

Likelihood Inflating Sampling Algorithm

Reihaneh Entezari^{*1}, Radu V. Craiu^{†1}, and Jeffrey S. Rosenthal^{‡1}

¹Department of Statistical Sciences, University of Toronto

June 30, 2017

Abstract

Markov Chain Monte Carlo (MCMC) sampling from a posterior distribution corresponding to a massive data set can be computationally prohibitive since producing one sample requires a number of operations that is linear in the data size. In this paper, we introduce a new communication-free parallel method, the *Likelihood Inflating Sampling Algorithm (LISA)*, that significantly reduces computational costs by randomly splitting the dataset into smaller subsets and running MCMC methods *independently* in parallel on each subset using different processors. Each processor will be used to run an MCMC chain that samples sub-posterior distributions which are defined using an “inflated” likelihood function. We develop a strategy for combining the draws from different sub-posteriors to study the full posterior of the Bayesian Additive Regression Trees (BART) model. The performance of the method is tested using simulated data and a large socio-economic study.

Keywords: Bayesian Additive Regression Trees (BART), Bayesian Inference, Big data, Consensus Monte Carlo, Markov Chain Monte Carlo (MCMC).

^{*}entezari@utstat.utoronto.ca

[†]craiu@ustat.toronto.edu web: <http://www.utstat.toronto.edu/craiu/>

[‡]jeff@math.toronto.edu web: <http://probability.ca/jeff/>

1 Introduction

Markov Chain Monte Carlo (MCMC) methods are essential for sampling highly complex distributions. They are of paramount importance in Bayesian inference as posterior distributions are generally difficult to characterize analytically (e.g., Brooks et al., 2011; Craiu and Rosenthal, 2014). When the posterior distribution is based on a massive sample of size N , posterior sampling can be computationally prohibitive since for some widely-used samplers at least $O(N)$ operations are needed to draw one MCMC sample. Additional issues include memory and storage bottlenecks where datasets are too large to be stored on one computer.

A common solution relies on parallelizing the computation task, i.e. dividing the load among a number of parallel *workers*, where a worker can be a processing unit, a computer, etc. Given the abundant availability of processing units, such strategies can be extremely efficient as long as there is no need for frequent communication between workers. Some have discussed parallel MCMC methods (Wilkinson, 2006; Rosenthal, 2000; Laskey and Myers, 2003) such that each worker runs on the full dataset. However, these methods do not resolve memory overload, and also face difficulties in assessing the number of burn-in iterations for each processor.

A truly parallel approach is to divide the dataset into smaller groups and run parallel MCMC methods on each subset using different workers. Such techniques benefit from not demanding space on each computer to store the full dataset. Generally, one needs to avoid frequent communication between workers, as it is time consuming. In a typical divide and conquer strategy the data is partitioned into non-overlapping sub-sets, called *shards*, and each shard is analyzed by a different worker. For such strategies some essential MCMC-related questions are: 1) how to define the sub-posterior distributions for each shard, and 2) how to combine the MCMC samples obtained from each sub-posterior so that we can recover the same information that would have been obtained by sampling the full

posterior distribution. Existing communication-free parallel methods proposed by Scott et al. (2013), Neiswanger et al. (2013) and Wang and Dunson (2013) have in common the fact that the product of the unnormalized sub-posteriors is equal to the unnormalized full posterior distribution, but differ in the strategies used to combine the samples. Specifically, Neiswanger et al. (2013) approximate each sub-posterior using kernel density estimators, while Wang and Dunson (2013) use the Weierstrass transformation. The popular Consensus Monte Carlo (CMC) method (Scott et al., 2013) relies on a weighted averaging approach to combine sub-posterior samples. The CMC relies on theoretical derivations that guarantee its validity when the full-data posterior and all sub-posteriors are Gaussian or mixtures of Gaussian.

We introduce a new communication-free parallel method, the *Likelihood Inflating Sampling Algorithm (LISA)*, that also relies on independent and parallel processing of the shards by different workers to sample the sub-posterior distributions. The latter are defined differently than in the competing approaches described above. In this paper, we develop techniques to combine the sub-posterior draws obtained for LISA in the case of Bayesian Additive Regression Trees (BART) (Chipman et al., 1998, 2010; Kapelner and Bleich, 2013) and compare the performance of our method with CMC.

Sections 2 and 3 contain a brief review of the CMC algorithm and the detailed description of LISA, respectively. Section 4 illustrates the potential difference brought by LISA over CMC in a simple Bernoulli example, and includes a simple application of LISA to linear regression models. Section 5 contains the justification for a modified and improved version of LISA for BART. Numerical experiments and the analysis of socio-economic data presented in Section 6 examine the computational performance of the algorithms proposed here and compare it with CMC. We end the paper with some ideas for future work. The Appendix contains theoretical derivations and descriptions of the steps used when running BART.

2 Review of Consensus Monte Carlo

In this paper we assume that of interest is to generate samples from $\pi(\theta|\vec{Y}_N)$, the posterior distribution θ given the iid sample $\vec{Y}_N = \{Y_1, \dots, Y_N\}$ of size N . The assumption is that N is large enough to prohibit running a standard MCMC algorithm in which draws from π are obtained on a single computer. We use the notation $\pi(\theta|\vec{Y}_N) \propto f(\vec{Y}_N|\theta)p(\theta)$, where $f(\vec{Y}_N|\theta)$ is the likelihood function corresponding to the observed data \vec{Y}_N and $p(\theta)$ is the prior. Major issues with MCMC posterior sampling for big data can be triggered because a) the data sample is too large to be stored on a single computer, or b) each chain update is too costly, e.g. if π is sampled via a Metropolis-Hastings type of algorithm each update requires N likelihood calculations.

In order to reduce the computational costs, the CMC method of Scott et al. (2013) partitions the sample into K batches (i.e. $\vec{Y}_N = \cup_{j=1}^K Y^{(j)}$) and uses the workers independently and in parallel to sample each sub-posterior. More precisely, the j -th worker ($j = 1, \dots, K$) will generate samples from the j -th sub-posterior distribution defined as:

$$\pi_{j,CMC}(\theta|Y^{(j)}) \propto f(Y^{(j)}|\theta)p(\theta)^{1/K}.$$

Note that the prior for each batch is considered to be $p_j(\theta) = [p(\theta)]^{1/K}$ such that $p(\theta) = \prod_{j=1}^K p_j(\theta)$ and thus the overall full-data unnormalized posterior distribution which we denote as $\pi_{Full}(\theta|\vec{Y}_N)$ is equal to the product of unnormalized sub-posterior distributions, i.e.

$$\pi_{Full}(\theta|\vec{Y}_N) \propto \prod_{j=1}^K \pi_{j,CMC}(\theta|Y^{(j)}).$$

When the full posterior is Gaussian, the weighted averages of the sub-samples from all batches can be used as full-data posterior draws. That is, assuming $\theta_1^{(k)}, \dots, \theta_S^{(k)}$ are S

sub-samples from the k th worker then the s -th approximate full posterior draw will be:

$$\theta_s = \left(\sum_k w_k \right)^{-1} \sum_k w_k \theta_s^{(k)}$$

where the weights $w_k = \Sigma_k^{-1}$ are optimal for Gaussian models with $\Sigma_k = \text{Var}(\theta|y^{(k)})$.

In the next section we introduce an alternative method to define the sub-posteriors in each batch.

3 Likelihood Inflating Sampling Algorithm (LISA)

LISA is an alternative to CMC that also benefits from the random partition of the dataset followed by independently processing each batch on a different worker. Assuming that the data have been divided into K batches of approximately equal size n , we define the sub-posterior distributions for each machine by adjusting the likelihood function without making changes to the prior. Thus the j -th sub-posterior distribution will be:

$$\pi_{j,LISA}(\theta|Y^{(j)}) \propto [f(Y^{(j)}|\theta)]^K p(\theta).$$

Since the data are assumed to be iid, inflating the likelihood function K -times is intuitive because the sub-posterior from each batch of data will be a closer representation of the whole data posterior. We expect that sub-posteriors sampled by each worker will be closer to the full posterior thus improving the computational efficiency.

We indeed prove in a theorem below that under mild conditions, LISA's sub-posterior distributions are asymptotically closer to the full posterior than those produced by the CMC-type approach.

The Taylor's series expansion for a log-posterior density $\log \pi(\theta|\vec{Y}_N)$ around its poste-

rior mode $\hat{\theta}_N$ yields the approximation

$$\log \pi(\theta|\vec{Y}_N) \approx \log \pi(\hat{\theta}_N|\vec{Y}_N) - \frac{1}{2}(\theta - \hat{\theta}_N)^T \hat{I}_N(\theta - \hat{\theta}_N)$$

where $\hat{I}_N = -\frac{\partial^2 \log(\pi(\theta|\vec{Y}_N))}{\partial \theta \partial \theta^T} \Big|_{\theta=\hat{\theta}_N}$. Exponentiating both sides will result in

$$\pi(\theta|\vec{Y}_N) \approx \pi(\hat{\theta}_N|\vec{Y}_N) \exp \left[-\frac{1}{2}(\theta - \hat{\theta}_N)^T \hat{I}_N(\theta - \hat{\theta}_N) \right]$$

which shows asymptotic normality, i.e. $\hat{I}_N^{1/2}(\Theta - \hat{\theta}_N) \xrightarrow{D} N(0, I)$ as $N \rightarrow \infty$ where $\Theta \sim \pi(\cdot|\vec{Y}_N)$. Let $\hat{\theta}_{n,L}^{(j)}$ and $\hat{\theta}_{n,C}^{(j)}$ denote the j -th sub-posterior modes in LISA and CMC, respectively. Similarly, $\hat{I}_{n,L}^{(j)}$ and $\hat{I}_{n,C}^{(j)}$ denote the negative second derivative of the j -th log sub-posterior for LISA and CMC, respectively, when calculated at the mode. Then consider the assumptions,

A1: There exist θ_L, θ_C such that if we define $\epsilon_{n,L}^{(j)} = |\hat{\theta}_{n,L}^{(j)} - \theta_L|$ and $\epsilon_{n,C}^{(j)} = |\hat{\theta}_{n,C}^{(j)} - \theta_C|$, then $\max_{1 \leq j \leq K} \epsilon_{n,L}^{(j)} \rightarrow 0$ and $\max_{1 \leq j \leq K} \epsilon_{n,C}^{(j)} \rightarrow 0$ w.p. 1 as $n \rightarrow \infty$.

A2: $|\hat{I}_{n,L}^{(i)} - \hat{I}_{n,L}^{(j)}| \rightarrow 0$ and $|\hat{I}_{n,C}^{(i)} - \hat{I}_{n,C}^{(j)}| \rightarrow 0$ w.p. 1 $\forall i \neq j$ as $n \rightarrow \infty$.

A3: π_{Full} , $\pi_{j,LISA}$, and $\pi_{j,CMC}$ are unimodal distributions that have continuous derivatives of order 2.

Theorem 1. Assume that assumptions **A1** through **A3** hold and if $\Theta_{Full} \sim \pi_{Full}(\cdot|\vec{Y}_N)$ we also assume $\hat{I}_N^{1/2}(\Theta_{Full} - \hat{\theta}_N) \xrightarrow{D} N(0, I)$ as $N \rightarrow \infty$. If $\Theta_{j,LISA} \sim \pi_{j,LISA}(\cdot|Y^{(j)})$ and $\Theta_{j,CMC} \sim \pi_{j,CMC}(\cdot|Y^{(j)})$ then as $N \rightarrow \infty$

$$\hat{I}_N^{1/2}(\Theta_{j,LISA} - \hat{\theta}_N) \xrightarrow{D} N(0, I) \quad \text{and} \quad \hat{I}_N^{1/2}(\Theta_{j,CMC} - \hat{\theta}_N) \xrightarrow{D} N(0, KI) \quad \forall j \in \{1, \dots, K\}.$$

Proof. See Appendix. □

Theorem 1 shows the difference between sub-posterior distributions for CMC and LISA, with LISA’s sub-posterior distributions being asymptotically similar to the full posterior distribution. This suggests that draws from LISA sub-posteriors can be combined using uniform weights.

Remarks:

1. When data are iid we expect the shards to become more and more similar as N (and thus $n = N/K$) increases and assumption **A1** is expected to hold for general models.
2. Assumption **A2** in Theorem 1 holds due to the structural form of sub-posteriors in LISA and CMC.
3. The validity of using uniform weights with LISA’s sub-posterior draws is justified asymptotically, but we will see that this approximation can be exact in some examples, e.g. for a Bernoulli model with balanced batch samples, while in others modified weights can improve the performance of the sampler. In this respect LISA is similar to other embarrassing parallel strategies where one must carefully consider the model of interest in order to find the best way to combine the sub-posterior samples.

In the next section we will illustrate LISA in some simple examples and compare its performance to the full-data posterior sampling as well as CMC.

4 Motivating Examples

In this section we examine some simple examples where theoretical derivations can be carried out in detail. We emphasize the difference between LISA and CMC.

4.1 Bernoulli Random Variables

Consider y_1, \dots, y_N to be N i.i.d. Bernoulli random variables with parameter θ . Hence, we consider a prior $p(\theta) = \text{Beta}(a, b)$. Assuming that we know little about the size of θ we set $a = b = 1$ which corresponds to a $U(0, 1)$ prior. The resulting full-data posterior $\pi_{Full}(\theta | \vec{Y}_N)$ is $\text{Beta}(S + a, N - S + b)$ where $S = \sum_{i=1}^N y_i$ is the total number of ones. Suppose we divide the data into K batches with S_j number of ones in batch j , such that $S_j = \frac{S}{K} \forall j \in \{1, \dots, K\}$, i.e. the number of 1's are divided equally between batches. Then the j^{th} sub-posterior based on batch-data of size $n = \frac{N}{K}$ for each method will be:

- **CMC:**

$$\begin{aligned} \pi_{j,CMC}(\theta | Y^{(j)}) &= \text{Beta} \left(S_j + \frac{a-1}{K} + 1, n - S_j + \frac{b-1}{K} + 1 \right) \\ &= \text{Beta} \left(\frac{S}{K} + \frac{a-1}{K} + 1, \frac{N-S}{K} + \frac{b-1}{K} + 1 \right) \end{aligned}$$

- **LISA:**

$$\begin{aligned} \pi_{j,LISA}(\theta | Y^{(j)}) &= \text{Beta}(S_j K + a, (n - S_j)K + b) \\ &= \text{Beta}(S + a, N - S + b) \end{aligned}$$

which implies

$$\pi_{j,LISA}(\theta | Y^{(j)}) = \pi_{Full}(\theta | \vec{Y}_N) \quad \forall j \in \{1, \dots, K\}.$$

In this simple case any one of LISA's sub-posterior distributions is equal to the full posterior distribution if the batches are balanced, i.e. the number of 1's are equally split across all batches. Thus, LISA's sub-samples from any batch will represent correctly the full posterior. On the other hand, the draws from the CMC sub-posterior distributions

will need to be recombined to obtain a representative sample from the true full posterior $\pi_{Full}(\theta|\vec{Y}_N)$.

However, when the number of ones is unequally distributed among the batches it is not easy to pick the winner between CMC and LISA as both require a careful weighting of each batch sub-posterior samples.

In the remaining part of this paper, we will mainly focus on the performance of LISA when it is applied to the Bayesian Additive Regression Trees (BART) model. Interestingly, we discover that using a minor modification inspired by running LISA on the simpler Bayesian Linear Regression model we can approximate the full posterior. The idea behind the modification is described in the next section.

4.2 Bayesian Linear Regression

Consider a standard linear regression model

$$Y = X\beta + \epsilon \tag{1}$$

where $\beta \in \mathbf{R}^p$, $X \in \mathbf{R}^{N \times p}$ and $Y, \epsilon \in \mathbf{R}^N$ with $\epsilon \sim \mathcal{N}_N(0, \sigma^2 \mathbf{I}_N)$. To simplify the presentation we consider the improper prior

$$p(\beta, \sigma^2) \propto \sigma^{-2}. \tag{2}$$

Straightforward calculations show that the conditional posterior distributions for the full data are

$$\pi_{Full}(\sigma^2|Y, X) = \text{Inv-Gamma}\left(\frac{N-p}{2}, \frac{s^2(N-p)}{2}\right) \tag{3}$$

$$\pi_{Full}(\beta|\sigma^2, Y, X) = N(\hat{\beta}, \sigma^2(X^T X)^{-1}) \tag{4}$$

where $\hat{\beta} = (X^T X)^{-1} X^T Y$ and $s^2 = \frac{(Y - X\hat{\beta})^T (Y - X\hat{\beta})}{N-p}$.

An MCMC sampler designed to sample from $\pi_{Full}(\beta, \sigma^2 | Y, X)$ will iteratively sample σ^2 using (3) and then β via (4). If we denote β_{Full} the r.v. with density $\pi_{Full}(\beta | Y, X)$ then, using the iterative formulas for conditional mean and variance we obtain

$$E[\beta_{Full} | Y, X] = (X^T X)^{-1} X^T Y$$

and

$$\text{Var}(\beta_{Full} | Y, X) = (X^T X)^{-1} \frac{(N-p)/2}{(N-p)/2-1} s^2 = (X^T X)^{-1} s^2 + O(N^{-1}). \quad (5)$$

We examine below the statistical properties of the samples produced by LISA. If the data are divided into K equal batches of size $n = N/K$, let us denote $Y^{(j)}$ and $X^{(j)}$ the response vector and model matrix from the j th batch, respectively.

With the prior given in (2), the sub-posteriors produced by LISA have the following conditional densities

$$\pi_j(\sigma^2 | Y^{(j)}, X^{(j)}) = \text{Inv-Gamma} \left(\frac{N-p}{2}, \frac{K s_j^2 (n-p)}{2} \right) \quad (6)$$

$$\pi_j(\beta | \sigma^2, Y^{(j)}, X^{(j)}) = N(\hat{\beta}_j, \frac{\sigma^2}{K} (X^{(j)T} X^{(j)})^{-1}), \quad (7)$$

where $\hat{\beta}_j = (X^{(j)T} X^{(j)})^{-1} X^{(j)T} Y^{(j)}$ and $s_j^2 = \frac{(Y^{(j)} - X^{(j)}\hat{\beta}_j)^T (Y^{(j)} - X^{(j)}\hat{\beta}_j)}{n-p}$ for all $1 \leq j \leq K$.

A simple Gibbs sampler designed to sample from $\pi_j(\beta, \sigma^2 | Y^{(j)}, X^{(j)})$ will iteratively sample σ^2 from (6) and then β from (7).

It can be shown using the iterative formulas for conditional means and variances that

$$E[\beta | Y^{(j)}, X^{(j)}] = \hat{\beta}_j$$

and

$$\text{Var}(\beta|Y^{(j)}, X^{(j)}) = (X^{(j)T} X^{(j)})^{-1} \frac{s_j^2(n-p)/2}{(N-p)/2-1} = (X^{(j)T} X^{(j)})^{-1} \frac{s_j^2(n-p)}{(N-p)} + O(N^{-1}).$$

In order to combine the sub-posterior samples we propose using the weighted average

$$\beta_{LISA} = \left(\sum_{j=1}^K W_j \right)^{-1} \sum_{j=1}^K W_j \beta_j, \quad (8)$$

where $\beta_j \sim \pi_j(\beta|Y^{(j)}, X^{(j)})$ and $W_j = \frac{X^{(j)T} X^{(j)}}{\sigma^2}$. Since $\sum_{j=1}^K X^{(j)T} X^{(j)} = X^T X$ we get

$$E[\beta_{LISA}|Y, X] = \hat{\beta} = (X^T X)^{-1} X^T Y \quad (9)$$

and

$$\text{Var}(\beta_{LISA}|Y, X) = (X^T X)^{-1} \frac{n-p}{N-p} \left[\sum_{j=1}^K s_j^2 (X^{(j)T} X^{(j)}) \right] (X^T X)^{-1} \approx (X^T X)^{-1} \frac{n-p}{N-p} s^2, \quad (10)$$

where the last approximation in (10) is based on the assumption that $s_j^2 \approx s^2$ as both are unbiased estimators for σ^2 based on n and, respectively, N observations. It is apparent that the variance computed in (10) is roughly K times smaller than the target given in (5).

In order to avoid underestimating the variance of the posterior distribution we propose a modified LISA sampling algorithm which consists of the following steps:

$$\begin{aligned} \sigma^2 &\sim \text{Inv-Gamma} \left(\frac{N-p}{2}, \frac{K s_j^2 (n-p)}{2} \right) \\ \tilde{\sigma} &= \sqrt{K} \sigma \\ \tilde{\beta} &\sim N(\hat{\beta}_j, \frac{\tilde{\sigma}^2}{K} (X^{(j)T} X^{(j)})^{-1}) = N(\hat{\beta}_j, \sigma^2 (X^{(j)T} X^{(j)})^{-1}). \end{aligned}$$

The intermediate step simply adjusts the variance samples so that

$$\text{Var}(\tilde{\beta}|Y^{(j)}, X^{(j)}) = (X^{(j)T} X^{(j)})^{-1} \frac{s_j^2 K(n-p)/2}{(N-p)/2 - 1} = (X^{(j)T} X^{(j)})^{-1} \frac{s_j^2 K(n-p)}{(N-p)} + O(N^{-1}).$$

In turn, if we define

$$\beta_{\text{modLISA}} = \left(\sum_{j=1}^K W_j \right)^{-1} \sum_{j=1}^K W_j \tilde{\beta}_j, \quad (11)$$

then $E[\beta_{\text{modLISA}}|Y, X] = (X^T X)^{-1} X^T Y$ and

$$\text{Var}(\beta_{\text{modLISA}}|Y, X) = (X^T X)^{-1} \frac{K(n-p)}{N-p} \left[\sum_{j=1}^K s_j^2 (X^{(j)T} X^{(j)}) \right] (X^T X)^{-1} \approx (X^T X)^{-1} \frac{K(n-p)}{N-p} s^2. \quad (12)$$

Note that when the regression has only an intercept, i.e. X consists of a column of 1's, the weights $W_j \propto (\sigma^2)^{-1}$.

While both (11) and (8) produce samples that have the correct mean, from equations (5), (10) and (12) we can see that the weighted average of the modified LISA samples have the variance closer to the desired target.

In the next section, we will examine LISA's performance on a more complex model, the Bayesian Additive Regression Trees (BART). The discussion above will guide our construction of a modified version of LISA for BART.

5 Bayesian Additive Regression Trees (BART)

Consider the nonparametric regression model:

$$y_i = f(x_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma^2) \quad i.i.d.$$

where $x_i = (x_{i1}, \dots, x_{ip})$ is a p -dimensional vector of inputs and f is approximated by a sum of m regression trees:

$$f(x) \approx \sum_{j=1}^m g(x; T_j, M_j)$$

where T_j denotes a binary tree consisting of a set of interior node decision rules and a set of terminal nodes. $M_j = \{\mu_{1j}, \dots, \mu_{bj}\}$ is the set of parameter values associated with the b terminal nodes of T_j . In addition, $g(x; T_j, M_j)$ is the function that maps each x to a $\mu_{ij} \in M_j$. Thus the regression model is approximated by a sum-of-trees model

$$y_i = \sum_{j=1}^m g(x_i; T_j, M_j) + \epsilon_i, \quad \epsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$$

Let $\theta := ((T_1, M_1), \dots, (T_m, M_m), \sigma^2)$ denote the vector of model parameters. Below, we briefly describe the prior specifications stated in Chipman et al. (2010) and Chipman et al. (1998).

Prior Specifications:

- Prior Independence and Symmetry:

$$p((T_1, M_1), \dots, (T_m, M_m), \sigma) = \left[\prod_j p(M_j | T_j) p(T_j) \right] p(\sigma)$$

where $p(M_j | T_j) = \prod_i p(\mu_{ij} | T_j)$.

- Recommended number of trees: $m=200$ (Chipman et al., 2010) and $m=50$ (Kapelner and Bleich, 2013)
- Tree prior $p(T_j)$, is characterised by three aspects:

1. The probability that a node at depth $d = 0, 1, \dots$ is non-terminal, which is

assumed to have the form $\alpha(1+d)^{-\beta}$, where $\alpha \in (0,1)$ and $\beta \geq 0$. (recommended values are $\alpha = 0.95$ and $\beta = 2$)

2. The distribution on the splitting variable assignments at each interior node which is recommended to have a uniform distribution.
 3. The distribution on the splitting rule assignment in each interior node, conditional on the splitting variable which is also recommended to have a uniform distribution.
- The conditional prior for μ_{ij} is $\mathcal{N}(\mu_\mu, \sigma_\mu^2)$ such that:

$$\begin{cases} m\mu_\mu - k\sqrt{m}\sigma_\mu = y_{min} \\ m\mu_\mu + k\sqrt{m}\sigma_\mu = y_{max} \end{cases}$$

with $k = 2$ recommended.

- The prior for σ^2 is Inv-Gamma($\frac{\nu}{2}, \frac{\nu\lambda}{2}$) where $\nu = 3$ is recommended and λ is chosen such that $p(\sigma < \hat{\sigma}) = q$ with recommended $q = 0.9$ and sample variance $\hat{\sigma}$.

Hence the posterior distribution will have the form:

$$\pi(\theta) = \pi(\theta|Y, X) \propto \underbrace{\left\{ (\sigma^2)^{-\frac{n}{2}} e^{-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - \sum_{j=1}^m g(x_i; M_j, T_j))^2} \right\}}_{\text{Likelihood}} \times \underbrace{\left\{ (\sigma^2)^{-\frac{\nu}{2}-1} e^{-\frac{\nu\lambda}{2\sigma^2}} \left[\prod_{j=1}^m \sigma_\mu^{-b_j} (2\pi)^{-\frac{b_j}{2}} e^{-\frac{1}{2\sigma_\mu^2} \sum_{k=1}^{b_j} (\mu_{kj} - \mu_\mu)^2} p(T_j) \right] \right\}}_{\text{Prior}}. \quad (13)$$

Gibbs Sampling is used to sample from this posterior distribution. The algorithm iterates between the following steps:

- $\sigma^2 \mid (T_1, M_1), \dots, (T_m, M_m), Y, X \propto \text{Inv-Gamma}(\rho, \gamma)$
 where $\rho = \frac{\nu+n}{2}$ and $\gamma = \frac{1}{2} [\sum_{i=1}^n (y_i - \sum_{j=1}^m g(x_i; M_j, T_j))^2 + \lambda\nu]$.

- $(T_j, M_j) \mid T_{(j)}, M_{(j)}, \sigma, Y, X$ which is the same as drawing from the conditional $(T_j, M_j) \mid R_j, \sigma$ where $T_{(j)}$ denotes all trees except the j -th tree, and residual R_j is defined as:

$$R_j = g(x; M_j, T_j) + \epsilon = y - \sum_{k \neq j} g(x; M_k, T_k).$$

The sampling of (T_j, M_j) is performed in two steps:

1. $T_j \mid R_j, \sigma$ and
2. $M_j \mid T_j, R_j, \sigma$.

Step 2 involves sampling from each component of M_j using

$$\mu_{ij} \mid T_j, R_j, \sigma \sim \mathcal{N} \left(\frac{\frac{\sigma^2}{\sigma_\mu^2} \mu_\mu + n_i \bar{R}_{j(i)}}{\frac{\sigma^2}{\sigma_\mu^2} + n_i}, \frac{\sigma^2}{\frac{\sigma^2}{\sigma_\mu^2} + n_i} \right)$$

where $\bar{R}_{j(i)}$ denotes the average residual (computed without tree j) at terminal node i with total number of observations n_i . The conditional density of T_j in step 1 can be expressed as:

$$p(T_j \mid R_j, \sigma) \propto p(T_j) \int p(R_j \mid M_j, T_j, \sigma) p(M_j \mid T_j, \sigma) dM_j. \quad (14)$$

The Metropolis-Hastings (MH) algorithm is then applied to draw T_j from (14) with four different proposal moves on trees:

- **GROW:** growing a terminal node (with probability 0.25);
- **PRUNE:** pruning a pair of terminal nodes (with probability 0.25);
- **CHANGE:** changing a non-terminal rule (with probability 0.4) (Kapelner and Bleich, 2013, change rules only for parent nodes with terminal children);

- **SWAP:** swapping a rule between parent and child (with probability 0.1) (This proposal move was removed by Kapelner and Bleich, 2013).

Detailed derivations involving the Metropolis-Hastings acceptance ratios are described in the Appendix.

Two existing packages in R, "BayesTree" and "bartMachine", can be used to run BART on any dataset, but as the sample size increases, these packages tend to run slower. In these situations we expect methods such as LISA or CMC to become useful, and for a fair illustration of the advantages gained we have used our own R implementation of BART and applied the same structure to implement LISA and CMC algorithm for BART. The Metropolis-Hastings acceptance ratios for LISA and CMC are also reported in the Appendix.

As discussed by Scott et al. (2013), the approximation to the posterior produced by the CMC algorithm can be poor. Thus, for comparison reasons, we applied both LISA and CMC to BART using a simulated dataset (described further) with $K = 30$ batches. Given Theorem 1, since LISA's sub-posterior distributions are asymptotically equivalent to the full posterior distribution, we examined its performance by uniformly taking sub-samples from all its batches as an approximation to full posterior samples. We will see further that LISA with uniform weights produces higher prediction accuracy compared to CMC. However, they both perform poorly in approximating the posterior samples as they generate larger trees and under-estimate σ^2 , which results in over-dispersed posterior distributions.

The following sub-section discusses a modified version of LISA for BART which will have significant improvement in performance.

5.1 Modified LISA for BART

The under estimation of σ^2 when applying LISA to BART is similar to the problem encountered when using LISA for the linear regression model discussed in Section 4.2. This is not a coincidence since BART is also a linear regression model, albeit one where the set of independent variables is determined through a highly sophisticated process. We will show below that when applying a similar variance adjustment to the one discussed in Section 4.2, the Modified LISA (modLISA) for BART will exhibit superior computational and statistical efficiency compared to either LISA or CMC.

Just like in the regression model we “correct” the sampling algorithm by adjusting the residual variance. We start with the conditional distribution of tree j from expression (14) which takes the form

$$p(T_j | R_j, \sigma) \propto p(T_j) \int p(R_j | M_j, T_j, \sigma) p(M_j | T_j, \sigma) dM_j.$$

Note that only the conditional distribution of the residuals, $R_j | M_j, T_j, \sigma$ is affected by the modifications brought by LISA. The Metropolis-Hastings acceptance ratios for tree proposals contain three parts: the transition ratio, the likelihood ratio and the tree structure ratio. The modifications brought by LISA will influence only the likelihood ratio which is constructed from the conditional distributions of residuals. Consider the likelihood ratio for GROW proposal in LISA (full details are presented in the Appendix)

$$\frac{P(R | T_*, \sigma^2)}{P(R | T, \sigma^2)} = \sqrt{\frac{\sigma^2(\sigma^2 + Kn_l\sigma_\mu^2)}{(\sigma^2 + Kn_{l_L}\sigma_\mu^2)(\sigma^2 + Kn_{l_R}\sigma_\mu^2)}} \times \exp \left\{ \frac{K^2\sigma_\mu^2}{2\sigma^2} \left[\frac{(\sum_{i=1}^{n_{l_L}} R_{l_L,i})^2}{\sigma^2 + Kn_{l_L}\sigma_\mu^2} + \frac{(\sum_{i=1}^{n_{l_R}} R_{l_R,i})^2}{\sigma^2 + Kn_{l_R}\sigma_\mu^2} - \frac{(\sum_{i=1}^{n_l} R_{l,i})^2}{\sigma^2 + Kn_l\sigma_\mu^2} \right] \right\} \quad (15)$$

where n_l is the total number of observations from batch-data that end up in terminal node

l . The newly grown tree, T_* , splits terminal node l into two terminal nodes (children) l_L and l_R , which will also divide n_l to n_{l_L} and n_{l_R} which are the corresponding number of observations in each new terminal node. By factoring out K in (15), we can rewrite it as

$$\frac{P(R | T_*, \sigma^2)}{P(R | T, \sigma^2)} = \sqrt{\frac{\frac{\sigma^2}{K}(\frac{\sigma^2}{K} + n_l \sigma_\mu^2)}{(\frac{\sigma^2}{K} + n_{l_L} \sigma_\mu^2)(\frac{\sigma^2}{K} + n_{l_R} \sigma_\mu^2)}} \times \exp \left\{ \frac{\sigma_\mu^2}{2\frac{\sigma^2}{K}} \left[\frac{(\sum_{i=1}^{n_{l_L}} R_{l_L,i})^2}{\frac{\sigma^2}{K} + n_{l_L} \sigma_\mu^2} + \frac{(\sum_{i=1}^{n_{l_R}} R_{l_R,i})^2}{\frac{\sigma^2}{K} + n_{l_R} \sigma_\mu^2} - \frac{(\sum_{i=1}^{n_l} R_{l,i})^2}{\frac{\sigma^2}{K} + n_l \sigma_\mu^2} \right] \right\}. \quad (16)$$

Expression (16) shows a similar residual variance that is K times smaller in each batch, and hence following the discussion in Section 4.2, to achieve similar variance, we need to modify LISA for BART by adding the intermediate step $\tilde{\sigma}^2 = K\sigma^2$ when updating *trees* in each batch, and then taking a weighted average combination of sub-samples (similar to Bayesian linear regression). As in Section 4.2, we don't apply any changes when updating σ^2 . All our numerical experiments show that modLISA also generates accurate predictions in BART, since the modification corrects the bias in the posterior draws of σ^2 and properly calibrates the size of the trees.

The BART algorithm will split the covariate space into disjoint subsets and on each subset a regression with only an intercept is fitted. Therefore, as suggested by the discussion in 4.2 the weight assigned to each batch will be proportional to the estimate of σ^2 in that batch. In the following sections we examine the improvement brought by modLISA when compared to LISA and CMC.

6 Numerical Experiments

6.1 The Friedman’s function

We have simulated data of size $N = 20,000$ from Friedman’s test function (Friedman, 1991)

$$f(x) = 10 \sin(\pi x_1 x_2) + 20(x_3 - 0.5)^2 + 10x_4 + 5x_5,$$

where the covariates $x = (x_1, \dots, x_{10})$ are simulated independently from a $U(0, 1)$ and $y \sim \mathcal{N}(f(x), \sigma^2)$ with $\sigma^2 = 9$. Note that five of the ten covariates are unrelated to the response variable. We have also generated test data containing 5000 cases. We apply BART to this simulated dataset using the default hyperparameters stated in Section 5 with $m = 50$ to generate posterior draws of (T, M, σ^2) that, in turn, yield posterior draws for $f(x)$ using the approximation $\hat{f}(x) \approx \sum_{j=1}^m g(x; \hat{T}_j, \hat{M}_j)$ for each $x = (x_1, \dots, x_{10})$. Since in this case the true f is known, one can compute the root mean squared error (RMSE) using average posterior draws of $\hat{f}(x)$ for each x (i.e. $\overline{\hat{f}(x)}$), as an estimate to measure its performance, i.e. $\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (f(x_i) - \overline{\hat{f}(x_i)})^2}$. It is known that SingleMachine BART may mix poorly when it is run on an extremely large dataset with small residual variance. However since the data simulated is of reasonable size and σ is not very small the SingleMachine BART is expected to be a good benchmark for comparison (see discussion in Pratola et al., 2016).

6.1.1 Comparison of modLISA with Competing Methods

We have implemented modLISA, LISA, and CMC for BART with $K = 30$ batches on the simulated data for 5000 iterations with a total of 1000 posterior draws. Table 1 shows results from all methods including the SingleMachine which runs BART on the full dataset using only one machine. Results are averaged over three different realizations of train and test data, and include the Train and Test RMSE for each method, along with tree sizes,

σ^2 estimates and their 95% Credible Intervals (CI). The summaries presented in Table 1 show that although LISA has better prediction performance than CMC, it does a terrible job at estimating σ^2 , its estimate being orders of magnitude smaller than the one produced by CMC. CMC and LISA both generate larger trees compared to SingleMachine, with CMC generating trees that are ten times larger than LISA’s. One can see that modLISA with weighted averages dominates both CMC and LISA across all performance indicators since it yields the smallest RMSE, the smallest tree size, and less biased σ^2 estimates. Generally, modLISA generates results that are by far the closest to the ones produced by SingleMachine.

Method	TrainRMSE	TestRMSE	Tree Nodes	Avg $\hat{\sigma}^2$	95% CI for σ^2
<i>CMC</i>	2.73	2.94	602	1.91	[1.45 , 2.88]
<i>LISA (unif wgh)</i>	1.18	1.19	55	0.001	[0.0009 , 0.0011]
<i>modLISA (wgh avg)</i>	0.57	0.59	7	7.97	[7.87 , 8.08]
<i>SingleMachine</i>	0.55	0.56	7	9.04	[8.85 , 9.21]

Table 1: Comparing Train & Test RMSE, tree sizes, and average post burn-in $\hat{\sigma}^2$ with 95% CI in each method for $K = 30$ to SingleMachine BART (all results are averaged over three different realizations of data).

Method	GROW	PRUNE	CHANGE
<i>CMC</i>	21%	0.03%	34%
<i>LISA</i>	1.8%	0.5%	1.6%
<i>modLISA</i>	20%	26%	19%
<i>SingleMachine</i>	9%	10%	6%

Table 2: Average acceptance rates of tree proposal moves.

The size of trees produced by each method is in sync with the average acceptance rates of each tree proposal move shown in Table 2. It is noticeable the difference between CMC and LISA ’s average acceptance rates between growing a tree and pruning one. On the other hand, modLISA has overall larger acceptance rates with the smallest relative absolute difference between growing and pruning probabilities compared to LISA and

CMC (6/20 = 23.1% for modLISA, 98.6% for CMC, and 72.2% for LISA) and is closest to SingleMachine (10%). Overall, modLISA induced a significant reduction in tree sizes by preserving a balance between growing and pruning trees which also improves exploring the posterior distribution.

Method	TrainPredCov	TestPredCov	TrainCredCov	TestCredCov
<i>CMC</i>	45.71 %	47.83 %	81.95 %	99.99 %
<i>LISA (unif wgh)</i>	1.54 %	1.54 %	100 %	100 %
<i>modLISA (wgh avg)</i>	92.93 %	92.91 %	60.88 %	58.45 %
<i>SingleMachine</i>	94.67 %	94.65 %	71.58 %	71.54 %

Table 3: Average coverage for 95% credible intervals constructed for training (TrainCredCov) and test (TestCredCov) data and 95% prediction intervals constructed for training (TrainPredCov) and test (TestPredCov) data. The prediction interval coverage is estimated based on 1000 iid samples, $N = 20,000$ and $K = 30$. All results are averaged over three different realizations of data.

For a more clear comparison of the methods, Table 3 shows the average coverage of 95% credible intervals (CI) for predictors $f(x)$ and 95% prediction intervals (PI) for future responses y . The calculations are made for the values of y and $f(x)$ in the training and test data sets.

The coverage for CI is given by the averaging for all training or test data of

$$\frac{\#\{f(x_i) \in \hat{I}_{f(x_i)} : 1 \leq i \leq N\}}{N}$$

where $\hat{I}_{f(x_i)}$ is the CI for $f(x_i)$ estimated based on the MCMC draws from π .

The coverage of the PI corresponding to a pair $(y_i, f(x_i))$ is given by the proportion of 1000 iid samples generated from the true generative model $N(f(x_i), \sigma^2)$ that fall between its limits, i.e. the average over training or test data of

$$\frac{\#\{\tilde{y}_j \in \hat{J}_{y_i} : \tilde{y}_j \stackrel{iid}{\sim} N(f(x_i), \sigma^2) 1 \leq j \leq 1000\}}{1000},$$

where \hat{J}_{y_i} is the PI for y_i . The PI coverage in modLISA and SingleMachine are very close to nominal and vastly outperform the PI’s produced using LISA or CMC.

One can see that coverages of the CI built via CMC and LISA are high, which is not surprising since both algorithms produce over-dispersed approximations to the conditional distributions of $f(x)$. Our observation is that the CI for LISA and CMC are too wide to be practically useful. Also, modLISA and SingleMachine have much lower CI coverage than nominal which, as pointed out by one of the referees, is also expected due to the systematic bias induced by the discrepancy between the functional forms of the true predictor (continuous) and of the one fitted by BART (piecewise constant). Thus, the CI for $f(x)$ will exhibit poor coverage as they are centered around a biased estimate of $f(x)$.


In order to verify that this is indeed the case we have generated a dataset of size 20,000 from the piecewise constant function:

$$f(x) = \mathbf{1}_{[0,0.2)}(x_1) + 2 \cdot \mathbf{1}_{[0.2,0.4)}(x_1) + 3 \cdot \mathbf{1}_{[0.4,0.6)}(x_1) + 4 \cdot \mathbf{1}_{[0.6,0.8)}(x_1) + 5 \cdot \mathbf{1}_{[0.8,1)}(x_1)$$

where $\mathbf{1}_{[a,b)}(x) = 1$ if $x \in [a, b)$ and 0 otherwise, $x = (x_1, \dots, x_{10}) \in (0, 1)^{10}$ is a ten-dimensional input vector, with $x_i \sim \text{Uniform}(0, 1)$, and $y \sim \mathcal{N}(f(x), 9)$. Additional 5000 data have also been simulated as test cases. Table 4 summarizes the analysis with $K = 30$ and confirms a sharp decrease in RMSEs even though the noise has the same variance $\sigma^2 = 9$. We note that the coverages of CI build under modLISA and SingleMachine are much higher.


Method	TestRMSE	TestCredCov
<i>CMC</i>	1.35	100 %
<i>LISA (unif wgh)</i>	0.94	100 %
<i>modLISA (wgh avg)</i>	0.24	90.16 %
<i>SingleMachine</i>	0.15	98.76 %

Table 4: Comparing test data RMSE and coverage of 95% credible intervals for piecewise $f(x)$ with $N = 20,000$ and $K = 30$.



(a) Test $x^* = 1000$, $f(x^*) = 13.8$

(b) Test $x^* = 2000$, $f(x^*) = 14.4$



(c) Training $x = 999$, $f(x) = 19.8$

(d) Training $x = 2001$, $f(x) = 11.2$

Figure 1: Empirical distribution functions of $\hat{f}(x)$ obtained from MCMC samples produced by modLISA (red line), LISA (green line), CMC (blue line), and SingleMachine BART (black line) for two different pairs of training and test data. In this example $K = 30$.

6.1.2 Comparison with SingleMachine BART

In order to investigate the closeness of posterior samples in each method to the SingleMachine BART, we have plotted in Figure 1 the empirical distribution functions of $\hat{f}(x)$ generated from each algorithm for two pairs of observations in the training and test dataset. One can see that the empirical distribution functions in LISA and CMC don't match the ones from SingleMachine, and look over-dispersed. However, the empirical

distribution functions in modLISA weighted average look much closer to SingleMachine with a slight shift in location.




Figure 2: Blue lines: Fitted polynomial trends (for both train and test data) of average squared difference between empirical distribution functions of SingleMachine and the following: (a) CMC for training (solid line) and test (dot dashed line) data, (b) LISA with uniform weights for training (solid line) and test (dot dashed line) data and (c) modLISA with weighted average for training (solid line) and test (dot dashed line) data. The difference is plotted against the mean prediction $\hat{f}(x)$ produced by SingleMachine. Grey areas represent the 95% credible intervals constructed from 100 independent replicates.

In order to assess the performance of the sampling procedures considered, we use the Cramér-von Mises distance to assess the difference between empirical distribution functions. This distance is defined to be $\omega^2 = \int_{-\infty}^{\infty} (F_n(x) - F(x))^2 dF(x)$ where in our case we assume $F(x) = F_{BART}(x)$ to be the empirical distribution function generated from posterior samples in SingleMachine BART and $F_n(x)$ is similarly computed for the alternative method that is considered for comparison.

Using a set of $T = 1000$ equispaced points, we compute the average squared difference between the single machine and all other alternative methods for each observation in the dataset. To illustrate, for LISA we estimate ω using $\hat{\omega}_{LISA}^2 = \frac{1}{T} \sum_{j=1}^T (F_{LISA}(t_j) - F_{BART}(t_j))^2$.

Figure 2 is comparing the fitted polynomial trends of $\hat{\omega}^2$ (in each method) versus mean predicted $\hat{f}(x)$ in SingleMachine with their corresponding 95% credible regions (for both train and test data). Clearly in LISA and modLISA, there are small variations around the trends with no significant changes in values of $\hat{\omega}^2$ among different mean predicted $\hat{f}(x)$, which specifies consistency within different train or test observations. In addition, the gap between trends from train and test data indicate that the average distance between LISA/modLISA and SingleMachine's distributions are smaller for test data compared to train data. Furthermore, there are still small variations seen around CMC's trends, but with slight changes in values of $\hat{\omega}^2$ among different mean predicted $\hat{f}(x)$, especially for the test dataset which indicates inconsistency within different observations.

To emphasize the difference in performance between modLISA and its competitors, Figure 3 shows all the fitted polynomial trends without their credible regions for the train and test data. One can see that there is a large gap between $\hat{\omega}^2$ values in modLISA weighted average and other alternative methods (for both train and test data), with modLISA having the lowest value. Thus the weighted average of samples produced by modLISA yields the closest results to SingleMachine. This can also be justified by comparing average $\hat{\omega}^2$ over all train observations for each trend which is calculated to be

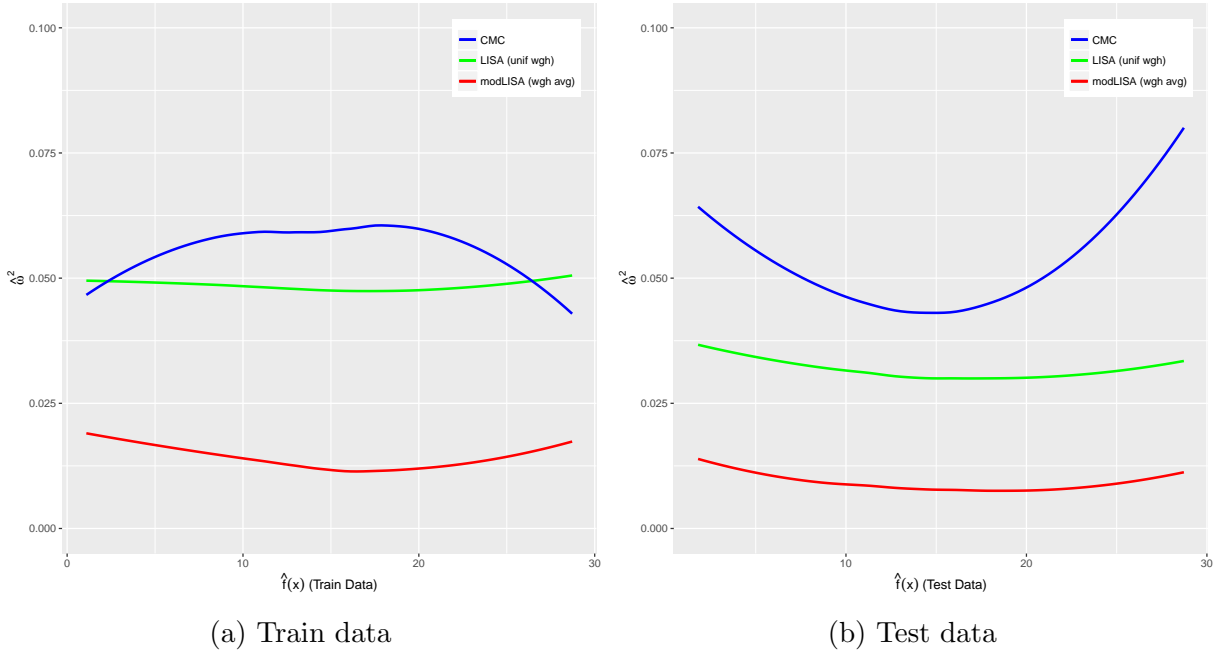


Figure 3: Comparing fitted polynomial trends of average squared difference in empirical distribution functions of each method and SingleMachine, as functions of mean predicted $\hat{f}(x)$ in SingleMachine (for both train and test data).

0.013 for modLISA that is significantly smaller than 0.059, 0.048 for CMC, and LISA, respectively. Similarly, the average $\hat{\omega}^2$ over test data are 0.008, 0.047, and 0.031 for modLISA, CMC, and LISA respectively, which again the smallest value is seen in modLISA. We conclude that modLISA weighted average sample yields the closest representation of the BART posterior and exhibits the best performance compared to alternative methods.

At last we compare run time per iteration for each method so we can draw some conclusions regarding the overall efficiency.

6.1.3 Run Time Comparisons

The main goal of methods such as LISA and CMC was to reduce run times regarding big data applications. Here we have compared average run times per iteration (from one processor) for each method using our implementation of BART.

As it is seen in Table 5, modLISA, LISA and CMC with $K = 30$ are all faster com-

Method	Avg Time per iteration (Secs)	Speed-up
<i>CMC</i>	11.99	31%
<i>LISA</i>	5.04	71%
<i>modLISA</i>	1.81	90%
<i>SingleMachine</i>	17.28	—

Table 5: Running times for CMC, LISA, modLISA and SingleMachine when $K = 30$.

pared to SingleMachine since they are influenced by the smaller subsets of data used. However, since LISA and CMC generate much larger trees, they become slower compared to modLISA which is the fastest method. We have also reported the speed-up percentages with respect to SingleMachine, which is defined to be $(1 - t/17.28) \times 100\%$ where t is the average time per iteration in each method. Clearly, CMC shows the smallest speed-up (31%) while modLISA has the highest (90%).

6.2 Additional Considerations

6.2.1 Effect of N (number of training data) on Posterior Accuracy

To see how the number of training data (N) can effect the posterior accuracy, we have examined the performance of all methods when N is increased to 60,000 while we keep the same number of batches $K = 30$. Tables 6 shows the results of 1000 posterior samples generated from fitting the BART model to the training set with additional 5000 data considered as test cases.

Unsurprisingly, Tables 1 and 6 show that the RMSE for training and test data in LISA, modLISA, and SingleMachine decrease as N increases. More importantly, while LISA and CMC estimates for σ^2 get worse, modLISA generates more accurate estimates of σ^2 with a larger N .

Trees have a stable size in modLISA, but tend to grow larger in CMC and LISA, as N increases. Table 7 shows that coverage of PI decreases in CMC and LISA, but increases in

Method	TrainRMSE	TestRMSE	Tree Nodes	Avg $\hat{\sigma}^2$	95% CI for σ^2
<i>CMC</i>	2.85	5.56	983	0.48	[0.30 , 0.66]
<i>LISA (unif wgh)</i>	1.17	1.19	125	0.0003	[0.00031 , 0.00035]
<i>modLISA (wgh avg)</i>	0.41	0.42	7	8.82	[8.79 , 8.86]
<i>SingleMachine</i>	0.41	0.41	11	9.04	[8.94 , 9.16]

Table 6: Tree sizes, estimates and 95% credible intervals for σ^2 , RMSE for training data (TrainRMSE) of size $N = 60,000$ and for test data (TestRMSE) of size 5,000 for each method run with $K = 30$.

Method	TrainPredCov	TestPredCov	TrainCredCov	TestCredCov
<i>CMC</i>	25.74 %	17.28 %	51.37 %	85.92 %
<i>LISA (unif wgh)</i>	0.84 %	0.84 %	100 %	100 %
<i>modLISA (wgh avg)</i>	94.54 %	94.53 %	53.68 %	52.68 %
<i>SingleMachine</i>	94.83 %	94.84 %	57.79 %	58.90 %

Table 7: Average coverage for 95% credible intervals constructed for training (TrainCredCov) and test (TestCredCov) data and 95% prediction intervals constructed for training (TrainPredCov) and test (TestPredCov) data. The prediction interval coverage is estimated based on 1000 iid samples, $N = 60,000$ and $K = 30$.

modLISA and SingleMachine for larger training data. We find it particularly promising that modLISA competes with SingleMachine for larger N . Note that, coverage of CI in LISA and CMC are still unreliable because of their over-dispersion, while in modLISA and SingleMachine they decrease as N increases, which is reasonable since larger sample size creates narrow CI that are around a biased $f(x)$ estimate, as discussed in the previous section. Overall, as N increases, modLISA seems to be a more reliable method as it shows a better performance compared to all other alternatives.

6.2.2 Effect of K (number of batches) on Posterior Accuracy

To examine the effect of K on posterior accuracy, we have generated 1000 posterior draws for training data of size $N = 20,000$ and $K = 10$. The test data sample is of size 5,000.

The results are shown in Table 8.

As K decreases, the performance of LISA and CMC drops while modLISA generates

Method	TrainRMSE	TestRMSE	Tree Nodes	Avg $\hat{\sigma}^2$	95% CI for σ^2
<i>CMC</i>	2.92	3.18	951	0.73	[0.57 , 0.90]
<i>LISA (unif wgh)</i>	1.70	1.78	131	0.001	[0.0010 , 0.0012]
<i>modLISA (wgh avg)</i>	0.46	0.47	7	8.69	[8.61 , 8.77]
<i>SingleMachine</i>	0.55	0.56	7	9.04	[8.85 , 9.21]

Table 8: Tree sizes, estimates and 95% credible intervals for σ^2 , RMSE for training data (TrainRMSE) of size $N = 20,000$ and for test data (TestRMSE) of size 5,000 for each method run with $K = 10$.

Method	TrainPredCov	TestPredCov	TrainCredCov	TestCredCov
<i>CMC</i>	31.08 %	29.83 %	48.18 %	99.80 %
<i>LISA (unif wgh)</i>	1.44 %	1.43 %	99.98 %	99.96 %
<i>modLISA (wgh avg)</i>	94.30 %	94.29 %	71.08 %	70.32 %
<i>SingleMachine</i>	94.67 %	94.65 %	71.58 %	71.54 %

Table 9: Average coverage for 95% credible intervals constructed for training (TrainCredCov) and test (TestCredCov) data and 95% prediction intervals constructed for training (TrainPredCov) and test (TestPredCov) data. The prediction interval coverage is estimated based on 1000 iid samples, $N = 20,000$ and $K = 10$.

stronger results, which is intuitively expected as each batch is larger and closer to the full sample when K is smaller. We also note the improvement of modLISA over SingleMachine in terms of RMSE. In addition, Table 9 shows that the PI and CI coverages for modLISA and SingleMachine are very close.

6.3 Varying the Underlying Model – Different $f(x)$

Consistency in performance of modLISA can also be seen when the underlying model is changed. For instance, we also considered a sample of size 20,000 using

$$f(x) = 3\sqrt{x_1} - 2x_2^2 + 5x_3x_4, \quad (17)$$

where $x = (x_1, \dots, x_4)$ is a four-dimensional input vector that is simulated independently from a $U(0, 1)$ and $y \sim \mathcal{N}(f(x), \sigma^2)$ with $\sigma^2 = 1$. Additional 5000 data have also been

simulated as test cases. Similarly, by fitting this newly simulated dataset to each method with $K = 30$, we have generated 1000 posterior samples with results averaged across three different realizations of data shown in Tables 10 and 11.

Method	TrainRMSE	TestRMSE	Tree Nodes	Avg $\hat{\sigma}^2$	95% CI for σ^2
<i>CMC</i>	0.89	0.76	614	0.21	[0.18 , 0.34]
<i>LISA (unif wgh)</i>	0.32	0.33	57	0.0001	[0.000083 , 0.000103]
<i>modLISA (wgh avg)</i>	0.11	0.11	7	0.88	[0.87 , 0.89]
<i>SingleMachine</i>	0.14	0.14	7	1.00	[0.99 , 1.03]

Table 10: Tree sizes, estimates and 95% credible intervals for σ^2 , RMSE for training data (TrainRMSE) of size $N = 20,000$ generated from (17) and for test data (TestRMSE) of size 5,000 for each method run with $K = 30$. Results are averaged over three different data replications.

Method	TrainPredCov	TestPredCov	TrainCredCov	TestCredCov
<i>CMC</i>	49.74 %	52.97 %	84.16 %	100 %
<i>LISA (unif wgh)</i>	1.50 %	1.49 %	100 %	100 %
<i>modLISA (wgh avg)</i>	93.07 %	93.18 %	82.88 %	83.50 %
<i>SingleMachine</i>	94.82 %	94.81 %	79.13 %	78.47 %

Table 11: Average coverage for 95% credible intervals constructed for training (TrainCredCov) and test (TestCredCov) data and 95% prediction intervals constructed for training (TrainPredCov) and test (TestPredCov) data generated from (17). The prediction interval coverage is estimated based on 1000 iid samples, $N = 20,000$ and $K = 30$. Results are averaged over three different data replications.

Again modLISA outperforms all alternative methods, and its performance is closest to SingleMachine. This confirms the previous simulation results and allows us to conclude that modLISA is a more reliable method for BART models with large datasets.

In the next section we will apply modLISA weighted average BART to a large socio-economic study.

6.4 Real Data Analysis

The American Community Survey (ACS) is a growing survey from the US Census Bureau and the Public Use Microdata Sample (PUMS) is a sample of responses to ACS which consists of various variables related to people and housing units (see US Bureau of Census, 2013). Considering the person-level data from PUMS 2013, we would like to predict a person’s total income based on variables such as sex, age, education, class of worker, living state, and citizenship status. We have collected information related to people who are employed and have total income of at least \$5000 with education level of either Bachelor’s degree, Master’s degree, or a PhD which resulted in 437,297 observations. We randomly divided the dataset into approximately 80% training and 20% testing sets, with $K = 100$ batches considered for splitting the training data to apply modLISA. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium (Loken et al., 2010) using 100 cores, each running on 3,500 observations. Considering the logarithm of total income for each person as the response variable, we have ran modLISA with weighted average and SingleMachine BART on this dataset for 1500 iterations (since SingleMachine is very slow) and discarded the first 1000 draws which resulted in 500 posterior samples. Table 12 contains the results of Test RMSE as well as average post burn-in σ^2 estimates and tree sizes.

Method	TestRMSE	Avg $\hat{\sigma}^2$	Tree Nodes	Speed-up
<i>modLISA (wgh avg)</i>	0.71	0.488	7	90%
<i>SingleMachine</i>	0.70	0.485	23	–

Table 12: Performance summaries computed from 1000 posterior samples generated from modLISA with $K = 100$ and SingleMachine BART on PUMS 2013 test data.

One can see that Test RMSE in modLISA is similar to the one from SingleMachine, but with a 90% speed-up of modLISA over SingleMachine. The speed-up can be explained by the larger acceptance probabilities and by the smaller tree sizes reported in Tables 13 and

Method	GROW	PRUNE	CHANGE
<i>modLISA</i>	10 %	11 %	14 %
<i>SingleMachine</i>	8%	7%	7%

Table 13: Average acceptance rates of tree proposal moves.

12, respectively. The 90% speedup is important for applications like the one considered here, as it takes more than a day to simulate 1,500 samples from the posterior using *SingleMachine*. The result indicate the potential of the proposed method for reducing computational costs while producing accurate predictions.

7 Discussion

The challenge of using MCMC algorithms to sample posterior distributions obtained from a massive sample of observations is a serious one.

In this paper, we introduced a new method based on the idea of randomly dividing the data into batches and drawing samples from each of the resulting sub-posteriors independently and in parallel on different machines. We propose a novel way to define the sub-posteriors and we develop a strategy to combine the samples produced by each batch analysis for the important class of Bayesian Additive Regression Trees Models. For this model, the proposed methodology performs very well and shows reduction in computation time that are as high as 90%.

In future work we would like to find a procedure for combining the sub-posterior samples that will make *LISA* easy to adapt to a wide variety of models. We also hope that our paper will stimulate the research into this type of divide-and-conquer approaches for Big Data MCMC and will expand the research on how to construct the batch-specific sub-posteriors along with novel strategies of combining or weighting the samples obtained from each batch analysis.

8 Acknowledgement

We thank the Editor, Grace Yi, the Guest Editor, Richard Lockhart, the Associate Editor and three anonymous referees for helpful suggestions that have greatly improved the paper. This work has been supported by NSERC of Canada grants to RVC and JSR.

References

- S. Brooks, A. Gelman, G. L. Jones, and X. Meng, editors. *Handbook of Markov Chain Monte Carlo*. Taylor & Francis, 2011.
- Hugh A Chipman, Edward I George, and Robert E McCulloch. Bayesian CART model search. *Journal of the American Statistical Association*, 93(443):935–948, 1998.
- Hugh A Chipman, Edward I George, and Robert E McCulloch. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, pages 266–298, 2010.
- Radu V Craiu and Jeffrey S Rosenthal. Bayesian computation via Markov chain Monte Carlo. *Annual Review of Statistics and Its Application*, 1:179–201, 2014.
- Jerome H Friedman. Multivariate adaptive regression splines. *The annals of statistics*, pages 1–67, 1991.
- Adam Kapelner and Justin Bleich. bartmachine: Machine learning with Bayesian additive regression trees. *arXiv preprint arXiv:1312.2171*, 2013.
- Kathryn Blackmond Laskey and James W Myers. Population Markov chain Monte Carlo. *Machine Learning*, 50(1-2):175–196, 2003.
- Chris Loken, Daniel Gruner, Leslie Groer, Richard Peltier, Neil Bunn, Michael Craig, Teresa Henriques, Jillian Dempsey, Ching-Hsing Yu, Joseph Chen, et al. Scinet: lessons

- learned from building a power-efficient top-20 system and data centre. In *Journal of Physics: Conference Series*, volume 256, page 012026. IOP Publishing, 2010.
- Willie Neiswanger, Chong Wang, and Eric Xing. Asymptotically exact, embarrassingly parallel MCMC. *arXiv preprint arXiv:1311.4780*, 2013.
- Matthew T Pratola et al. Efficient Metropolis–Hastings proposal mechanisms for Bayesian regression tree models. *Bayesian Analysis*, 11(3):885–911, 2016.
- Jeffrey S Rosenthal. Parallel computing and Monte Carlo algorithms. *Far east journal of theoretical statistics*, 4(2):207–236, 2000.
- Steven L Scott, Alexander W Blocker, Fernando V Bonassi, H Chipman, E George, and R McCulloch. Bayes and big data: The consensus Monte Carlo algorithm. In *International Journal of Management Science and Engineering Management*, 11(2):78–88, 2016.
- US Bureau of Census. 2013 ACS 1-YEAR PUMS data. URL <http://www.census.gov/programs-surveys/acs/data/pums.html>.
- Xiangyu Wang and David B Dunson. Parallelizing MCMC via Weierstrass sampler. *arXiv preprint arXiv:1312.4605*, 2013.
- Darren J Wilkinson. Parallel Bayesian computation. *Statistics Textbooks and Monographs*, 184:477, 2006.