladder with $\tau_1 > \dots > \tau_\pi \equiv 1$. Issues related to the choice of the temperature ladder can be found in Miasojedow *et al.* (2012). Let $\boldsymbol{\theta}^i$ denote a sample simulated from $f_i(\cdot)$. The samples $\boldsymbol{\theta}^1, \dots, \boldsymbol{\theta}^\pi$ form a population denoted by $\Xi = \{\boldsymbol{\theta}^1, \dots, \boldsymbol{\theta}^\pi\}$. We further assume that all samples in the same population Ξ are mutually independent, and they have the joint density

$$f(\Xi) \propto \exp\left\{-\sum_{i=1}^n \mathbb{H}(\boldsymbol{\theta}^i|\mathcal{D})/\tau_i\right\}. \quad (36)$$

Tempering BMH works by sampling from $f(\Xi)$ with the full-data energy $\mathbb{H}(\boldsymbol{\theta}|\mathcal{D})$ approximated by the average of the energies $\mathbb{H}(\boldsymbol{\theta}^i|D_1), \dots, \mathbb{H}(\boldsymbol{\theta}^i|D_k)$ calculated from the bootstrap samples $D_1, \dots, D_k$. The posterior samples of interest can be collected at the target temperature level. Tempering BMH consists of two operators, local updating and exchange.

For BNN, the local updating operators include “birth”, “death”, and “Metropolis” moves as defined in Liang (2003). The “birth” move is to add a connection to the current network, the “death” move is to delete a connection from the current network, and the “Metropolis” move is to update the connection weights while keeping the current network structure unchanged. The exchange operator is to swap two samples, say $\boldsymbol{\theta}^i$ and $\boldsymbol{\theta}^j$, without changing the associated temperature ladder. The acceptance probability of these operators can be written as in Liang (2003), but with the full-data energy function $\mathbb{H}(\boldsymbol{\theta}|\mathcal{D})$ replaced by the average energy function calculated from bootstrap samples. Refer to Liang (2003) for more details of the local updating operators.

In summary, each iteration of the tempering BMH algorithm consists of the following steps:

Algorithm 2

1. Draw k bootstrap samples, $D_1, \dots, D_k$, with/without replacement from the entire training dataset $\mathcal{D}$.
2. Try to update each sample of the current population $(\boldsymbol{\theta}_t^1, \dots, \boldsymbol{\theta}_t^\pi)$ by the local updating operators, where t indexes iterations, and the energy function is

calculated by averaging the energy values calculated from the bootstrap samples $D_1, \dots, D_k$.

3. Try to exchange θ_t^i with θ_t^j for $n-1$ pairs (i, j) with i being sampled uniformly on $\{1, \dots, n\}$ and $j = i \pm 1$ with probability ω_{ij} , where $\omega_{i,i+1} = \omega_{i,i-1} = 0.5$ for $1 < i < \pi$ and $\omega_{1,2} = \omega_{\pi,\pi-1} = 1$.

The validity of this algorithm follows from the validity of BMH with repeated bootstrap samples. From the viewpoint of population MCMC, both the local updating and exchange steps can be viewed as specific MH updates with respect to the joint density $f(\Xi)$. In this algorithm, the user-specified parameters include the population size π , the temperature ladder τ , and some parameters used in the local updating operators. In this paper, we set $\pi = 20$ and set the temperature ladder to be a series of equal harmonic difference, i.e., $1/\tau_{i+1} - 1/\tau_i = c$ for a constant c , with $t_1 = 5$ and $\tau_\pi = 1$. The parameters in the local updating operators are tuned for different problems such that the local updating moves to have a moderate acceptance rate, e.g., a value between 0.2 and 0.4 as suggested by Gelman et al. (1996).

The implementation of this algorithm is depicted in Figure 2, where the energy function is calculated in parallel. The parallel implementation involves two message passing interface (MPI) functions, namely, broadcast and reduce. The former is to broadcast a message from the master processor to slave processors, and the latter is to reduce values on all processors to a single value on the master processor. Our numerical results indicate that the parallel implementation of BMH can result in a great saving of computational time for some problems.

4.3 Forest Cover Type Data

The goal of this study is to predict forest cover types from cartographic variables in the forested areas with minimal human-caused disturbances. The data were taken from four wilderness areas located in the Roosevelt National Forest of northern Colorado. It consisted of 581,012 observations. Each observation was obtained from the US Geological Survey (USGS) digital elevation model data based on the raster cells of size 30 meters by 30 meters, and it consisted of 54 cartographic explanatory variables including 10 quantitative variables, 4 binary wilderness area variables, and 40 binary soil type variables. These observations have been classified into seven classes according to their cover types. The respective class sizes are 211840, 283301, 35754, 2747, 9493, 17367, and 20510. The full dataset is available at UCI machine learning repository <http://archive.ics.uci.edu/ml/>.

In pre-processing the data, we first normalize each of the quantitative variables to have mean 0 and variance 1, and then divide the full dataset into two parts, training data and test data. As discussed in Liang (2003), the normalization makes the magnitude of the weights connected to different input units comparable and thus facilitates the BNN training. Since the proportions of different classes of observations are unbalanced, which ranges from 0.47% to 48.76%, the training data is constructed by randomly drawing 75% observations from each class of observations. The remaining observations were used as the test data. As a result, the training data consists of 435,759 observations, and the test data consists of 145,253 observations.

For this example, we tried a BNN with 54 input units, 10 hidden units, and 7 output units. Under this setting, the full BNN contains 1005 connections. The hyperparameter λ was set to 50. Tempering BMH was then applied to train the BNN with $k = 20$ and $k = 40$ processors. For each value of k , the algorithm was run 10 times independently. In the first five runs the bootstrap samples were drawn with replacement, and in the next five runs the bootstrap samples were drawn without replacement. Each run consisted of 25,500 iterations, where the first 500 iterations were discarded for the burn-in process and the remaining 25,000 iterations were used for inference. The bootstrap sample sizes for the seven types of observations are 794, 1062, 268, 52, 107, 130, and 154, respectively; and the corresponding subsampling rates are 0.5%, 0.5%, 1.0%, 2.5%, 1.5%, 1% and 1%, respectively. Although the subsampling rates for some small classes of observations are high, the aggregated subsampling rate for the training data is only 0.59%. Such a low subsampling rate implies that tempering BMH can potentially result in a great saving of computational time for this example. We note that the above settings are not necessarily optimal. They are just used to demonstrate that tempering BMH can be effectively used to train BNNs for big data problems. In fact, for unbalanced classification problems, how to choose the proportions for different classes of observations in forming the training dataset is an interesting problem.

Table 1 summarizes the training and prediction results of BNN for this example. As shown in the table, the prediction rate is consistently higher than 72% for all settings of tempering BMH. We note that this rate is significantly higher than the rate 70.52% (with a 95% confidence interval of 70.26–70.80%) obtained by Blackard and Dean (1999) using a non-Bayesian neural network model. Blackard and Dean (1999) employed a neural network with 120 hidden units, which consists of 7447 connections and is much larger than the networks we used. On average, the networks we used consisted of only about 260 connections as reported in Table 1. Regarding the CPU time, we note that its variation is mainly due to the convergence time of different runs. Since our simulations always start with the fully connected networks, a run will cost more CPU time if it converges slowly. As the secondary source, the variation can also be due to the between-processor communications.

To show the advantage of tempering BMH for this example, we have applied parallel tempering to train the BNN with a single-threaded simulation on an Intel Nehalem simulation server which uses Intel(R) Xeon(R) CPU X5560@2.80GHz. At each local updating step of parallel tempering, the whole training dataset is scanned once. Hence, this algorithm runs extremely slow. The first 5536 iterations have taken 688 CPU hours, although the processor used in the Intel Nehalem server is much faster (approximately 2.5 times) than the processor used in the cluster machine. To finish 25000 iterations, it will take about 3000 CPU hours (125 days) on the Intel Nehalem server. Compared 3000 hours to 30 hours, it shows a great advantage of the parallelized BMH algorithm for big data problems.

5 BMH on Distributed Architectures

In this paper, we have illustrated the BMH algorithm using three examples for which the full dataset can be stored in the memory of a single processor. For this kind of problems, the bootstrap samples can be independently drawn from each processor. In practice, we may encounter the problems for which the whole dataset is so big that it has to be stored in the

memories of multiple processors. For these problems, we propose two subsampling schemes, subsampling with data exchange and exclusive subsampling, which are described as follows.

5.1 Subsampling with Data Exchange

Suppose that the whole dataset has been divided into k “random” subsets, with each being of about the same size and stored in the memory of a single processor. These subsets are called random subsets because each forms a subsample of the whole data. This can be done by running a Map-Reduce job (Dean and Ghemawat, 2008) for initial data subsetting and distributing, which is essentially to randomly assign each observation to one of k processors of the distributed architecture. Thus, at the initial iteration of BMH, the bootstrap samples, in

particular, $\binom{n}{m}$ -bootstrapping, can be drawn in parallel with each from a single processor. Although the bootstrap samples obtained in this way are not mutually independent (they are mutually exclusive), they are identically distributed and ready to be applied to the BMH algorithm. To maintain this property for successive iterations, we update the random subsets at each iteration in the following procedure. Let $\mathcal{D}_i$ denote the i th subset, $i = 1, \dots, k$. One iteration of the following contains the following steps.

Algorithm 3: For each i ,

1. *pick j from the set $\{1, 2, \dots, i-1, i+1, \dots, k\}$ at random;*
2. *exchange M randomly selected observations between $\mathcal{D}_i$ and $\mathcal{D}_j$, where M can be a pre-specified or random number.*

It follows from the standard theory of MCMC (see e.g. Geyer, 1991) that the above procedure ensures each subset stored in a single processor to be a random subset of the

whole dataset. This implies that the $\binom{n}{m}$ -bootstrap samples can be drawn in parallel within single processors at each iteration of BMH. We suspect that the m/n -bootstrap samples drawn in this way may be more preferred than those drawn from the full dataset due to the less correlation between them. If each subset is too big to be stored in the memory of a single processor, a swapping operation between the memory and hard disc of each processor can be added in the above procedure. As implied by the following exclusive subsampling scheme, which corresponds to the case $M = 0$, the choice of M will not affect the validity of BMH. However, an appropriate value of M should improve the mixing of the data between different processors and thus the convergence of BMH. We suggest to determine the value of M by trial-and-error for real problems.

To demonstrate the performance of this algorithm, we consider a simulation example, where $n = 200,000$, $k = 100$, and $\mathcal{D}_i$ initially contains 2,000 points uniformly generated from the rectangle $[(i-1), i] \times [0, 100]$. We expect that after a small number of iterations of the algorithm, the points in each $\mathcal{D}_i$ are uniformly distributed in $[0, 100] \times [0, 100]$. Figure 3 shows that after about 10 iterations, the samples in each $\mathcal{D}_i$ have been roughly uniformly distributed in $[0, 100] \times [0, 100]$. We set $M = 500$ in the simulation.

5.2 Exclusive Resampling

As in the scheme of subsampling with data exchange, we assume that the whole dataset has been divided into k “random” subsets with each being of about the same size and stored in the memory of a single processor. The scheme of exclusive subsampling proposes to draw bootstrap samples independently from each subset, which avoids data exchange between different processors. The effectiveness of this scheme can be justified as follows.

Let $\mathcal{D}_1, \dots, \mathcal{D}_k$ denote the k random subsets, let $n_1, \dots, n_k$ denote their respective sizes, and

let D_i denote a bootstrap sample drawn from $\mathcal{D}_i$. Let $n = \sum_{i=1}^k n_i$ and $n_{\min} = \min_{1 \leq i \leq k} n_i$.

Define $\mathbf{D}_s = (D_1, \dots, D_k)$, and define $U_{m,n_i}(\mathcal{D}_i|\boldsymbol{\theta}) = \binom{n_i}{m}^{-1} \sum_{D_i \in \mathfrak{D}} \log \tilde{f}(D_i|\boldsymbol{\theta})$, where $\mathfrak{D}$

denotes the space of D_i and it contains all the possible $\binom{n_i}{m}$ subsamples of size m . Since $D_1, \dots, D_k$ are mutually independent, it follows from Chebyshev’s weak law of large numbers (see e.g., Serfling, 1980, p.27) that under condition (A), we have

$$l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - \frac{1}{k} \sum_{i=1}^k U_{m,n_i}(\mathcal{D}_i|\boldsymbol{\theta}) \xrightarrow{p} 0, \quad \text{as } k \rightarrow \infty \text{ and } n_{\min} \rightarrow \infty,$$

provided that as $k \rightarrow \infty$, $m/k \rightarrow 0$ holds. By noting the following result

$$\text{Var} \left(\frac{1}{k} \sum_{i=1}^k U_{m,n_i}(\mathcal{D}_i|\boldsymbol{\theta}) \right) = \frac{m^2}{k} \left(\sum_{i=1}^k \frac{1}{n_i} \right) \text{Var}(\log f(X|\boldsymbol{\theta})) \leq \frac{m^2}{n_{\min}} \text{Var}(\log f(X|\boldsymbol{\theta})),$$

we have, as $m^2/n_{\min} \rightarrow 0$ and $n_{\min} \rightarrow \infty$,

$$l_{m,n,k}(\mathbf{D}_s|\boldsymbol{\theta}) - \log(g_m(\mathcal{D}|\boldsymbol{\theta})) \xrightarrow{p} 0.$$

Then Theorem 2 can be established for the exclusive $\binom{n}{m}$ -bootstrapping scheme. For the exclusive m/n -bootstrapping scheme, the convergence can be established in a similar way.

6 Discussion

In this paper, we have proposed the BMH algorithm as a basic MCMC algorithm for Bayesian analysis of big data. The BMH algorithm is workable on parallel and distributed architectures and avoids repeated scans of the full dataset in iterations, and is thus feasible for big data problems. Compared to the popular divide-and-combine method, BMH is generally more efficient as it can asymptotically integrate the whole data information into a single simulation run. The BMH algorithm is very flexible. Like the Metropolis-Hastings

algorithm, it can serve as a basic building block for developing advanced MCMC algorithms that are feasible for big data problems. This has been illustrated in the paper by the tempering BMH algorithm, which is a combination of parallel tempering and the BMH algorithm. BMH provides a simple yet effective way of uncertainty quantification for big data problems.

In Section 3.3, we study the asymptotic normality of $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$ under the assumption both m and n are sufficiently large. In practice, since both m and n are finite, we may encounter the problems for which $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$ is multimodal even if Θ_* is a singleton. In this case, one may re-run BMH with a larger value of m . Alternatively, one may try an importance sampling procedure for inference of $\pi_n(\boldsymbol{\theta}|\mathcal{D})$. It is easy to see that, under the prior setting (27), $\log \tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$ is approximately reduced to a fraction of $\log \pi_n(\boldsymbol{\theta}|\mathcal{D})$, i.e., $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D}) \propto [\pi_n(\boldsymbol{\theta}|\mathcal{D})]^{m/n}$. Let $\hat{\pi}_m(\boldsymbol{\theta})$ denote the fitted parametric density function for the BMH samples. For example, one may use the software MCLUST (Fraley and Raftery, 2002) to fit the samples with a mixture Gaussian distribution. When the dimension of $\boldsymbol{\theta}$ is low, the samples can also be fitted with a nonparametric density. Then one can weight each sample by

$$w(\boldsymbol{\theta}) = [\hat{\pi}_m(\boldsymbol{\theta})]^{n/m-1},$$

and estimate $E_{\pi_n} b(\boldsymbol{\theta})$ by

$$E_{\pi_n}[\widehat{b(\boldsymbol{\theta})}] = \frac{\sum_{i=1}^T w(\boldsymbol{\theta}_i) b(\boldsymbol{\theta}_i)}{\sum_{i=1}^T w(\boldsymbol{\theta}_i)}, \quad (37)$$

where $b(\boldsymbol{\theta})$ is integrable, and $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_T$ denotes T approximately independent samples generated by BMH. When the ratio n/m is large, the estimator $E_{\pi_n}[\widehat{b(\boldsymbol{\theta})}]$ will be dominated by very a few components of $\hat{\pi}_m(\boldsymbol{\theta})$.

Let T denote the number of iterations of BMH. Then the overall computational complexity of BMH can be expressed as $\mathcal{O}(mkT)$, which is the same for both subsampling schemes. Note that the computational complexity of BMH does not directly depend on n , although m , k and T can all increase with n . As shown in Section 3, for BMH, we can set $m = \mathcal{O}(n^\gamma)$ and $k = \mathcal{O}(n^{\gamma+\varepsilon_0})$ for $\gamma < 1/2$ and any $\varepsilon_0 > 0$. Hence, the overall computational complexity of BMH is $\mathcal{O}(n^{2\gamma+\varepsilon_0}T)$. In a parallel implementation, the time complexity of BMH is $\mathcal{O}(n^\gamma T)$. For a parallel implementation of the MH algorithm on the full data, the time complexity is $\mathcal{O}(n^{1-\gamma-\varepsilon_0}T)$ if $n^{\gamma+\varepsilon_0}$ nodes are used. Since γ usually takes a very small value and ε_0 is nearly zero, BMH can be much faster than MH, even when the latter is implemented in parallel. Here we have assumed that BMH and MH have the same convergence time, which might be true only approximately.

BMH aims to provide a numerical approximation to the full data posterior for big data problems. The BMH approximation preserves all features of the full data posterior, such as marginal and correlation structures, which can be inferred from its samples. As shown in

Figure 2 and Figure 4 of the Supplemental Material, the approximation can be rather accurate (up to a known scale factor). We note that the variational Bayesian method also seeks an approximation to the posterior, but in an analytic form. To be precise, variational Bayes is to approximate the posterior $\pi_{\theta}(\boldsymbol{\theta}|\mathcal{D})$ using a joint distribution which can be factored into a product of marginal distributions, where the “closeness” of approximation is measured using the Kullback-Leibler divergence. As pointed out by Fox and Roberts (2012), this method approximates the true posterior distribution jointly, and the individual marginals can be poor approximations to the true marginals; hence, it is inappropriate to use the approximated distribution to explore the marginal or correlation structure of the true posterior distribution. We also note that variation Bayes is not designed for big data problems. It can be very time consuming to solve for the parameters of the approximated distribution when the data size is big, where an EM algorithm-like iterative procedure is often involved.

The BMH algorithm proposes to replace the full data log-likelihood by a Monte Carlo average of the log-likelihoods that are calculated in parallel from multiple bootstrap samples. As an alternative strategy, one may try to replace the likelihood function by its Monte Carlo average in simulations. Although, in theory, this averaged likelihood method may go through under suitable conditions, our numerical results show that it can be much less efficient than BMH in terms of accuracy of the resulting parameter estimates. Compared to the averaged likelihood method, BMH has some significant advantages. As explained previously, it follows from Jensen’s inequality that for any $m \in \mathbb{N}$, the mode of $g_m(\mathcal{D}|\boldsymbol{\theta})$ is identical to $\boldsymbol{\theta}_*$; that is, the *maximum mean log-likelihood estimator* is identical to the true parameter. This makes $I_{m,n,k}(\mathcal{D}_s|\boldsymbol{\theta})$ perform like a regular log-likelihood function when n and k are large. However, this property does not hold for the maximum mean likelihood estimator. We note that the *maximum mean log-likelihood estimator* has been explored in Liang et al. (2013) using a resampling-based stochastic approximation method in the context of large geostatistical data.

In this paper, we have considered only the use of BMH in parameter estimation. Applying BMH to model selection is straightforward. For example, BMH can be combined with the reversible jump MCMC algorithm (Green, 1995) in a similar way to tempering BMH for tackling the problem of model selection:

Algorithm 4

1. Draw k bootstrap samples, $D_1, \dots, D_k$, with/without replacement from the entire training dataset $\mathcal{D}$.
2. Run the reversible jump MCMC algorithm for a few steps, where the energy function is calculated by averaging the energy values calculated from the bootstrap samples $D_1, \dots, D_k$.

In addition to the parameter estimation and model selection problems, BMH can also be applied to optimization problems by running it under the framework of simulated annealing (Kirkpatrick et al., 1983). This leads to the annealing BMH algorithm. A further exploration for the performance of these algorithms is of great interest.

In this paper, BMH is developed for *i.i.d* observations. Extending BMH to dependent data is possible. If the observations in $\mathcal{D}$ are generally dependent, say, they are generated from a dependent process, then we can define

$$\tilde{f}(D_i|\boldsymbol{\theta})=f(X_1^{(i)}, \dots, X_m^{(i)}|\boldsymbol{\theta}),$$

which treats D_i as a sample from the original process. In this case, only the $\binom{n}{m}$ -bootstrapping scheme can be used, as $\tilde{f}(D_i|\boldsymbol{\theta})$ might not be well defined if D_i includes repeated observations. It can be seen from Section 3.1 that the key to establish the validity of BMH is the weak consistency of U -statistics with increasing degrees. Given the weak consistency of the U -statistics, it is straightforward to show that the stationary distribution of BMH is asymptotically equivalent to $\tilde{\pi}_m(\boldsymbol{\theta}|\mathcal{D})$.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank the editor, associate editor and two referees for their constructive comments which have led to significant improvement of this paper. Liang's research was partially supported by the NSF grant DMS-1545022 and DMS-1545738 and the NIH grant R01-GM117597.

References

- Ahn, S., Korattikara, A., Welling, M. Bayesian posterior sampling via stochastic gradient Fisher scoring. Proceedings of the International Conference on Machine Learning; 2012.
- Andrieu C, Roberts GO. The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics*. 2009; 37:697–725.
- Bardenet, R., Doucet, A., Holmes, C. Towards scaling up MCMC: An adaptive subsampling approach. Proceedings of the International Conference on Machine Learning; 2014.
- Bickel PJ, Götze F, van Zwet WR. Resampling fewer than n observations: Gains, losses, and remedies for losses. *Statistica Sinica*. 1997; 7:1–31.
- Blackard JA, Dean DJ. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture*. 1999; 24:131–151.
- Chen CF. On asymptotic normality of limiting density functions with Bayesian implications. *Journal of the Royal Statistical Society, Series B*, bf. 1985; 47:540–546.
- Chen T, Fox EB, Guestrin C. Stochastic gradient Hamiltonian Monte Carlo. 2014 arXiv: 1402.4102v1.
- Cybenko G. Approximations by superpositions of sigmoidal functions. *Mathematics of Control, Signals, and Systems*. 1989; 2:303–314.
- Dawid AP. On the limiting normality of posterior distributions. *Proc Camb Phil Soc*. 1970; 67:625–633.
- Dean J, Ghemawat S. Mapreduce: simplified data processing on large clusters. *Comm ACM*. 2008; 51:107–113.
- Fox CW, Roberts SJ. A tutorial on variational Bayesian inference. *Artif Intell Rev*. 2012; 38:85–95.
- Fraley C, Raftery AE. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*. 2002; 97:611–631.

- Gelman, A., Roberts, RO., Gilks, WR. Efficient Metropolis jumping rules. In: Bernardo, JM. Berger, JO. Dawid, AP., Smith, AFM., editors. Bayesian Statistics. Vol. 5. New York: Oxford University Press; 1996. p. 599-607.
- Geyer, CJ. Markov chain Monte Carlo maximum likelihood. In: Keramigas, EM., editor. Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface. 1991. p. 153-163.
- Green PJ. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*. 1995; 82:711-732.
- Heyde CC, Johnstone IM. On asymptotic posterior normality for stochastic processes. *Journal of the Royal Statistical Society, Series B*. 1979; 41:184-189.
- Hoeffding W. A nonparametric test for independence. *Ann Math Statist*. 1948; 19:546-557.
- Hoeffding, W. Institute of Statistics Mimeo Series No. 302. University of North Carolina; 1961. The strong law of large numbers for U -statistics.
- Hornik K. Approximation capabilities of multilayer feedforward networks. *Neural Networks*. 1991; 4:251-257.
- Hukushima K, Nemoto K. Exchange Monte Carlo method and application to spin glass simulations. *J Phys Soc Jpn*. 1996; 65:1604-1608.
- Kass, RE., Tierney, L., Kadane, JB. The validity of posterior expansions based on Laplace's method. In: Geisser, S. Hodges, JS. Press, SJ., Zellner, A., editors. *Essays in Honor of George Barnard*. Amsterdam: North-Holland; 1990. p. 473-488.
- Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science*. 1983; 220:671-680. [PubMed: 17813860]
- Kleiner A, Talwalkar A, Sarkar P, Jordan MI. A Scalable Bootstrap for Massive Data. 2012 arXiv: 1112.5016v2.
- Korattikara, A., Chen, Y., Welling, M. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. *Proceedings of the International Conference on Machine Learning (ICML)*; 2014; 2014. See also arXiv:1304.5299v4
- Lee, AJ. *U-Statistics: Theory and Practice*. Marcel Dekker, Inc; New York: 1990.
- Liang F. An effective Bayesian neural network classifier with a comparison study to support vector machine. *Neural Computation*. 2003; 15:1959-1989.
- Liang F, Cheng Y, Song Q, Park J, Yang P. A resampling-based stochastic approximation method for analysis of large geostatistical data. *Journal of the American Statistical Association*. 2013; 108:325-339.
- Liang F, Jin IH. A Monte Carlo Metropolis-Hastings algorithm for sampling from distributions with intractable normalizing constants. *Neural Computation*. 2013; 25:2199-2234. [PubMed: 23607562]
- Lin N, Xi R. Aggregated estimating equation estimation. *Statistics and Its Interface*. 2011; 4:73-83.
- Lindley, DV. *Proceedings of the Fourth Berkeley Symposium*. Vol. 1. Berkeley, CA: University of California Press; 1961. The use of prior probability distributions in statistical inference and decisions; p. 453-468.
- Lindley, DV. Approximate Bayesian methods. In: Bernardo, JM. Degroot, MH. Lindley, DV., Smith, AFM., editors. *Bayesian Statistics*. Valencia, Spain: University Press; 1980. p. 223-245.
- Miasojedow B, Moulines E, Vihola M. Adaptive Parallel Tempering Algorithm. 2012 arXiv: 1205.1076.
- Miyata Y. Fully exponential Laplace approximations using asymptotic modes. *Journal of the American Statistical Association*. 2004; 99:1037-1049.
- Politis DN, Romano JP. Large Sample Confidence Regions Based on Subsamples under Minimal Assumptions. *Ann Statist*. 1994; 22:2031-2050.
- Serfling, RJ. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, Inc; 1980.
- Stephens M. Dealing with label switching in mixture models. *Journal of the Royal Statistical Society, Series B*. 2000; 62:795-809.
- Tibshirani R. Regression shrinkage and selection via the LASSO. *J Roy Statist Soc Ser B*. 1996; 58:267-288.

- Tierney L. Markov chains for exploring posterior distributions (with discussion). *Annals of Statistics*. 1994; 22:1701–1762.
- Walker AM. On the asymptotic behavior of posterior distributions. *Journal of the Royal Statistical Society, Series B*. 1969; 31:80–88.
- Welling, M., Teh, YW. Bayesian learning via stochastic gradient Langevin dynamics. *Proceedings of the International Conference on Machine Learning*; 2011.

Figure 1.

A fully connected one hidden layer MLP network with four input units (I_1 , I_2 , I_3 , I_4), one bias unit (B), three hidden units (H_1 , H_2 , H_3), and three output units (O_1 , O_2 , O_3). The arrows show the direction of data feeding.

Figure 2.
The flowchart of the tempering BMH algorithm with 3 processors.

Figure 3.

(a) Scatterplot of the samples initially contained in $\mathcal{D}_1$; (b) scatterplot of the samples contained in $\mathcal{D}_1$ after 5 iterations; (c) scatterplot of the samples contained in $\mathcal{D}_1$ after 10 iterations.

Training and prediction rates of BNN for the forest cover type data, where the average network size, training rate, and prediction rate were calculated by averaging over five independent runs of tempering BMH with the standard deviations given in the parentheses. The CPU time was recorded on a cluster machine which uses Quad-Core AMD Opteron(tm) Processor 2380@2.5Ghz.

Table 1

Bootstrapping	k	Average network Size	Training rate(%)	Prediction rate(%)	CPU(h)
$m n$	20	259.0 (1.18)	72.2 (0.17)	72.4 (0.17)	32.0 (2.9)
	40	259.0 (0.95)	72.2 (0.15)	72.4 (0.12)	33.7 (3.2)
$\begin{pmatrix} n \\ m \end{pmatrix}$	20	260.0 (1.52)	72.3 (0.07)	72.4 (0.07)	28.5 (2.5)
	40	263.5 (0.83)	72.3 (0.15)	72.4 (0.16)	31.9 (3.0)