



sgmcmc: An R Package for Stochastic Gradient Markov Chain Monte Carlo

Jack Baker
Lancaster University

Paul Fearnhead
Lancaster University

Emily B. Fox
University of Washington

Christopher Nemeth
Lancaster University

Abstract

This paper introduces the R package **sgmcmc**; which can be used for Bayesian inference on problems with large data sets using stochastic gradient Markov chain Monte Carlo (SGMCMC). Traditional Markov chain Monte Carlo (MCMC) methods, such as Metropolis-Hastings, are known to run prohibitively slowly as the data set size increases. SGMCMC solves this issue by only using a subset of data at each iteration. SGMCMC requires calculating gradients of the log-likelihood and log-priors, which can be time consuming and error prone to perform by hand. The **sgmcmc** package calculates these gradients itself using automatic differentiation, making the implementation of these methods much easier. To do this, the package uses the software library **TensorFlow**, which has a variety of statistical distributions and mathematical operations as standard, meaning a wide class of models can be built using this framework. SGMCMC has become widely adopted in the machine learning literature, but less so in the statistics community. We believe this may be partly due to lack of software; this package aims to bridge this gap.

Keywords: R, stochastic gradient Markov chain Monte Carlo, big data, MCMC, stochastic gradient Langevin dynamics, stochastic gradient Hamiltonian Monte Carlo, stochastic gradient Nosé-Hoover thermostat.

1. Introduction

This article introduces **sgmcmc**, an R package (R Core Team 2019) for scalable Bayesian inference on a wide variety of models using stochastic gradient Markov chain Monte Carlo (SGMCMC, Baker, Nemeth, Fearnhead, and Fox 2019). A disadvantage of most traditional

Markov chain Monte Carlo (MCMC) methods are that they require calculations over the full data set at each iteration; meaning the methods are prohibitively slow for large data sets. SGMCMC methods provide a solution to this issue. The methods use only a subset of the full data set, known as a mini batch, at each MCMC iteration. While the methods no longer target the true posterior, they instead produce accurate approximations to the posterior at a reduced computational cost.

The **sgmcmc** package (Baker *et al.* 2019) implements a number of popular SGMCMC samplers including stochastic gradient Langevin dynamics (SGLD; Welling and Teh 2011), stochastic gradient Hamiltonian Monte Carlo (SGHMC; Chen, Fox, and Guestrin 2014) and stochastic gradient Nosé-Hoover thermostats (SGNHT; Ding, Fang, Babbush, Chen, Skeel, and Neven 2014). Recent work has shown how control variates can be used to reduce the computational cost of SGMCMC algorithms (Baker, Fearnhead, Fox, and Nemeth 2017; Nagapetyan, Duncan, Hasenclever, Vollmer, Szpruch, and Zygalakis 2017). For each of the samplers implemented in the package, there is also a corresponding control variate sampler providing improved sampling efficiency.

Performing statistical inference on a model using SGMCMC requires calculating the gradient of the log-likelihood and log-priors. Calculating gradients by hand is often time consuming and error prone. One of the major advantages of **sgmcmc** is that gradients are calculated within the package using automatic differentiation (Griewank and Walther 2008). This means that users need only specify the log-likelihood function and log-prior for their model. The package calculates the gradients using TensorFlow (Abadi *et al.* 2015).

TensorFlow is a machine learning system that operates at large scale and in heterogeneous environments. The **tensorflow** (Allaire and Tang 2019) package offers an interface from R to TensorFlow. TensorFlow is an efficient library for numerical computation which can take advantage of a wide variety of architectures, as such, **sgmcmc** keeps much of this efficiency. Both **sgmcmc** and **tensorflow** are available on Comprehensive R Archive Network (CRAN), so **sgmcmc** can be installed by using the standard `install.packages` function. Though after the **tensorflow** (Allaire and Tang 2019) package has been installed, the extra `install_tensorflow()` function needs to be run, which installs the required Python (Van Rossum *et al.* 2011) implementation of TensorFlow¹. The **sgmcmc** package also has a website with vignettes, tutorials and an API reference².

SGMCMC methods have become popular in the machine learning literature but less so in the statistics community. We partly attribute this to the lack of available software. To the best of our knowledge, there are currently no R packages available for SGMCMC, probably the most popular programming language within the statistics community. The only package we are aware of which implements scalable MCMC is the Python (Van Rossum *et al.* 2011) package **Edward** (Tran, Kucukelbir, Dieng, Rudolph, Liang, and Blei 2016). This package implements both SGLD and SGHMC, but does not implement SGNHT or any of the control variate methods.

Section 2 introduces MCMC and discusses the software currently available for implementing MCMC algorithms, we discuss the scenarios where **sgmcmc** is designed to be used. In Section 3 we review the methodology behind the SGMCMC methods implemented in **sgmcmc**. Section 4

¹More information on installing **tensorflow** (Allaire and Tang 2019) for R can be found at <https://tensorflow.rstudio.com/>.

²See the **sgmcmc** website at <https://stor-i.github.io/sgmcmc>.

provides a brief introduction to TensorFlow (Abadi *et al.* 2015) and `tensorflow` (Allaire and Tang 2019). Section 5 overviews the structure of the package, as well as details of how the algorithms are implemented. Section 6 presents a variety of example simulations. Finally, Section 7 provides a discussion on benefits and drawbacks of the implementation, as well as how we plan to extend the package in the future.

2. Introduction to MCMC and available software

Suppose we have a data set of size N , with data $\mathbf{x} = \{x_i\}_{i=1}^N$, where $x_i \in \mathcal{X}$ for some space \mathcal{X} . We denote the probability density of x_i as $p(x_i|\theta)$, where $\theta \in \Theta \subseteq \mathbb{R}^p$ are model parameters to be inferred. We assign a prior density $p(\theta)$ to the parameters and the resulting posterior is then $p(\theta|\mathbf{x}) = p(\theta) \prod_{i=1}^N p(x_i|\theta)/Z$.

Often the posterior can only be calculated up to a constant of proportionality Z . In practice Z is rarely analytically tractable; so MCMC provides a way to construct a Markov chain using only the unnormalized posterior density $h(\theta) := p(\theta) \prod_{i=1}^N p(x_i|\theta)$. The Markov chain is designed so that its stationary distribution is the posterior $p(\theta|\mathbf{x})$. The result (once the chain has converged) is a sample $\{\theta_t\}_{t=1}^T$ from the posterior, though this sample is not independent. A downside of these traditional MCMC algorithms is that they require the evaluation of the unnormalized density $h(\theta)$ at every iteration. This results in an $O(N)$ cost per iteration. Thus MCMC becomes prohibitively slow on large data sets.

The Metropolis-Hastings algorithm is a type of MCMC algorithm. New proposed samples θ' are drawn from a proposal distribution $q(\theta'|\theta)$ and then accepted with probability,

$$\min \left\{ 1, \frac{p(\theta'|\mathbf{x})q(\theta|\theta')}{p(\theta|\mathbf{x})q(\theta'|\theta)} \right\}. \quad (1)$$

Notice that the normalising constant Z cancels in (1), so we can interchange the posterior $p(\theta|\mathbf{x})$ with $h(\theta)$. The efficiency of the Metropolis-Hastings algorithm is dependent on the choice of proposal distribution, q .

There are a number of proposals for the Metropolis-Hastings algorithm which can have a very high acceptance rate. Some examples are the Metropolis-adjusted Langevin algorithm (MALA; see, e.g., Roberts and Rosenthal 1998) and Hamiltonian Monte Carlo (HMC; see Neal 2010). The reason these proposals achieve such high acceptance rates is that they approximate a continuous diffusion process whose stationary distribution is $p(\theta|\mathbf{x})$. As an example, consider the MALA algorithm. The MALA algorithm uses a Euler discretisation of the Langevin diffusion as the proposal,

$$q(\theta'|\theta) = \mathcal{N} \left(\theta' | \theta + \frac{\epsilon}{2} \nabla_{\theta} \log p(\theta|\mathbf{x}), \epsilon \mathbf{I} \right), \quad (2)$$

where $\mathcal{N}(\theta|\mu, \Sigma)$ denotes a multivariate normal density evaluated at θ with mean μ and variance Σ ; \mathbf{I} is simply the identity matrix; ϵ is a tuning parameter referred to as the step size. Discretising the diffusion introduces an approximation error, which is corrected by the Metropolis-Hastings acceptance step (1). This means that as $\epsilon \rightarrow 0$, we tend back towards the exact, continuous diffusion and the acceptance rate is 1. However this would result in a Markov chain that never moves. Thus picking ϵ is a balance between a high acceptance rate and good mixing.




Figure 1: KL divergence (left) and run time (right) of the standard Stan algorithm and the `sgldcv` algorithm of the **sgmcmc** package when each are used to sample from data following a standard normal distribution as the number of observations are increased.

There are a number of general purpose samplers for MCMC that fulfill different purposes to **sgmcmc**. The most popular samplers are Stan, BUGS and JAGS (Carpenter *et al.* 2017; Lunn, Thomas, Best, and Spiegelhalter 2000; Lunn, Spiegelhalter, Thomas, and Best 2009; Plummer 2003). The samplers BUGS and JAGS implement automated Gibbs sampling. These samplers work with both continuous and discrete parameter spaces and can be highly efficient. However because the samplers rely on Gibbs sampling, conjugate priors need to be used; also the samplers are not efficient when there is high correlation between the parameters (Carpenter *et al.* 2017). The Stan software implements state of the art HMC techniques, which means non-conjugate priors can be used, and that the sampler is more robust when there is correlation between parameters. However the package cannot perform inference on discrete parameters, and requires that these are integrated out of the model.

The **sgmcmc** package aims to fill a gap when the data set is large enough that other general purpose MCMC samplers such as Stan, BUGS and JAGS cannot be run or run prohibitively slowly. In the R interface packages **rstan**, **R2WinBUGS** and **rjags** (Stan Development Team 2018; Sturtz, Ligges, and Gelman 2005; Plummer 2018), properly specified models will define a Markov chain whose stationary distribution is the posterior distribution. However a major problem with these methods are that as the number of observations gets large the algorithms run slowly. This has become a problem as data set sizes have been increasing. The reason these methods are slow when running on large data sets are because they require a calculation over the full data set at each iteration. The methods implemented in **sgmcmc** aim to account for this issue by only using a subset of the data set at each iteration. The main downside being that the stationary distribution is no longer the true posterior, just a close approximation. However, as the data set size increases, the main tuning constant in **sgmcmc**, known as the step size, can be set smaller and the approximation to the posterior improves. Compared to Stan, BUGS and JAGS; **sgmcmc** offers significant computational advantages for Bayesian modeling with large data sets, but like **rstan**, a downside of **sgmcmc** is it requires that discrete parameters are integrated out. The package also requires more tuning than other general purpose samplers since satisfactory results for tuning these methods are still under development.

Figure 1 demonstrates in which scenarios practitioners may find **sgmcmc** useful. The standard Stan sampler and the `sgldcv` algorithm of the **sgmcmc** package are used to sample from the posterior of data drawn from a standard normal $N(0, 1)$ with a $N(0, 10)$ prior. The Kullback-Leibler (KL) divergence between the MCMC sample and the true posterior is calculated, and

the plots show how this and the run time changes as the number of observations are increased. Since both Stan and TensorFlow models need to be compiled, we recompile the models each time they are run to keep the comparison fair; but it is worth mentioning that the Stan run time is much quicker for the small observation models if precompiled. We can see the run time of Stan increasing rapidly as the number of observations is increased, while the run time of the SGMCMC algorithm increases more slowly. This is a very simple model, used so that the KL divergence can be calculated exactly. As the model complexity increases the run time of Stan can quickly become unmanageable for large data sets. On the other hand, we can see that the KL divergence of the SGMCMC algorithm for this example is poor compared with Stan for a small number of observations. However, as the data set size increases, the KL divergence for the SGMCMC algorithm becomes much more reasonable compared with Stan. Thus the `sgmcmc` package is best used when the data set size is slowing down the run time of the other general purpose algorithms, and practitioners can safely trade-off a small amount of accuracy in order to gain significant speed-ups.

3. Stochastic gradient MCMC

Many popular MCMC proposal distributions, including HMC and MALA, described in (2), require the calculation of the gradient of the log posterior at each iteration, which is an $O(N)$ calculation. Stochastic gradient MCMC methods get around this by replacing the true gradient with the following unbiased estimate

$$\nabla_{\theta} \widehat{\log p(\theta_t | \mathbf{x})} := \nabla_{\theta} \log p(\theta_t) + \frac{N}{n} \sum_{i \in S_t} \nabla_{\theta} \log p(x_i | \theta_t), \quad (3)$$

calculated on some subset of the all observations $S_t \subset \{1, \dots, N\}$, known as a mini batch, with $|S_t| = n$.

Calculating the Metropolis-Hastings acceptance step (1) is another $O(N)$ calculation. To get around this, SGMCMC methods set the tuning constants such that the acceptance rate will be high, and remove the acceptance step from the algorithm altogether. By ignoring the acceptance step, and estimating the log posterior gradient, the per iteration cost of SGMCMC algorithms is $O(n)$, where n is the mini batch size. However, the algorithm no longer targets the true posterior but an approximation. There has been a body of theory exploring how these methods perform in different settings. Similar to MALA, the algorithms rely on a step size parameter ϵ . Some of the algorithms have been shown to weakly converge as $\epsilon \rightarrow 0$.

3.1. Stochastic gradient Langevin dynamics

SGLD, first introduced by [Welling and Teh \(2011\)](#), is an SGMCMC approximation to the MALA algorithm. By substituting (3) into the MALA proposal equation (2), we arrive at the following update for θ

$$\theta_{t+1} = \theta_t + \frac{\epsilon_t}{2} \nabla_{\theta} \widehat{\log p(\theta_t | \mathbf{x})} + \zeta_t, \quad (4)$$

where $\zeta_t \sim \mathcal{N}(0, \epsilon_t)$.

[Welling and Teh \(2011\)](#) noticed that as $\epsilon_t \rightarrow 0$ this update will sample from the true posterior. Although the algorithm mixes slower as ϵ gets closer to 0, so setting the step size is a trade-off

between convergence speed and accuracy. This motivated [Welling and Teh \(2011\)](#) to suggest setting ϵ_t to decrease to 0 as t increases. There is a body of theory that investigates the SGLD approximation to the true posterior (see, e.g., [Teh, Thiéry, and Vollmer 2016](#); [Sato and Nakagawa 2014](#); [Vollmer, Zygalkis, and Teh 2016](#)). In particular, SGLD is found to converge weakly to the true posterior distribution asymptotically as $\epsilon_t \rightarrow 0$. The mean squared error of the algorithm is found to decrease at best with rate $O(T^{-1/3})$. In practice, the algorithm tends to mix poorly when ϵ is set to decrease to 0 ([Vollmer et al. 2016](#)), so in our implementation we use a fixed step size which has been shown to mix better empirically. Theoretical analysis for this case is provided in [Vollmer et al. \(2016\)](#). The tuning constant ϵ , referred to as the step size is a required argument in the package. It affects the performance of the algorithm considerably.

3.2. Stochastic gradient Hamiltonian Monte Carlo

The stochastic gradient Hamiltonian Monte Carlo algorithm (SGHMC; [Chen et al. 2014](#)) is similar to SGLD, but instead approximates the HMC algorithm ([Neal 2010](#)). To ensure SGHMC is $O(n)$, the same unbiased estimate to the log posterior gradient is used (3). SGHMC augments the parameter space with momentum variables ν and samples approximately from a joint distribution $p(\theta, \nu | \mathbf{x})$, whose marginal distribution for θ is the posterior of interest. The algorithm performs the following updates at each iteration

$$\begin{aligned}\theta_{t+1} &= \theta_t + \nu_t, \\ \nu_{t+1} &= (1 - \alpha)\nu_t + \epsilon \nabla_{\theta} \widehat{\log p(\theta_{t+1} | \mathbf{x})} + \zeta_t,\end{aligned}$$

where $\zeta_t \sim \mathcal{N}(0, 2[\alpha - \hat{\beta}_t]\epsilon)$; ϵ and α are tuning constants and $\hat{\beta}_t$ is proportional to an estimate of the Fisher information matrix. In our current implementation, we simply set $\hat{\beta}_t := 0$, as in the experiments of the original paper by [Chen et al. \(2014\)](#). In future implementations, we aim to estimate $\hat{\beta}_t$ using a Fisher scoring estimate similar to [Ahn, Korattikara, and Welling \(2012\)](#). Often the dynamics are simulated repeatedly L times before the state is stored, at which point ν is resampled. The parameter L can be included in our implementation. The tuning constant ϵ is the step size and is a required argument in our implementation, as for SGLD its value affects performance considerably. The constant α tends to be fixed at a small value in practice. As a result, in our implementation it is an optional argument with default value 0.01.

3.3. Stochastic gradient Nosé-Hoover thermostat

[Ding et al. \(2014\)](#) suggest that the quantity $\hat{\beta}_t$ in SGHMC is difficult to estimate in practice. They appeal to analogues between these proposals and statistical physics in order to suggest a set of updates which do not need this estimation to be made. Once again [Ding et al. \(2014\)](#) augment the space with momentum parameters ν . They replace the tuning constant α with a dynamic parameter α_t known as the thermostat parameter. The algorithm performs the following updates at each iteration

$$\theta_{t+1} = \theta_t + \nu_t, \tag{5}$$

$$\nu_{t+1} = (1 - \alpha_t)\nu_t + \epsilon \nabla_{\theta} \widehat{\log p(\theta_{t+1} | \mathbf{x})} + \zeta_t, \tag{6}$$

$$\alpha_{t+1} = \alpha_t + \left[\frac{1}{p} (\nu_{t+1})^{\top} (\nu_{t+1}) - \epsilon \right]. \tag{7}$$

Here $\zeta_t \sim \mathcal{N}(0, 2a\epsilon)$; ϵ and a are tuning parameters to be chosen and p is the dimension of θ . The update for α in (7) requires a vector dot product, it is not obvious how to extend this when θ is higher order than a vector, such as a matrix or tensor. In our implementation, when θ is a matrix or tensor we use the standard inner product in those spaces (Abraham, Marsden, and Ratiu 1988). The tuning constant ϵ is the step size and is a required argument in our implementation, as again its value affects performance considerably. The constant a , similarly to α in SGHMC, tends to be fixed at a small value in practice (Ding *et al.* 2014). As a result, in our implementation it is an optional argument with default value 0.01.

3.4. Stochastic gradient MCMC with control variates

A key feature of SGMCMC methods is replacing the log posterior gradient calculation with an unbiased estimate. The unbiased gradient estimate, which can be viewed as a noisy version of the true gradient, can have high variance when estimated using a small mini batch of the data. Increasing the mini batch size will reduce the variance of the gradient estimate, but increase the per iteration computational cost of the SGMCMC algorithm. Recently control variates (Ripley 2009) have been used to reduce the variance in the gradient estimate of SGMCMC (Dubey, Reddi, Williamson, Póczos, Smola, and Xing 2016; Nagapetyan *et al.* 2017; Baker *et al.* 2017). Using these improved gradient estimates have been shown to lead to improvements in the mean squared error (MSE) of the algorithm (Dubey *et al.* 2016), as well as its computational cost (Nagapetyan *et al.* 2017; Baker *et al.* 2017).

We implement the formulation of Baker *et al.* (2017), who replace the gradient estimate $\nabla_{\theta} \widehat{\log p(\theta|\mathbf{x})}$ with

$$\nabla_{\theta} \widetilde{\log p(\theta|\mathbf{x})} := \nabla_{\theta} \log p(\hat{\theta}|\mathbf{x}) + \nabla_{\theta} \widehat{\log p(\theta|\mathbf{x})} - \nabla_{\theta} \widehat{\log p(\hat{\theta}|\mathbf{x})}, \quad (8)$$

where $\hat{\theta}$ is an estimate of the posterior mode. This method requires the burn-in phase of MCMC to be replaced by an optimization step which finds $\hat{\theta} := \operatorname{argmax}_{\theta} \log p(\theta|\mathbf{x})$. There is then an $O(N)$ preprocessing step to calculate $\nabla_{\theta} \log p(\hat{\theta}|\mathbf{x})$, after which the chain can be started from $\hat{\theta}$ resulting in a negligible mixing time. Baker *et al.* (2017) and Nagapetyan *et al.* (2017) have shown that there are considerable improvements to the computational cost of SGLD when (8) is used in place of (3). In particular they showed that standard SGLD requires setting the mini batch size n to be $O(N)$ for arbitrarily good performance; while using control variates requires an $O(N)$ preprocessing step, but after that a batch size of $O(1)$ can be used to reach the desired performance. Baker *et al.* (2017) also showed empirically that this particular formulation can lead to a reduction in the burn-in time compared with standard SGLD and the formulation of Dubey *et al.* (2016), which tended to get stuck in local stationary points. This is because in complex scenarios the optimization step is often faster than the burn-in time of SGMCMC. The package `sgmcmc` includes control variate versions of all the SGMCMC methods implemented: SGLD, SGHMC and SGNHT.

4. Brief tensorflow introduction

TensorFlow (Abadi *et al.* 2015) is a software library released by Google. The tool was initially designed to easily build deep learning models; but the efficient design and excellent implementation of automatic differentiation (Griewank and Walther 2008) has made it useful more

generally. This package is built on TensorFlow, and while we have tried to make the package as easy as possible to use, some knowledge of TensorFlow is unavoidable; especially when declaring the log-likelihood and log-prior functions, or for high dimensional chains which will not fit into memory. With this in mind, we provide a brief introduction to TensorFlow and in particular to the R interface package **tensorflow** (Allaire and Tang 2019) in this section. This should provide enough knowledge for the rest of the article. A more detailed introduction to TensorFlow for R can be found in Allaire and Tang (2019).

Any procedure written in **tensorflow** follows three main steps. The first step is to declare all the variables, constants and equations required to run the algorithm. In the background, these declarations enable **tensorflow** to build a graph of the possible operations, and how they are related to other variables, constants and operations. Once everything has been declared, the TensorFlow session is begun and all the variables and operations are initialized. Then the previously declared operations can be run as required; typically these will be sequential and will be run in a for loop.

4.1. Declaring tensorflow tensors

Everything in **tensorflow**, including operations, are represented as a tensor; which is basically a multi-dimensional array. There are a number of ways of creating tensors:

- `tf$constant(value)` – create a constant tensor with the same shape and values as `value`. The input `value` is generally an R array, vector or scalar. The most common use for this in the context of the package is for declaring constant parameters when declaring log-likelihood and log-prior functions.
- `tf$Variable(value)` – create a tensor with the same shape and values as `value`. Unlike `tf$constant`, this type of tensor can be changed by a declared operation. The input `value` is generally an R array, vector or scalar.
- `tf$placeholder(datatype, shape)` – create an empty tensor of type `datatype` and dimensions `shape` which can be fed all or part of a data set, this is useful when declaring operations which rely on data which can change. When you have storage constraints (see Section 5.2) you can use a placeholder to declare test functions that rely on a test data set. The `datatype` should be a **tensorflow** data type, such as `tf$float32`; the `shape` should be an R vector or scalar, such as `c(100, 2)`.
- `operation` – an operation declares an operation on previously declared tensors. These use **tensorflow** defined functions, such as those in its math library. This is essentially what you are declaring when coding the `logLik` and `logPrior` functions. The `params` input consists of a list of `tf$Variables`, representing the model parameters to be inferred. The `dataset` input consists of a list of `tf$placeholder` tensors, representing the mini batch of data fed at each iteration. Your job is to declare functions that return an operation on these tensors that define the log-likelihood and log-prior.

4.2. tensorflow operations

tensorflow operations take other tensors as input and manipulate them to reach the desired output. Once the TensorFlow session has begun, these operations can be run as needed, and

will use the current values for its input tensors. For example, we could declare a normal density tensor which manipulates a `tf$Variable` tensor representing the parameters and a `tf$placeholder` tensor representing the current data point. The tensor could then use the **tensorflow** `'tf$contrib$distributions$MultivariateNormalDiag'` object to return a tensor object of the current value for a normal density given the current parameter value and the data point that is fed to the placeholder. We can break this example down into three steps. First we declare the tensors that we require:

```
R> library("tensorflow")
R> loc <- tf$Variable(rep(0, 2))
R> dataPoint <- tf$placeholder(tf$float32, c(2))
R> scaleDiag <- tf$constant(c(1, 1))
R> distn <- tf$contrib$distributions$MultivariateNormalDiag(loc, scaleDiag)
R> dens <- distn$prob(dataPoint)
```

Here we have declared a `tf$Variable` tensor to hold the location parameter; a `tf$placeholder` tensor which will be fed the current data point; the scale parameter is fixed so we declare this as a `tf$constant` tensor. Next we declare the operation which takes the inputs we just declared and returns the normal density value. The first line which is assigned to `distn` creates a `'MultivariateNormalDiag'` object, which is linked to the `loc` and `scaleDiag` tensors. Then `dens` evaluates the density of this distribution. The `dens` variable is now linked to the tensors `dataPoint` and `scaleDiag`, so if it is evaluated it will use the current values of those tensors to calculate the density estimate. Next we begin the **TensorFlow** session:

```
R> sess <- tf$Session()
R> sess$run(tf$global_variables_initializer())
```

The two lines we just ran starts the **TensorFlow** session and initializes all the tensors we just declared. The **TensorFlow** graph has now been built and no new tensors can now be added. This means that all operations need to be declared before they can be evaluated. Now the session is started we can run the operation `dens` we declared given current values for `dataPoint` and `loc` as follows:

```
R> set.seed(1)
R> x <- rnorm(2)
R> out <- sess$run(dens, feed_dict = dict(dataPoint = x))
R> cat(paste0("Density value for x is ", out, ".\n"))
```

Density value for x is 0.128611162304878.

Since `dataPoint` is a placeholder, we need to feed it values each time. In the block of code above we feed `dataPoint` a random value simulated from a standard normal. The `sess$run` expression then evaluates the current normal density value given `loc` and `dataPoint`.

As mentioned earlier, this is essentially what is happening when you are writing the `logLik` and `logPrior` functions. These functions will be fed a list of `'tf$Variable'` objects to the `params` input, and a list of `'tf$placeholder'` objects to the `dataset` input. The output of the function will then be declared as a **TensorFlow** operation. This allows the gradient to be calculated automatically, while maintaining the efficiencies of **TensorFlow**.

Function	Algorithm
<code>sgld</code>	Stochastic gradient Langevin dynamics
<code>sghmc</code>	Stochastic gradient Hamiltonian Monte Carlo
<code>sgnht</code>	Stochastic gradient Nosé-Hoover thermostat
<code>sgldcv</code>	Stochastic gradient Langevin dynamics with control variates
<code>sghmccv</code>	Stochastic gradient Hamiltonian Monte Carlo with control variates
<code>sgnhtcv</code>	Stochastic gradient Nosé-Hoover thermostat with control variates

Table 1: Outline of 6 main functions implemented in **sgmcmc**.

tensorflow implements a lot of useful functions to make building these operations easier. For example a number of distributions are implemented at `tf$contrib$distributions`³ (e.g., `tf$contrib$distributions$Normal` and `tf$contrib$distributions$Gamma`). TensorFlow also has a comprehensive math library which provides a variety of useful tensor operations⁴. For examples of writing TensorFlow operations see the worked examples in Section 5 or the **sgmcmc** vignettes.

5. Package structure and implementation

The package has 6 main functions. The first three: `sgld`, `sghmc` and `sgnht` will implement SGLD, SGHMC and SGNHT, respectively. The other three: `sgldcv`, `sghmccv` and `sgnhtcv` implement the control variate versions of SGLD, SGHMC and SGNHT, respectively. All of these are built on the **tensorflow** library for R, which enables gradients to be automatically calculated and efficient computations to be performed in high dimensions. The usage for these functions is very similar, with the only differences in input being tuning parameters. These main functions are outlined in Table 1. The functions `sgld`, `sghmc` and `sgnht` have the same required inputs: `logLik`, `dataset`, `params` and `stepsize`. These determine respectively: the log-likelihood function for the model; the data for the model; the parameters of the model; and the step size tuning constants for the SGMCMC algorithm. The input `logLik` is a function taking `dataset` and `params` as input, while the rest are defined as lists. Using lists in this way provides a lot of flexibility: allowing multiple parameters to be defined; use multiple data sets; and use different step sizes for each parameter, which is vital if the scalings are different. It also allows users to easily reference parameters and data sets in the `logLik` function by simply referring to the relevant names in the list.

The functions also have a couple of optional parameters that are particularly important, `logPrior` and `minibatchSize`. These respectively define the log-prior for the model; and the mini batch size, as it was defined in Section 3. By default, the `logPrior` is set to an uninformative uniform prior, which is fine to use for quick checks but will need to be set properly for more complex models. The `logPrior` is a function similar to `logLik`, but only takes `params` as input. The `minibatchSize` is a numeric, and can either be a proportion of the data set size, if it is set between 0 and 1, or an integer. The default value is 0.01, meaning that 1% of the full data set is used at each iteration.

³For full API details see https://www.tensorflow.org/api_docs/python/tf/contrib/distributions though note this is for Python, so the `.` object notation needs to be replaced by `$`, for example `tf.contrib.distributions.Normal` would be replaced by `tf$contrib$distributions$Normal`.

⁴See https://www.tensorflow.org/api_guides/python/math_ops.

Function inputs	Definition
<code>logLik</code>	Log-likelihood function taking <code>dataset</code> and <code>params</code> as inputs
<code>dataset</code>	R list containing data
<code>params</code>	R list containing model parameters
<code>stepsize</code>	R list of step sizes for each parameter
<code>optstepsize</code>	R numeric step size for control variate optimization step
<code>logPrior</code>	Function of the parameters; default $p(\theta) \propto 1$
<code>minibatchSize</code>	Size of mini batch per iteration as integer or proportion; default 0.01.
<code>nIters</code>	Number of MCMC iterations; default is 10,000.

Table 2: Outline of the key arguments required by the functions in Table 1.

The control variate functions have the same inputs as the non-control variate functions, with one more required input. The `optstepsize` input is a numeric that specifies the step size for the initial optimization step to find the $\hat{\theta}$ maximizing the posterior as defined in Section 3.4. For a full outline of the key inputs, see Table 2.

Often large data sets and high dimensional problems go hand in hand. In these high dimensional settings storing the full MCMC chain in memory can become an issue. For this situation, we provide functionality to run the chain one iteration at a time in a user defined loop. This enables the user to deal with the output at each iteration how they see fit. For example, they may wish to calculate a test function on the output to reduce the dimensionality of the chain; or they might calculate the required Monte Carlo estimates on the fly. We aim to extend this functionality to enable the user to define their own Gibbs updates alongside the SGMCMC procedure. This functionality is more involved, and requires more knowledge of TensorFlow, so we leave the details to the example in Section 5.2.

As mentioned earlier, `tensorflow` already has a variety of distributions implemented. These are located in two places: `tf$distributions` contains some common univariate distributions, while `tf$contrib$distributions` contains common multivariate distributions and more univariate distributions. Specifying the `logLik` and `logPrior` functions will require you to regularly specify distributions. Therefore, in the interest of cleanliness, we make the minor change of copying the distributions from `tf$contrib$distributions` to `tf$distributions`. The module `tf$contrib$distributions` will still work as expected. The copying across occurs as soon as the package is loaded into the workspace. This means any distribution `distn` in the `tensorflow` package can be accessed by referencing `tf$distributions$distn`.

For the rest of this section we go into more detail about the usage of the functions using a worked example. The package is used to infer the bias and coefficient parameters in a logistic regression model. Section 5.1 demonstrates standard usage by performing inference on the model using the `sgld` and `sgldcv` functions. Section 5.2 demonstrates usage in problems where the full MCMC chain cannot be fit into memory. The same logistic regression model is used throughout.

5.1. Example usage

In this example we use the functions `sgld` and `sgldcv` to infer the bias (or intercept) and coefficients of a logistic regression model. The `sgldcv` case is also available as a vignette. Suppose we have data $\mathbf{x}_1, \dots, \mathbf{x}_N$ of dimension d taking values in \mathbb{R}^d ; and response variables

y_1, \dots, y_N taking values in $\{0, 1\}$. Then a logistic regression model with coefficients β and bias β_0 will have likelihood

$$p(\mathbf{X}, \mathbf{y} | \beta, \beta_0) = \prod_{i=1}^N \left[\frac{1}{1 + e^{-\beta_0 - \mathbf{x}_i^\top \beta}} \right]^{y_i} \left[1 - \frac{1}{1 + e^{-\beta_0 - \mathbf{x}_i^\top \beta}} \right]^{1-y_i}. \quad (9)$$

We will use the `covertime` data set (Blackard and Dean 1999) which can be downloaded and loaded using the `sgmcmc` function `getDataset`, which downloads example data sets for the package. All datasets are also available in the supplementary material. The `covertime` data set uses mapping data to predict the type of forest cover. Our particular data set is taken from the `LIBSVM` library (Chang and Lin 2011), which converts the data to a binary problem, rather than multiclass. The data set consists of a matrix of dimension 581012×55 . The first column contains the labels \mathbf{y} , taking values in $\{0, 1\}$. The remaining columns are the explanatory variables \mathbf{X} , which have been scaled to take values in $[0, 1]$.

```
R> library("sgmcmc")
R> load("Data/covertime.rda")
```

Now we will split the data set into predictors and response and get the data set into the required R list format.

```
R> X <- covertype[, 2:ncol(covertype)]
R> y <- covertype[, 1]
R> dataset <- list("X" = X, "y" = y)
```

In the last line, we defined the data set as a list object which will be input to the relevant `sgmcmc` function. The list names can be arbitrary, but must be consistent with the variables declared in the `logLik` function (see below).

When accessing the data, it is assumed that observations are split along the first axis. In other words, `dataset$X` is a 2-dimensional matrix, and we assume that observation \mathbf{x}_i is accessed at `dataset$X[i,]`. Similarly, suppose `Z` was a 3-dimensional array, we would assume that observation i would be accessed at `Z[i, ,]`. Parameters are declared in a similar way, except the list entries are the desired parameter starting points. There are two parameters for this model, the bias β_0 and the coefficients β , which can be arbitrarily initialized to 0.

```
R> d <- ncol(dataset$X)
R> params <- list("bias" = 0, "beta" = matrix(rep(0, d), nrow = d))
```

The log-likelihood is specified as a function of the `dataset` and `params`, which are lists with the same names we declared earlier. The only difference is that the objects inside the lists will have automatically been converted to `tensorflow` objects. The `dataset` list will contain `tensorflow` placeholders. The `params` list will contain `tensorflow` variables. The `logLik` function should be a function that takes these lists as input and returns the log-likelihood value given the current parameters and data. This is done using `tensorflow` operations, as this allows the gradient to be automatically calculated.

For the logistic regression model (9), the log-likelihood is

$$\log p(\mathbf{X}, \mathbf{y} | \beta, \beta_0) = \sum_{i=1}^N y_i \log y_{\text{est}} + (1 - y_i) \log(1 - y_{\text{est}}),$$

where $y_{\text{est}} = [1 + e^{-\beta_0 - \mathbf{x}_i^\top \beta}]^{-1}$, which coded as a `logLik` function is defined as follows

```
R> logLik <- function(params, dataset) {
+   yEst <- 1 / (1 + tf$exp(- tf$squeeze(params$bias +
+     tf$matmul(dataset$X, params$beta))))
+   logLik <- tf$reduce_sum(dataset$y * tf$log(yEst) +
+     (1 - dataset$y) * tf$log(1 - yEst))
+   return(logLik)
+ }
```

For more information about the usage of these **tensorflow** functions, please see the **tensorflow** documentation⁵.

Next, the log-prior density, where we assume each β_j , for $j = 0, \dots, d$, has an independent Laplace prior with location 0 and scale 1, so $\log p(\beta) \propto -\sum_{j=0}^d |\beta_j|$. Similar to `logLik`, this is defined as a function, but with only `params` as input

```
R> logPrior <- function(params) {
+   logPrior <- - (tf$reduce_sum(tf$abs(params$beta)) +
+     tf$reduce_sum(tf$abs(params$bias)))
+   return(logPrior)
+ }
```

The final input that needs to be set is the `stepsize` for tuning the methods, this can be set as a list

```
R> stepsize <- list("beta" = 2e-5, "bias" = 2e-5)
```

Setting the same step size for all parameters is done as `stepsize = 2e-5`. This shorthand can also be used for any of the optional tuning parameters which need to be specified as lists. The step size parameter will generally require a bit of tuning in order to get good performance, for this we recommend using cross validation (see, e.g., [Friedman, Hastie, and Tibshirani 2001](#), Chapter 7).

Everything is now ready to run a standard SGLD algorithm, with `minibatchSize` set to 500. To keep things reproducible we will set the seed to 13.

```
R> output <- sgld(logLik, dataset, params, stepsize, logPrior = logPrior,
+   minibatchSize = 500, nIters = 10000, seed = 13)
```

The output of the function is also a list with the same names as the `params` list. Suppose a given parameter has shape (d_1, \dots, d_l) , then the output will be an array of shape $(\text{nIters}, d_1, \dots, d_l)$. So in our case, `output$beta[i, ,]` is the i th MCMC sample from the parameter β ; and `dim(output$beta)` is `c(10000, 54, 1)`.

In order to run a control variate algorithm such as `sgldcv`, we need one additional argument, which is the step size for the initial optimization step. The optimization uses the `TensorFlow GradientDescentOptimizer`. The step size should be quite similar to the step size for SGLD, though is often slightly larger. First, so that we can measure the performance of the chain, we shall set a seed and randomly remove some observations from the full `dataset` to form a `testset`. We also remove a short burn-in period of 1000 from the final output.

⁵Documentation for **tensorflow** for R is available at <https://tensorflow.rstudio.com/tensorflow/>.

```
R> set.seed(13)
R> testInd <- sample(nrow(dataset$X), 10^4)
R> testset <- list("X" = dataset$X[testInd, ], "y" = dataset$y[testInd])
R> data set <- list("X" = dataset$X[-testInd, ], "y" = dataset$y[-testInd])
R> output <- sgldev(logLik, dataset, params, 5e-6, 5e-6,
+   logPrior = logPrior, minibatchSize = 500, nIters = 11000, seed = 13)
R> output$beta <- output$beta[-c(1:1000), , ]
R> output$bias <- output$bias[-c(1:1000)]
```

A common performance measure for a classifier is the log loss. Given an observation with data \mathbf{x}_i and response y_i , logistic regression predicts that $y_i = 1$ with probability

$$\pi(\mathbf{x}_i, \beta, \beta_0) = \frac{1}{1 + e^{-\beta_0 - \mathbf{x}_i^T \beta}}.$$

Given a test set T of data response pairs (\mathbf{x}, y) , the log loss $s(\cdot)$, of a binary chain, is defined as

$$s(\beta, \beta_0, T) = -\frac{1}{|T|} \sum_{(\mathbf{x}, y) \in T} [y \log \pi(\mathbf{x}, \beta, \beta_0) + (1 - y) \log(1 - \pi(\mathbf{x}, \beta, \beta_0))]. \quad (10)$$

To check convergence of `sgldev` we will plot the log loss every 10 iterations, using the `testset` we removed earlier. Results are given in Figure 2. The plot shows the `sgldev` algorithm converging to stationarity after a short burn-in period. The burn-in period is short due to the initial optimization step. The following chunk of code will calculate the log loss of the output every 10 iterations, and then plot the result.

```
R> iterations <- seq(from = 1, to = 10^4, by = 10)
R> logLoss <- rep(0, length(iterations))
R> for (iter in 1:length(iterations)) {
+   j <- iterations[iter]
+   beta0_j <- output$bias[j]
+   beta_j <- output$beta[j, ]
+   for (i in 1:length(testset$y)) {
+     piCurr <- 1 / (1 + exp(- beta0_j - sum(testset$X[i, ] * beta_j)))
+     y_i <- testset$y[i]
+     logLossCurr <- - (y_i * log(piCurr) + (1 - y_i) * log(1 - piCurr))
+     logLoss[iter] <- logLoss[iter] + 1 / length(testset$y) * logLossCurr
+   }
+ }
R> plotFrame <- data.frame("Iteration" = iterations, "logLoss" = logLoss)
R> ggplot(plotFrame, aes(x = Iteration, y = logLoss)) +
+   geom_line(color = "maroon") + ylab("Log loss of test set")
```

5.2. Example usage: Storage constraints

Often large data sets and high dimensionality go hand in hand. Sometimes the dimensionality is so high that storage of the full MCMC chain in memory becomes an issue. There are a number of ways around this, including: calculating estimates of the desired posterior quantity




Figure 2: Log loss on a test set for parameters simulated using the `sgldcv` algorithm after 1000 iterations of burn-in. Logistic regression problem with the `covertype` data set.

on the fly; reducing the dimensionality of the chain using a test function; or just periodically saving the chain to the hard disk and starting from scratch. To deal with high storage costs, `sgmcmc` provides functionality to run SGMCMC algorithms step by step. This allows users to deal with the output as they see fit at each iteration.

In this section, we detail how to run SGMCMC chains step by step. To do this requires more knowledge of TensorFlow, including using TensorFlow sessions and building custom placeholders and tensors. For more details see the `tensorflow` for R documentation (Allaire and Tang 2019). The step by step procedure works similarly to a standard TensorFlow procedure: `tensorflow` variables, tensors and placeholders are declared; then the TensorFlow session is started and all the required tensors are initialized; finally the SGMCMC algorithm is run step by step in a user defined loop, and the user evaluates tensors as required.

To demonstrate this concept we keep things simple and use the logistic regression example introduced in the previous section. While this example can fit into memory, it allows us to demonstrate the concepts without getting bogged down in a complicated model. For a more realistic example, where the output does not fit into memory, see the Bayesian neural network model in Section 6.3.

We start by assuming the `dataset`, `params`, `logLik`, `logPrior` and `stepsize` objects have been created in exactly the same way as in the previous example (Section 5.1). Now suppose we want to make inference using stochastic gradient Langevin dynamics (SGLD), but we want to run it step by step. The first step is to initialize an ‘`sgld`’ object using the function `sgldSetup`. For every function in Table 1 there is a corresponding `*Setup` function, such as `sghmccvSetup` or `sgnhtSetup`. This function will create all the `tensorflow` objects required, as well as declare the dynamics of the SGMCMC algorithm. For our example we can run the following

```
R> sgld <- sgldSetup(logLik, dataset, params, stepsize, logPrior = logPrior,
+   minibatchSize = 500, seed = 13)
```

This ‘`sgld`’ object is a type of ‘`sgmcmc`’ object, it is an R S3 object, which is essentially a list with a number of entries. The most important of these entries for building SGMCMC algorithms is called `params`, which holds a list, with the same names as in the `params` that were fed to `sgldSetup`, but this list contains `tf$Variable` objects. This is how the tensors

are accessed which hold the current parameter values in the chain. For more details on the attributes of these objects, see the documentation for `sgldSetup`, `sgldcvSetup`, etc.

Now that we have created the ‘`sgld`’ object, we want to initialize the **tensorflow** variables and the SGMCMC algorithm chosen. For a standard algorithm, this will initialize the **tensorflow** graph and all the tensors that were created. For an algorithm with control variates (e.g., `sgldcv`), this will also find the $\hat{\theta}$ estimates of the parameters and calculate the full log posterior gradient at that point; as detailed in Section 3.4. The function used to do this is `initSess`,

```
R> sess <- initSess(sgld)
```

The `sess` returned by `initSess` is the current TensorFlow session, which is needed to run the SGMCMC algorithm of choice, and to access any of the tensors needed, such as `sgld$params`.

Now we have everything to run an SGLD algorithm step by step as follows

```
R> for (i in 1:10^4) {
+   sgmcmcStep(sgld, sess)
+   currentState <- getParams(sgld, sess)
+ }
```

Here the function `sgmcmcStep` will update `sgld$params` using a single update of SGLD, or whichever SGMCMC algorithm is given. The function `getParams` will return a list of the current parameters as R objects rather than as tensors.

This simple example of running SGLD step by step only stores the most recent value in the chain, which is useless for a Monte Carlo method. Also, for large scale examples, it is often useful to reduce the dimension of the chain by calculating some test function $g(\cdot)$ of θ at each iteration rather than the parameters themselves. This example will demonstrate how to store a test function at each iteration, as well as calculating the estimated posterior mean on the fly. We assume that a new R session has been started and the ‘`sgld`’ object has just been created using `sgldSetup` with the same properties as in the example above. We assume that no TensorFlow session has been created (i.e., `initSess` has not been run yet).

Before the TensorFlow session has been declared, the user is able to create their own custom tensors. This is useful, as test functions can be declared beforehand using the `sgld$params` variables, which allows the test functions to be quickly calculated by just evaluating the tensors in the current session. The test function used here is once again the log loss of a test set, as defined in (10).

Suppose we input `sgld$params` and the testset T to the `logLik` function. Then the log loss is actually $-\frac{1}{|T|}$ times this value. This means we can easily create a tensor that calculates the log loss by creating a list of placeholders that hold the test set, then using the `logLik` function with the testset list and `sgld$params` as input. We can do this as follows

```
R> testPlaceholder <- list()
R> testPlaceholder[["X"]] <- tf$placeholder(tf$float32, dim(testset[["X"]]))
R> testPlaceholder[["y"]] <- tf$placeholder(tf$float32, dim(testset[["y"]]))
R> testSize <- as.double(nrow(testset[["X"]]))
R> logLoss <- - logLik(sgld$params, testPlaceholder) / testSize
```

This placeholder is then fed the full `testset` each time the log loss is calculated. Now we will declare the TensorFlow session, and run the chain step by step. At each iteration, we will calculate the current Monte Carlo estimate of the parameters. The log loss will be stored every 100 iterations. We omit a plot of the log loss as it is similar to Figure 2. This first chunk of code corresponds to the burn-in, at this point we do not store parameter estimates, we just print the log loss every 100 iterations to check convergence of the chain.

```
R> sess <- initSess(sgld)
R> feedDict <- dict()
R> feedDict[[testPlaceholder[["X"]]]] <- testset[["X"]]
R> feedDict[[testPlaceholder[["y"]]]] <- testset[["y"]]
R> message("Burning-in chain...")
R> message("iteration\tlog loss")
R> for (i in 1:10^4) {
+   if (i %% 100 == 0) {
+     progress <- sess$run(logLoss, feed_dict = feedDict)
+     message(paste0(i, "\t", progress))
+   }
+   sgmcmcStep(sgld, sess)
+ }
```

In the previous chunk of code, the purpose of `feedDict` is to feed the whole `testset` to the `logLoss` tensor. In the next chunk of code we run the chain for 10^4 iterations. At each iteration we store a running average of the parameter estimates in the R list object `postMean`. Again we print the log loss of the chain every 100 iterations.

```
R> postMean <- getParams(sgld, sess)
R> logLossOut <- rep(0, 10^4 / 100)
R> message("Running SGMCMC...")
R> for (i in 1:10^4) {
+   sgmcmcStep(sgld, sess)
+   currentState <- getParams(sgld, sess)
+   for (paramName in names(postMean)) {
+     postMean[[paramName]] <- (postMean[[paramName]] * i +
+       currentState[[paramName]]) / (i + 1)
+   }
+   if (i %% 100 == 0) {
+     logLossOut[i/100] <- sess$run(logLoss, feed_dict = feedDict)
+     message(paste0(i, "\t", logLossOut[i/100]))
+   }
+ }
```

6. Simulations

In this section, we demonstrate the algorithms and performance by simulating from a variety of models using all the implemented methods and commenting on the performance of

each. These simulations are reproducible and available in the supplementary material and on Github⁶. For more usage tutorials similar to Sections 5.1 and 5.2, please see the vignettes on the package website⁷.

6.1. Gaussian mixture

In this model we assume our data set x_1, \dots, x_N is drawn i.i.d. from

$$X_i | \theta_1, \theta_2 \sim \frac{1}{2} \mathcal{N}(\theta_1, \mathbf{I}_2) + \frac{1}{2} \mathcal{N}(\theta_2, \mathbf{I}_2), \quad i = 1, \dots, N; \quad (11)$$

where θ_1, θ_2 are parameters to be inferred and \mathbf{I}_2 is the 2×2 identity matrix. We assume the prior $\theta_1, \theta_2 \sim \mathcal{N}(0, 10\mathbf{I}_2)$. To generate the synthetic data set, we simulate 10^3 observations from $\frac{1}{2} \mathcal{N}([0, 0]^\top, \mathbf{I}_2) + \frac{1}{2} \mathcal{N}([0.1, 0.1]^\top, \mathbf{I}_2)$. While this is a small number of observations, it allows us to compare the results to a full HMC scheme using the R implementation of Stan (Stan Development Team 2018; Carpenter *et al.* 2017). The full HMC scheme should sample from close to the true posterior distribution, so acts as a good surrogate for the truth. We compare each SGMCMC algorithm implemented to the HMC sample to compare performance. Larger scale examples are given in Sections 6.2 and 6.3. We ran all methods for 10^4 iterations, except SGHMC, since the computational cost is greater for this method due to the trajectory parameter L . We ran SGHMC for 2,000 iterations, using default trajectory $L = 5$, as this ensures the overall computational cost of the method is similar to the other methods. We used a burn-in step of 10^4 iterations, except for the control variate methods, where we used 10^4 iterations in the initial optimization step, with no burn-in. Again this ensures comparable computational cost across different methods.

The `logLik` and `logPrior` functions used for this model are as follows. Remember that any distribution can be referenced using `tf$distributions` once the package has been loaded in.

```
R> logLik <- function(params, dataset) {
+   SigmaDiag <- c(1, 1)
+   component1 <- tf$distributions$MultivariateNormalDiag(
+     params$theta1, SigmaDiag)
+   component2 <- tf$distributions$MultivariateNormalDiag(
+     params$theta2, SigmaDiag)
+   probs <- tf$distributions$Categorical(c(0.5, 0.5))
+   distn <- tf$distributions$Mixture(probs,
+     list(component1, component2))
+   logLik <- tf$reduce_sum(distn$log_prob(dataset$X))
+   return(logLik)
+ }
R> logPrior <- function(params) {
+   mu0 <- c(0, 0)
+   Sigma0Diag <- c(10, 10)
+   priorDistn <- tf$distributions$MultivariateNormalDiag(mu0,
+     Sigma0Diag)
```

⁶<https://github.com/jbaker92/sgmcmc-simulations>

⁷<https://stor-i.github.io/sgmcmc>




Figure 3: Plots of the approximate posterior for θ_1 simulated using each of the methods implemented by **sgmcmc**, compared with a full HMC run, treated as the truth, for the Gaussian mixture model (11).

```
+   logPrior <- priorDistn$log_prob(params$theta1) +
+     priorDistn$log_prob(params$theta2)
+   return(logPrior)
+ }
```

The following list determines the step sizes used for each method, the `optstepsize` used for control variate methods was `5e-5`.

```
R> stepsizeList <- list("sgld" = 5e-3, "sghmc" = 5e-4, "sgnht" = 3e-4,
+   "sgldcv" = 1e-2, "sghmccv" = 1.5e-3, "sgnhtcv" = 3e-3)
```

We set the seed to be 2 using the optional `seed` argument and use a mini batch size of 100. We also used a seed of 2 when generating the data (see the supplementary material for full details). Starting points were sampled from a standard normal. The results are plotted in Figure 3. The black contours represent the best guess at the true posterior, which was found using the standard HMC procedure in Stan. The colored contours that overlay the black contours are the approximations of each of the SGMCMC methods implemented by **sgmcmc**. This allows us to compare the SGMCMC estimates with the “truth” by eye.

In the simulation, we obtain two chains, one approximating θ_1 and the other approximating θ_2 . In order to examine how well the methods explore both modes, we take just θ_1 and compare this to the HMC run for θ_1 . The results are quite variable, and it demonstrates a point nicely: there seems to be a trade-off between predictive accuracy and exploration. Many methods have demonstrated good performance using predictive accuracy; where a test set is removed from the full data set to assess how well the fitted model performs on the test set. This is a useful technique for complex models, which are high dimensional and have a large number of data points, as they cannot be plotted, and an MCMC run to act as the “truth” cannot be fitted.

However, this example shows that it does not give the full picture. A lot of the methods which show improved predictive performance (e.g., control variate methods and especially **sgnht**) over **sgld** appear here to perform worse at exploring the full space. In this example, **sgld** performs the best at exploring both components, though it over-estimates posterior uncertainty. The algorithm **sghmc** also explores both components but somewhat unevenly. We find that **sgnht**, while being shown to have better predictive performance in the original paper (Ding *et al.* 2014), does not do nearly as well as the other algorithms at exploring the

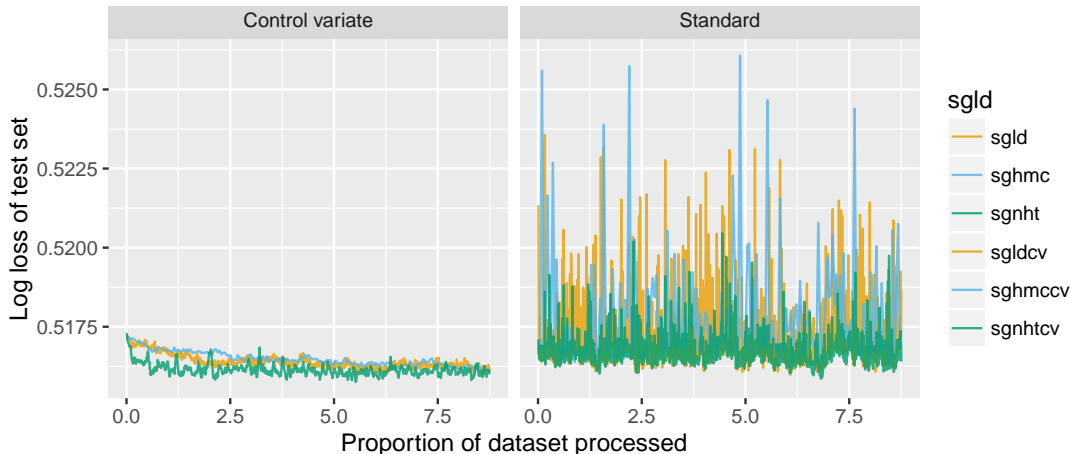


Figure 4: Plots of the log loss of a test set for β_0 and β simulated using each of the methods implemented by **sgmcmc**. Logistic regression problem with the covertype data set.

space and appears to collapse to the posterior mode. The control variate methods, shown in the following sections, and in [Baker et al. \(2017\)](#), appear to have better predictive performance than **sgld**, but do not explore both components either. For example, **sgldcv** explores the space the best but over-estimates uncertainty of the first component, since it relies on SGLD updates which also overestimates uncertainty. In contrast, **sgnhtcv** collapses to a posterior mode since it relies on the SGNHT updates which also collapse.

6.2. Bayesian logistic regression

In this section, we apply all the methods to the logistic regression example in Section 5.1. We compare the performance of the methods by calculating the log loss of a test set every 10 iterations, again as detailed in Section 5.1. The standard methods (**sgld**, **sghmc**, **sgnht**) were run for 10^4 iterations with an additional 10^4 iterations of burn-in; except for **sghmc** which has $5\times$ the computational cost, so is ran for 2,000 iterations with 2,000 iterations of burn-in. The control variate methods were run for 10^4 iterations with an additional 10^4 iterations for the initial optimization step, and no burn-in; again except for **sghmccv** which was run for 2,000 iterations. This means that all the methods should be somewhat comparable in terms of computation time. The following list determines the step sizes used for each method, the `optstepsize` used was $1e-6$.

```
R> stepsizes <- list("sgld" = 5e-6, "sghmc" = 1e-7, "sgnht" = 1e-7,
+   "sgldcv" = 1e-5, "sghmccv" = 1e-6, "sgnhtcv" = 5e-7)
```

We set the seed to be 1 for each of the simulations, and when generating the test data (see the supplementary material for reproducible code) and use a mini batch size of 500. Starting points are sampled from a standard normal.

Results are plotted in Figure 4. All of the algorithms show decent performance. Methods which use control variates have significantly better predictive performance; and result in chains with lower variance. **sghmc** has lower variance than **sgld** and **sgnht**, though this could be related to the high computational cost. One might envisage setting a lower trajectory L

would result in a chain with higher variance. `sg1dcv` takes longer to burn-in than the other control variate methods. The algorithm `sgld` has the highest variance by far; this could be related to our discussion in Section 6.1 on exploration versus accuracy.

6.3. Bayesian neural network

In this simulation, we demonstrate a very high dimensional chain. This gives a more realistic example of when we would want to run the chain step by step. The model is a two layer Bayesian neural network which is fit to the MNIST data set (LeCun and Cortes 2010). More step by step detail of how to fit this model can be found in the `sgmcmc` vignettes (Baker *et al.* 2019) or on the package website⁸. The MNIST data set consists of 28×28 pixel images of handwritten digits from zero to nine. The images are flattened to be a vector of length 784. The data set is available as a standard data set from the TensorFlow library, with a matrix of 55,000 training vectors and 10,000 test vectors, each with their corresponding labels. The data set can be constructed in a similar way to the logistic regression example of Section 5.1, using the standard data set in the package `mnist`.

```
R> library("sgmcmc")
R> load("Data/mnist.rda")
R> dataset <- list("X" = mnist$train$images, "y" = mnist$train$labels)
R> testset <- list("X" = mnist$test$images, "y" = mnist$test$labels)
```

We build the same neural network model as in the original SGHMC paper by Chen *et al.* (2014). Suppose Y_i takes values in $\{0, \dots, 9\}$, so is the output label of a digit, and \mathbf{x}_i is the input vector, with \mathbf{X} the full $N \times 784$ data set, where N is the number of observations. The model is then as follows

$$Y_i | \theta, \mathbf{x}_i \sim \text{Categorical}(\beta(\theta, \mathbf{x}_i)), \quad (12)$$

$$\beta(\theta, \mathbf{x}_i) = \sigma \left(\sigma \left(\mathbf{x}_i^\top B + b \right) A + a \right). \quad (13)$$

Here A, B, a, b are parameters to be inferred with $\theta = (A, B, a, b)$; $\sigma(\cdot)$ is the softmax function (a generalization of the logistic link function). A, B, a and b are matrices with dimensions: 100×10 , 784×100 , 1×10 and 1×100 respectively. Each element of these parameters is assigned a normal prior

$$\begin{aligned} A_{kl} | \lambda_A &\sim \mathcal{N}(0, \lambda_A^{-1}), & B_{jk} | \lambda_B &\sim \mathcal{N}(0, \lambda_B^{-1}), \\ a_l | \lambda_a &\sim \mathcal{N}(0, \lambda_a^{-1}), & b_k | \lambda_b &\sim \mathcal{N}(0, \lambda_b^{-1}), \\ j = 1, \dots, 784; & & k = 1, \dots, 100; & & l = 1, \dots, 10; \end{aligned}$$

where $\lambda_A, \lambda_B, \lambda_a$ and λ_b are hyperparameters. Finally, we assume

$$\lambda_A, \lambda_B, \lambda_a, \lambda_b \sim \text{Gamma}(1, 1).$$

The model contains a large number of high dimensional parameters, and unless there is sufficient RAM available, a standard chain of length 10^4 will not fit into memory. First, we shall create the `params` dictionary, and then code the `logLik` and `logPrior` functions. We can sample the initial λ parameters from a standard Gamma distribution, and the remaining parameters from a standard normal as follows

⁸<https://stor-i.github.io/sgmcmc/articles/nn.html>

```

R> d <- ncol(dataset$X)
R> params <- list()
R> params$A <- matrix(rnorm(10*100), ncol = 10)
R> params$B <- matrix(rnorm(d*100), ncol = 100)
R> params$a <- rnorm(10)
R> params$b <- rnorm(100)
R> params$lambdaA <- rgamma(1, 1)
R> params$lambdaB <- rgamma(1, 1)
R> params$lambdaa <- rgamma(1, 1)
R> params$lambdab <- rgamma(1, 1)

```

Next, we can declare the `logLik` function, as based on (13), using standard TensorFlow methods.

```

R> logLik <- function(params, dataset) {
+   beta <- tf$nn$softmax(tf$matmul(dataset$X, params$B) + params$b)
+   beta <- tf$nn$softmax(tf$matmul(beta, params$A) + params$a)
+   logLik <- tf$reduce_sum(dataset$y * tf$log(beta))
+   return(logLik)
+ }

```

The `logPrior` function can be declared as follows. Remember that any distribution can be referenced using `tf$distributions` once the package has been loaded in.

```

R> logPrior <- function(params) {
+   distLambda <- tf$distributions$Gamma(1, 1)
+   distA <- tf$distributions$Normal(0, tf$rsqrt(params$lambdaA))
+   logPriorA <- tf$reduce_sum(distA$log_prob(params$A)) +
+     distLambda$log_prob(params$lambdaA)
+   distB <- tf$distributions$Normal(0, tf$rsqrt(params$lambdaB))
+   logPriorB <- tf$reduce_sum(distB$log_prob(params$B)) +
+     distLambda$log_prob(params$lambdaB)
+   dista <- tf$distributions$Normal(0, tf$rsqrt(params$lambdaa))
+   logPriora <- tf$reduce_sum(dista$log_prob(params$a)) +
+     distLambda$log_prob(params$lambdaa)
+   distb <- tf$distributions$Normal(0, tf$rsqrt(params$lambdab))
+   logPriorb <- tf$reduce_sum(distb$log_prob(params$b)) +
+     distLambda$log_prob(params$lambdab)
+   logPrior <- logPriorA + logPriorB + logPriora + logPriorb
+   return(logPrior)
+ }

```

Similar to Section 5.2, we use the log loss as a test function. This time though it is necessary to update the definition, as the logistic regression example was a binary problem whereas now we have a multiclass problem. Given a test set T of pairs (\mathbf{x}, y) , now y can take values in $\{0, \dots, K\}$, rather than just binary values. To account for this we redefine the definition of log loss to be

$$s(\theta, T) = -\frac{1}{|T|} \sum_{\mathbf{x}, y \in T} \sum_{k=1}^K \mathbf{1}_{y=k} \log \beta_k(\theta, \mathbf{x}),$$




Figure 5: Plots of the log loss of a held out test set based on simulations using each of the methods implemented by **sgmcmc**. Bayesian neural network problem with MNIST data set.

where $\mathbf{1}_A$ is the indicator function, and $\beta_k(\theta, \mathbf{x})$ is the k th element of $\beta(\theta, \mathbf{x})$ as defined in (13).

As in Section 5.2, the log loss is simply $-\frac{1}{|T|}$ times the `logLik` function, if we feed it the `testset` rather than the `dataset`. This means the `logLoss` tensor can be declared in a similar way to Section 5.2.

```
R> testPlaceholder <- list()
R> testPlaceholder[["X"]] <- tf$placeholder(tf$float32, dim(testset[["X"]]))
R> testPlaceholder[["y"]] <- tf$placeholder(tf$float32, dim(testset[["y"]]))
R> testSize <- as.double(nrow(testset[["X"]]))
R> logLoss <- -logLik(sgld$params, testPlaceholder) / testSize
```

We can run the chain in exactly the same way as Section 5.2, and so omit the code for this. We ran 10^4 iterations of each of the algorithms in Table 1, calculating the log loss for each algorithm every 10 iterations. The standard algorithms have 10^4 iterations of burn-in while the control variate algorithms have no burn-in, but 10^4 iterations in the initial optimization step. Note that due to the trajectory parameter L , `sghmc` and `sghmccv` will have 5 times the computational cost of the other algorithms. Therefore, we ran these algorithms for 2,000 iterations instead, to make the computational cost comparable. We used the following list of step sizes

```
R> list("sgld" = 1e-4, "sghmc" = 1e-5, "sgnht" = 5e-6, "sgldcv" = 5e-5,
+      "sghmccv" = 1e-5, "sgnhtcv" = 5e-7)
```

Generally these are the step sizes which produce the smallest log loss; except when these chains did not seem to explore the space fully, in which case we increased the step size slightly. We set the seed to be 1 for each of the simulations, and when generating the test data (see the supplementary material for reproducible code).

The results are plotted in Figure 5. Again we see improvements in the predictive performance of the control variate methods. Among the standard methods, `sghmc` and `sgnht` have the best predictive performance; which is to be expected given the apparent trade-off between accuracy and exploration.

7. Discussion

We presented the R package **sgmcmc**, which enables Bayesian inference with large data sets using stochastic gradient Markov chain Monte Carlo. The package only requires the user to specify the log-likelihood and log-prior functions; and any differentiation required can be performed automatically. The package is based on TensorFlow (Abadi *et al.* 2015), an efficient library for numerical computation that can take advantage of many different architectures, including GPUs. The **sgmcmc** package keeps much of this efficiency. The package provides functionality to deal with cases where the full MCMC chain is too large to fit into memory. As the chain can be run step by step at each iteration, there is flexibility for these cases.

We implemented the methods on a variety of statistical models, many on realistic data sets. One of these statistical models was a neural network, for which the full MCMC chain would not fit into memory. In this case, we demonstrated building test functions and calculating the Monte Carlo estimates on the fly. We empirically demonstrated the predictive performance of the algorithms and the trade-off that appears to occur between predictive performance and exploration.

Many complex models for which SGMCMC methods have been found to perform well require Gibbs updates to be performed periodically (Patterson and Teh 2013; Li, Ahn, and Welling 2016). In the future, we would like to build functionality for user defined Gibbs steps that can be updated step by step alongside the SGMCMC algorithms. SGHMC has been implemented by setting the value $\hat{\beta}_t = 0$, as in the experiments of the original paper (Chen *et al.* 2014). In the future, we would like to implement a more sophisticated approach to set this value, such as using a similar estimate to Ahn *et al.* (2012).

Acknowledgments

The first author gratefully acknowledges the support of the EPSRC funded EP/L015692/1 STOR-i Centre for Doctoral Training.

This work was supported by EPSRC grant EP/K014463/1, ONR Grant N00014-15-1-2380 and NSF CAREER Award IIS-1350133.

References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, Zheng X (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. URL <https://www.tensorflow.org/>.
- Abraham R, Marsden JE, Ratiu R (1988). *Manifolds, Tensor Analysis, and Applications*, volume 2. Springer-Verlag.
- Ahn S, Korattikara A, Welling M (2012). “Bayesian Posterior Sampling via Stochastic Gradient Fisher Scoring.” In *Proceedings of the 29th International Conference on Machine Learning*, pp. 1591–1598.

- Allaire JJ, Tang Y (2019). *tensorflow: R Interface to TensorFlow*. R package version 1.13.1, URL <https://CRAN.R-project.org/package=tensorflow>.
- Baker J, Fearnhead P, Fox EB, Nemeth C (2017). “Control Variates for Stochastic Gradient MCMC.” arXiv:1706.05439 [stat.CO], URL <https://arxiv.org/abs/1706.05439>.
- Baker J, Nemeth C, Fearnhead P, Fox EB (2019). *sgmcmc: Stochastic Gradient Markov Chain Monte Carlo*. R package version 0.2.5, URL <https://CRAN.R-project.org/package=sgmcmc>.
- Blackard JA, Dean DJ (1999). “Comparative Accuracies of Artificial Neural Networks and Discriminant Analysis in Predicting Forest Cover Types from Cartographic Variables.” *Computers and Electronics in Agriculture*, **24**(3), 131–151. doi:10.1016/s0168-1699(99)00046-0.
- Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A (2017). “Stan: A Probabilistic Programming Language.” *Journal of Statistical Software*, **76**(1), 1–32. doi:10.18637/jss.v076.i01.
- Chang CC, Lin CJ (2011). “LIBSVM: A Library for Support Vector Machines.” *ACM Transactions on Intelligent Systems and Technology*, **2**(3), 1–27. doi:10.1145/1961189.1961199.
- Chen T, Fox EB, Guestrin C (2014). “Stochastic Gradient Hamiltonian Monte Carlo.” In *Proceedings of the 31st International Conference on Machine Learning*, pp. 1683–1691.
- Ding N, Fang Y, Babbush R, Chen C, Skeel RD, Neven H (2014). “Bayesian Sampling Using Stochastic Gradient Thermostats.” In *Advances in Neural Information Processing Systems 27*, pp. 3203–3211.
- Dubey KA, Reddi SJ, Williamson SA, Póczos B, Smola AJ, Xing EP (2016). “Variance Reduction in Stochastic Gradient Langevin Dynamics.” In *Advances in Neural Information Processing Systems 29*, pp. 1154–1162.
- Friedman J, Hastie T, Tibshirani R (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag. doi:10.1007/978-0-387-21606-5.
- Griewank A, Walther A (2008). *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, volume 2. SIAM.
- LeCun Y, Cortes C (2010). *MNIST Handwritten Digit Database*. URL <http://yann.lecun.com/exdb/mnist/>.
- Li W, Ahn S, Welling M (2016). “Scalable MCMC for Mixed Membership Stochastic Block-models.” In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pp. 723–731.
- Lunn D, Spiegelhalter D, Thomas A, Best N (2009). “The BUGS Project: Evolution, Critique and Future Directions.” *Statistics in Medicine*, **28**(25), 3049–3067. doi:10.1002/sim.3680.
- Lunn DJ, Thomas A, Best N, Spiegelhalter D (2000). “WinBUGS – A Bayesian Modelling Framework: Concepts, Structure, and Extensibility.” *Statistics and Computing*, **10**(4), 325–337. doi:10.1023/a:1008929526011.

- Nagapetyan T, Duncan A, Hasenclever L, Vollmer SJ, Szpruch L, Zygalkakis K (2017). “The True Cost of Stochastic Gradient Langevin Dynamics.” arXiv:1706.02692 [stat.ME], URL <https://arxiv.org/abs/1706.02692>.
- Neal RM (2010). “MCMC Using Hamiltonian Dynamics.” In *Handbook of Markov Chain Monte Carlo*, volume 1. Chapman & Hall.
- Patterson S, Teh YW (2013). “Stochastic Gradient Riemannian Langevin Dynamics on the Probability Simplex.” In *Advances in Neural Information Processing Systems 26*, pp. 3102–3110.
- Plummer M (2003). “**JAGS**: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling.” In K Hornik, F Leisch, A Zeileis (eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*. Technische Universität Wien, Vienna, Austria. URL <https://www.R-project.org/conferences/DSC-2003/Proceedings/Plummer.pdf>.
- Plummer M (2018). *rjags: Bayesian Graphical Models using MCMC*. R package version 4-8, URL <https://CRAN.R-project.org/package=rjags>.
- R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Ripley BD (2009). *Stochastic Simulation*, volume 1. John Wiley & Sons.
- Roberts GO, Rosenthal JS (1998). “Optimal Scaling of Discrete Approximations to Langevin Diffusions.” *Journal of the Royal Statistical Society B*, **60**(1), 255–268. doi:10.1111/1467-9868.00123.
- Sato I, Nakagawa H (2014). “Approximation Analysis of Stochastic Gradient Langevin Dynamics by Using Fokker-Planck Equation and Ito Process.” In *Proceedings of the 31st International Conference on Machine Learning*, pp. 982–990.
- Stan Development Team (2018). “**rstan**: R interface to Stan.” R package version 2.18.2, URL <http://mc-stan.org/>.
- Sturtz S, Ligges U, Gelman A (2005). “**R2WinBUGS**: A Package for Running WinBUGS from R.” *Journal of Statistical Software*, **12**(3), 1–16. doi:10.18637/jss.v012.i03.
- Teh YW, Thiéry AH, Vollmer SJ (2016). “Consistency and Fluctuations for Stochastic Gradient Langevin Dynamics.” *Journal of Machine Learning Research*, **17**(7), 1–33.
- Tran D, Kucukelbir A, Dieng AB, Rudolph M, Liang D, Blei DM (2016). “**Edward**: A Library for Probabilistic Modeling, Inference, and Criticism.” arXiv:1610.09787 [stat.CO], URL <https://arxiv.org/abs/1610.09787>.
- Van Rossum G, et al. (2011). *Python Programming Language*. URL <https://www.python.org/>.
- Vollmer SJ, Zygalkakis KC, Teh YW (2016). “Exploration of the (Non-)Asymptotic Bias and Variance of Stochastic Gradient Langevin Dynamics.” *Journal of Machine Learning Research*, **17**(159), 1–48.

Welling M, Teh YW (2011). “Bayesian Learning via Stochastic Gradient Langevin Dynamics.”
In *Proceedings of the 28th International Conference on Machine Learning*, pp. 681–688.

Affiliation:

Jack Baker
STOR-i Centre for Doctoral Training
Department of Mathematics and Statistics
Lancaster University
Lancaster, LA1 4YW, United Kingdom
E-mail: j.baker1@lancaster.ac.uk
URL: <http://lancs.ac.uk/~bakerj1/>

Christopher Nemeth, Paul Fearnhead
Department of Mathematics and Statistics
Lancaster University
Lancaster, LA1 4YW, United Kingdom
E-mail: c.nemeth@lancaster.ac.uk, p.fearnhead@lancaster.ac.uk
URL: <http://lancs.ac.uk/~nemeth/>, <http://www.maths.lancs.ac.uk/~fearnhea/>

Emily Fox
Department of Statistics
University of Washington
Box 354322
Seattle, WA, 98195-4322, United States of America
E-mail: ebfox@uw.edu
URL: <https://homes.cs.washington.edu/~ebfox/>