



LaplacesDemon: An R Package for Bayesian Inference

Byron Hall
STATISTICAT, LLC

Abstract

LaplacesDemon, usually referred to as Laplace’s Demon, is a contributed R package for Bayesian inference, and is freely available on the Comprehensive R Archive Network (CRAN). Laplace’s Demon allows Laplace Approximation and the choice of numerous MCMC algorithms to update a Bayesian model according to a user-specified model function. The user-specified model function enables Bayesian inference for any model form, provided the user specifies, or approximates, the likelihood. Laplace’s Demon also attempts to assist the user by creating and offering R code, based on a previous model update, that can be copy/pasted and executed. Posterior predictive checks and many other features are included as well. Laplace’s Demon seeks to be generalizable and user-friendly to Bayesians, especially Laplacians.

Keywords: ~ Adaptive, AM, Bayesian, Delayed Rejection, DR, DRAM, DRM, Gradient Ascent, Laplace Approximation, LaplacesDemon, Laplace’s Demon, Markov chain Monte Carlo, MCMC, Metropolis, Metropolis-within-Gibbs, Optimization, R, Random Walk, Random-Walk, Resilient Backpropagation, STATISTICAT, t-walk.

Bayesian inference is named after Reverend Thomas Bayes (1702-1761) for developing Bayes’ theorem, which was published posthumously after his death ([Bayes and Price 1763](#)). This was the first instance of what would be called inverse probability¹.

Unaware of Bayes, Pierre-Simon Laplace (1749-1827) independently developed Bayes’ theorem and first published his version in 1774, eleven years after Bayes, in one of Laplace’s first major works ([Laplace 1774](#), p. 366–367). In 1812, Laplace introduced a host of new ideas and mathematical techniques in his book, *Theorie Analytique des Probabilites*, ([Laplace 1812](#)). Before Laplace, probability theory was solely concerned with developing a mathematical analysis of games of chance. Laplace applied probabilistic ideas to many scientific and practical problems. Although Laplace is not the father of probability, Laplace may be considered the father of the field of probability.

¹‘Inverse probability’ refers to assigning a probability distribution to an unobserved variable, and is in essence, probability in the opposite direction of the usual sense. Bayes’ theorem has been referred to as “the principle of inverse probability”. Terminology has changed, and the term ‘Bayesian probability’ has displaced ‘inverse probability’. The adjective “Bayesian” was introduced by R. A. Fisher as a derogatory term.

In 1814, Laplace published his “Essai Philosophique sur les Probabilites”, which introduced a mathematical system of inductive reasoning based on probability (Laplace 1814). In it, the Bayesian interpretation of probability was developed independently by Laplace, much more thoroughly than Bayes, so some “Bayesians” refer to Bayesian inference as Laplacian inference. This is a translation of a quote in the introduction to this work:

“We may regard the present state of the universe as the effect of its past and the cause of its future. An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes” (Laplace 1814).

The ‘intellect’ has been referred to by future biographers as Laplace’s Demon. In this quote, Laplace expresses his philosophical belief in hard determinism and his wish for a computational machine that is capable of estimating the universe.

This article is an introduction to an R (R Development Core Team 2012) package called **LaplacesDemon** (Hall 2012), which was designed without consideration for hard determinism, but instead with a lofty goal toward facilitating high-dimensional Bayesian (or Laplacian) inference², posing as its own intellect that is capable of impressive analysis. The **LaplacesDemon** R package is often referred to as Laplace’s Demon. This article guides the user through installation, data, specifying a model, initial values, updating Laplace’s Demon, summarizing and plotting output, posterior predictive checks, general suggestions, discusses independence and observability, covers details of the algorithm, software comparisons, and discusses large data sets and speed.

Herein, it is assumed that the reader has basic familiarity with Bayesian inference, numerical approximation, and R. If any part of this assumption is violated, then suggested sources include the vignette entitled “Bayesian Inference” that comes with the **LaplacesDemon** package, Gelman, Carlin, Stern, and Rubin (2004), and Crawley (2007).

1. Installation

To obtain Laplace’s Demon, simply open R and install the **LaplacesDemon** package from a CRAN mirror:

```
> install.packages("LaplacesDemon")
```

A goal in developing Laplace’s Demon was to minimize reliance on other packages or software. Therefore, the usual `dep=TRUE` argument does not need to be used, because **LaplacesDemon** does not depend on anything other than base R. Once installed, simply use the `library` or `require` function in R to activate the **LaplacesDemon** package and load its functions into memory:

²Even though the **LaplacesDemon** package is dedicated to Bayesian inference, frequentist inference may be used instead with the same functions by omitting the prior distributions and maximizing the likelihood.

```
> library(LaplacesDemon)
```

2. Data

Laplace’s Demon requires data that is specified in a list. As an example, there is a data set called `demonsnacks` that is provided with the **LaplacesDemon** package. For no good reason, other than to provide an example, the log of `Calories` will be fit as an additive, linear function of the remaining variables. Since an intercept will be included, a vector of 1’s is inserted into design matrix \mathbf{X} .

```
> data(demonsnacks)
> N <- nrow(demonsnacks)
> J <- ncol(demonsnacks)
> y <- log(demonsnacks$Calories)
> X <- cbind(1, as.matrix(demonsnacks[,c(1,3:10)]))
> for (j in 2:J) {X[,j] <- CenterScale(X[,j])}
> mon.names <- c("LP", "sigma")
> parm.names <- as.parm.names(list(beta=rep(0,J), log.sigma=0))
> MyData <- list(J=J, X=X, mon.names=mon.names, parm.names=parm.names, y=y)
```

There are $J=10$ independent variables (including the intercept), one for each column in design matrix \mathbf{X} . However, there are 11 parameters, since the residual variance, σ^2 , must be included as well. The reason why it is called `log.sigma` will be explained later. Each parameter must have a name specified in the vector `parm.names`, and parameter names must be included with the data. This is using a function called `as.parm.names`. Also, note that each predictor has been centered and scaled, as per Gelman (2008). Laplace’s Demon provides a `CenterScale` function to center and scale predictors³.

Laplace’s Demon will consider using Laplace Approximation, and part of this consideration includes determining the sample size. The user must specify the number of observations in the data as either a scalar `n` or `N`. If these are not found by the `LaplaceApproximation` or `LaplacesDemon` functions, then it will attempt to determine sample size as the number of rows in `y` or `Y`.

3. Specifying a Model

Laplace’s Demon is capable of estimating any Bayesian model for which the likelihood is specified⁴. To use Laplace’s Demon, the user must specify a model. Let’s consider a linear regression model, which is often denoted as:

$$\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2)$$

³Centering and scaling a predictor is `x.cs <- (x - mean(x)) / (2*sd(x))`.

⁴Examples of more than 80 Bayesian models may be found in the “Examples” vignette that comes with the **LaplacesDemon** package. Likelihood-free estimation is also possible by approximating the likelihood, such as in Approximate Bayesian Computation (ABC).

$$\mu = \mathbf{X}\beta$$

The dependent variable, \mathbf{y} , is normally distributed according to expectation vector μ and scalar variance σ^2 , and expectation vector μ is equal to the inner product of design matrix \mathbf{X} and parameter vector β .

For a Bayesian model, the notation for the residual variance, σ^2 , has often been replaced with the inverse of the residual precision, τ^{-1} . Here, σ^2 will be used. Prior probabilities are specified for β and σ (the standard deviation, rather than the variance):

$$\beta_j \sim \mathcal{N}(0, 1000), \quad j = 1, \dots, J$$

$$\sigma \sim \mathcal{HC}(25)$$

Each of the J β parameters is assigned an uninformative⁵ prior probability distribution that is normally-distributed according to $\mu = 0$ and $\sigma^2 = 1000$. The large variance or small precision indicates a lot of uncertainty about each β , and is hence an uninformative distribution. The residual standard deviation σ is half-Cauchy-distributed according to its hyperparameter, `scale=25`.

To specify a model, the user must create a function called `Model`. Here is an example for a linear regression model:

```
> Model <- function(parm, Data)
+   {
+     ### Parameters
+     beta <- parm[1:Data$J]
+     sigma <- exp(parm[Data$J+1])
+     ### Log(Prior Densities)
+     beta.prior <- dnormv(beta, 0, 1000, log=TRUE)
+     sigma.prior <- dhalfcauchy(sigma, 25, log=TRUE)
+     ### Log-Likelihood
+     mu <- tcrossprod(beta, Data$X)
+     LL <- sum(dnorm(Data$y, mu, sigma, log=TRUE))
+     ### Log-Posterior
+     LP <- LL + sum(beta.prior) + sigma.prior
+     Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(LP,sigma), yhat=mu,
+       parm=parm)
+     return(Modelout)
+   }
```

Laplace's Demon iteratively maximizes the logarithm of the unnormalized joint posterior density as specified in this `Model` function. In Bayesian inference, the logarithm of the unnormalized joint posterior density is proportional to the sum of the log-likelihood and logarithm of the prior densities:

$$\log[p(\Theta|\mathbf{y})] \propto \log[p(\mathbf{y}|\Theta)] + \log[p(\Theta)]$$

⁵'Non-informative' may be more widely used than 'uninformative', but here that is considered poor English, such as saying something is 'non-correct' when there's a word for that ... 'incorrect'.

where Θ is a set of parameters, \mathbf{y} is the data, \propto means ‘proportional to’⁶, $p(\Theta|\mathbf{y})$ is the joint posterior density, $p(\mathbf{y}|\Theta)$ is the likelihood, and $p(\Theta)$ is the set of prior densities.

During each iteration in which Laplace’s Demon is maximizing the logarithm of the unnormalized joint posterior density, Laplace’s Demon passes two arguments to `Model`: `parm` and `Data`, where `parm` is short for the set of parameters, and `Data` is a list of data. These arguments are specified in the beginning of the function:

```
Model <- function(parm, Data)
```

Then, the `Model` function is evaluated and the logarithm of the unnormalized joint posterior density is calculated as `LP`, and returned to Laplace’s Demon in a list called `Modelout`, along with the deviance (`Dev`), a vector (`Monitor`) of any variables desired to be monitored in addition to the parameters, \mathbf{y}^{rep} (`yhat`) or replicates of \mathbf{y} , and the parameter vector `parm`. All arguments must be returned. Even if there is no desire to observe the deviance and any monitored variable, a scalar must be placed in the second position of the `Modelout` list, and at least one element of a vector for a monitored variable. This can be seen in the end of the function:

```
LP <- LL + sum(beta.prior) + sigma.prior
Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(LP,sigma),
  yhat=mu, parm=parm)
return(Modelout)
```

The rest of the function specifies the parameters, log of the prior densities, and calculates the log-likelihood. Since design matrix \mathbf{X} has $J=10$ column vectors (including the intercept), there are 10 `beta` parameters and a `sigma` parameter for the residual standard deviation.

Since Laplace’s Demon passes a vector of parameters called `parm` to `Model`, the function needs to know which parameter is associated with which element of `parm`. For this, the vector `beta` is declared, and then each element of `beta` is populated with the value associated in the corresponding element of `parm`. The reason why `sigma` is exponentiated will, again, be explained later.

```
beta <- parm[1:Data$J]
sigma <- exp(parm[Data$J+1])
```

To work with the log of the prior densities and according to the assigned names of the parameters and hyperparameters, they are specified as follows:

```
beta.prior <- dnormv(beta, 0, 1000, log=TRUE)
sigma.prior <- dhalfcauchy(sigma, 25, log=TRUE)
```

It is important to reparameterize all parameters to be real-valued. For example, a positive-only parameter such as variance should be allowed to range from $-\infty$ to ∞ , and be transformed in the `Model` function with the `exp` function, which will force it to positive values. A parameter θ that needs to be bounded in the model, such as in the interval $[1,5]$, can be transformed to that range with a logistic function, such as $1 + 4[\exp(\theta)/(\exp(\theta) + 1)]$. Alternatively, each parameter may be constrained in the `Model` function, such as with the `interval` function. Laplace’s Demon will attempt to increase or decrease the value of each parameter to maximize `LP`, without consideration for the distributional form of the parameter. In the

⁶For those unfamiliar with \propto , this symbol simply means that two quantities are proportional if they vary in such a way that one is a constant multiplier of the other. This is due to an unspecified constant of proportionality in the equation. Here, this can be treated as ‘equal to’.

above example, the residual standard deviation `sigma` receives a half-Cauchy distributed prior of the form:

$$\sigma \sim \mathcal{HC}(25)$$

In this specification, `sigma` cannot be negative. By reparameterizing `sigma` as

```
sigma <- exp(parm[Data$J+1])
```

Laplace’s Demon will increase or decrease `parm[Data$J+1]`, which is effectively `log(sigma)`. Now it is possible for Laplace’s Demon to decrease `log(sigma)` below zero without causing an error or violating its half-Cauchy distributed specification.

Finally, everything is put together to calculate LP, the logarithm of the unnormalized joint posterior density. The expectation vector `mu` is the inner product of the design matrix, `Data$X`, and the transpose of the vector `beta`. Expectation vector `mu`, vector `Data$y`, and scalar `sigma` are used to estimate the sum of the log-likelihoods, where:

$$\mathbf{y} \sim \mathcal{N}(\mu, \sigma^2)$$

and as noted before, the logarithm of the unnormalized joint posterior density is:

$$\log[p(\Theta|\mathbf{y})] \propto \log[p(\mathbf{y}|\Theta)] + \log[p(\Theta)]$$

```
mu <- tcrossprod(Data$X, t(beta))
LL <- sum(dnorm(Data$y, mu, sigma, log=TRUE))
LP <- LL + sum(beta.prior) + sigma.prior
```

Specifying the model in the `Model` function is the most involved aspect for the user of Laplace’s Demon. But it has been designed so it is also incredibly flexible, allowing a wide variety of Bayesian models to be specified. Missing values are also easy to estimate (see the “Examples” vignette).

4. Initial Values

Laplace’s Demon requires a vector of initial values for the parameters. Each initial value is a starting point for the estimation of a parameter. When all initial values are set to zero, Laplace’s Demon will optimize initial values using a resilient backpropagation algorithm in the `LaplaceApproximation` function. Laplace Approximation is asymptotic with respect to sample size, so it is inappropriate in this example with a sample size of 39 and 11 parameters. Laplace’s Demon will not use Laplace Approximation when the sample size is not at least five times the number of parameters. Otherwise, the user may prefer to optimize initial values in the `LaplaceApproximation` function before using the `LaplacesDemon` function. When Laplace’s Demon receives initial values that are not all set to zero, it will begin to update each parameter.

In this example, there are 11 parameters. With no prior knowledge, it is a good idea either to randomize each initial value, such as with the `GIV` function (which stands for “generate initial values”), or set all of them equal to zero and let the `LaplaceApproximation` function optimize the initial values, provided there is sufficient sample size. Here, the `LaplaceApproximation`

function will be introduced in the `LaplacesDemon` function, so the first 10 parameters, the `beta` parameters, have been set equal to zero, and the remaining parameter, `log.sigma`, has been set equal to `log(1)`, which is equal to zero. This visually reminds me that I am working with the log of `sigma`, rather than `sigma`, and is merely a personal preference. The order of the elements of the vector of initial values must match the order of the parameters associated with each element of `parm` passed to the `Model` function.

```
> Initial.Values <- c(rep(0,J), log(1))
```

5. Laplace's Demon

Compared to specifying the model in the `Model` function, the actual use of Laplace's Demon is very easy. Since Laplace's Demon is stochastic, or involves pseudo-random numbers, it's a good idea to set a 'seed' for pseudo-random number generation, so results can be reproduced. Pick any number you like, but there's only one number appropriate for a demon⁷:

```
> set.seed(666)
```

As with any R package, the user can learn about a function by using the `help` function and including the name of the desired function. To learn the details of the **LaplacesDemon** function, enter:

```
> help(LaplacesDemon)
```

Here is one of many possible ways to begin:

```
> Fit <- LaplacesDemon(Model, Data=MyData, Initial.Values,
+   Covar=NULL, Iterations=2000, Status=100, Thinning=2,
+   Algorithm="AMWG", Specs=list(Periodicity=10))
```

In this example, an output object called `Fit` will be created as a result of using the **LaplacesDemon** function. `Fit` is an object of class `demonoid`, which means that since it has been assigned a customized class, other functions have been custom-designed to work with it. Laplace's Demon offers Laplace Approximation and numerous MCMC algorithms (which are explained in section 11). The above example did not use Laplace Approximation due to small sample size, and instead used the Adaptive Metropolis-within-Gibbs (AMWG) algorithm for updating.

This example tells the **LaplacesDemon** function to maximize the first component in the list output from the user-specified `Model` function, given a data set called `Data`, and according to several settings.

- The `Initial.Values` argument requires a vector of initial values for the parameters.

⁷Demonic references are used only to add flavor to the software and its use, and in no way endorse beliefs in demons. This specific pseudo-random seed is often referred to, jokingly, as the 'demon seed'.

- The `Covar=NULL` argument indicates that a user-specified variance vector or covariance matrix has not been supplied, so the algorithm will begin with its own estimate.
- The `Iterations=2000` argument indicates that the `LaplacesDemon` function will update 2,000 times before completion.
- The `Status=100` argument indicates that a status message will be printed to the R console every 100 iterations.
- The `Thinning=2` argument indicates that only every `n`th iteration will be retained in the output, and in this case, every 2nd iteration will be retained.
- The `Algorithm` argument requires the abbreviated name of the MCMC algorithm in quotes. In this case, `AMWG` is short for the Adaptive Metropolis-within-Gibbs algorithm.
- Finally, the `Specs` argument contains specifications for each algorithm named in the `Algorithm` argument. The `AMWG` algorithm requires only one specification: `Periodicity`. In this case, the `Periodicity=10` argument indicates that once adaptation begins (and `AMWG` begins immediately), the algorithm will adapt every 10 iterations. Specifications differ by algorithm. Other adaptive algorithms typically require `Adaptive` to be specified, which indicates that a non-adaptive version of the adaptive algorithm will begin so that it can gather information, and that it will become adaptive at the specified iteration. For example, `Adaptive=500` indicates that the algorithm will begin to adapt at the 500th iteration. Other specifications may also be required, depending on the algorithm.

By running the `LaplacesDemon` function, the following output was obtained:

```
> Fit <- LaplacesDemon(Model, Data=MyData, Initial.Values,
+   Covar=NULL, Iterations=2000, Status=100, Thinning=2,
+   Algorithm="AMWG", Specs=list(Periodicity=10))
```

```
Laplace's Demon was called on Mon Jun  4 11:46:19 2012
```

```
Performing initial checks...
```

```
Algorithm: Adaptive Metropolis-within-Gibbs
```

```
Laplace's Demon is beginning to update...
```

```
Iteration: 100,   Proposal: Componentwise
Iteration: 200,   Proposal: Componentwise
Iteration: 300,   Proposal: Componentwise
Iteration: 400,   Proposal: Componentwise
Iteration: 500,   Proposal: Componentwise
Iteration: 600,   Proposal: Componentwise
Iteration: 700,   Proposal: Componentwise
Iteration: 800,   Proposal: Componentwise
Iteration: 900,   Proposal: Componentwise
Iteration: 1000,  Proposal: Componentwise
Iteration: 1100,  Proposal: Componentwise
```

```
Iteration: 1200, Proposal: Componentwise
Iteration: 1300, Proposal: Componentwise
Iteration: 1400, Proposal: Componentwise
Iteration: 1500, Proposal: Componentwise
Iteration: 1600, Proposal: Componentwise
Iteration: 1700, Proposal: Componentwise
Iteration: 1800, Proposal: Componentwise
Iteration: 1900, Proposal: Componentwise
Iteration: 2000, Proposal: Componentwise
```

```
Assessing Stationarity
Assessing Thinning and ESS
Creating Summaries
Creating Output
```

Laplace's Demon has finished.

Laplace's Demon finished quickly, though it had a small data set ($N=39$), few parameters ($K=11$), and the model was very simple. The output object, `Fit`, was created as a list. As with any R object, use `str()` to examine its structure:

```
> str(Fit)
```

To access any of these values in the output object `Fit`, simply append a dollar sign and the name of the component. For example, here is how to access the observed acceptance rate:

```
> Fit$Acceptance.Rate
```

```
[1] 0.4488636
```

6. Summarizing Output

The output object, `Fit`, has many components. The (copious) contents of `Fit` can be printed to the screen with the usual R functions:

```
> Fit
> print(Fit)
```

Both return the same output, which is:

```
> Fit
```

Call:

```
LaplacesDemon(Model = Model, Data = MyData, Initial.Values = Initial.Values,
  Covar = NULL, Iterations = 2000, Status = 100, Thinning = 2,
```

Algorithm = "AMWG", Specs = list(Periodicity = 10))

Acceptance Rate: 0.44886

Adaptive: 2

Algorithm: Adaptive Metropolis-within-Gibbs

Covariance Matrix: (NOT SHOWN HERE; diagonal shown instead)

```
[1] 0.07254651 0.15313539 0.52611710 0.32452284 0.26883273 0.08655878
[7] 0.36971329 0.14092320 0.30733137 0.52198351 0.02791225
```

Covariance (Diagonal) History: (NOT SHOWN HERE)

Deviance Information Criterion (DIC):

All Stationary

Dbar 91.333 91.333

pD 379.806 379.806

DIC 471.139 471.139

Delayed Rejection (DR): 0

Initial Values:

```
[1] 0 0 0 0 0 0 0 0 0 0 0
```

Iterations: 2000

Log(Marginal Likelihood): NA

Minutes of run-time: 0.11

Model: (NOT SHOWN HERE)

Monitor: (NOT SHOWN HERE)

Parameters (Number of): 11

Periodicity: 10

Posterior1: (NOT SHOWN HERE)

Posterior2: (NOT SHOWN HERE)

Recommended Burn-In of Thinned Samples: 1

Recommended Burn-In of Un-thinned Samples: 2

Recommended Thinning: 48

Status is displayed every 100 iterations

Summary1: (SHOWN BELOW)

Summary2: (SHOWN BELOW)

Thinned Samples: 1000

Thinning: 2

Summary of All Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	5.02226985	0.2884288	0.02650105	215.18847	4.8253944
beta[2]	-0.49813076	0.3966606	0.06861114	82.64569	-1.2073272
beta[3]	-0.59171623	0.7238099	0.16408563	27.31984	-1.9587287
beta[4]	-0.05955862	0.5692229	0.10713934	53.65730	-1.2369838
beta[5]	-0.52257030	0.5205659	0.13948247	18.46695	-1.8337525
beta[6]	-0.50089658	0.2953064	0.03375876	131.99351	-1.0344104

beta[7]	2.33021125	0.6084677	0.12232181	46.55444	1.0841906
beta[8]	0.65251502	0.3752885	0.06200709	82.12562	-0.1364323
beta[9]	-0.17187270	0.5503216	0.12967205	30.34461	-1.1536858
beta[10]	1.81150568	0.7271973	0.20140623	17.30472	0.6038835
log.sigma	-0.35187059	0.1677690	0.01869313	135.77891	-0.5997676
Deviance	91.33298240	27.5610776	1.32724857	105.15859	74.8429399
LP	-93.08538028	13.7811811	0.66363206	105.15050	-120.5291468
sigma	0.72238843	0.1958220	0.02040342	130.13890	0.5278232
	Median		UB		
beta[1]	5.03798788	5.24629592			
beta[2]	-0.52312007	0.29513706			
beta[3]	-0.61718994	0.68827737			
beta[4]	-0.04278835	1.05910022			
beta[5]	-0.43862326	0.33894182			
beta[6]	-0.50079053	0.14795765			
beta[7]	2.38422530	3.40702248			
beta[8]	0.63603579	1.37584623			
beta[9]	-0.20618005	1.11747094			
beta[10]	1.78630651	3.48992133			
log.sigma	-0.36665800	-0.02735787			
Deviance	85.09887675	146.20979961			
LP	-89.97078909	-84.83955081			
sigma	0.69430263	1.08931120			

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	5.02226985	0.2884288	0.02650105	215.18847	4.8253944
beta[2]	-0.49813076	0.3966606	0.06861114	82.64569	-1.2073272
beta[3]	-0.59171623	0.7238099	0.16408563	27.31984	-1.9587287
beta[4]	-0.05955862	0.5692229	0.10713934	53.65730	-1.2369838
beta[5]	-0.52257030	0.5205659	0.13948247	18.46695	-1.8337525
beta[6]	-0.50089658	0.2953064	0.03375876	131.99351	-1.0344104
beta[7]	2.33021125	0.6084677	0.12232181	46.55444	1.0841906
beta[8]	0.65251502	0.3752885	0.06200709	82.12562	-0.1364323
beta[9]	-0.17187270	0.5503216	0.12967205	30.34461	-1.1536858
beta[10]	1.81150568	0.7271973	0.20140623	17.30472	0.6038835
log.sigma	-0.35187059	0.1677690	0.01869313	135.77891	-0.5997676
Deviance	91.33298240	27.5610776	1.32724857	105.15859	74.8429399
LP	-93.08538028	13.7811811	0.66363206	105.15050	-120.5291468
sigma	0.72238843	0.1958220	0.02040342	130.13890	0.5278232
	Median		UB		
beta[1]	5.03798788	5.24629592			
beta[2]	-0.52312007	0.29513706			
beta[3]	-0.61718994	0.68827737			
beta[4]	-0.04278835	1.05910022			
beta[5]	-0.43862326	0.33894182			

```

beta[6]    -0.50079053    0.14795765
beta[7]    2.38422530    3.40702248
beta[8]    0.63603579    1.37584623
beta[9]    -0.20618005    1.11747094
beta[10]   1.78630651    3.48992133
log.sigma  -0.36665800   -0.02735787
Deviance   85.09887675   146.20979961
LP         -89.97078909  -84.83955081
sigma      0.69430263    1.08931120

```

Several components are labeled as NOT SHOWN HERE, due to their size, such as the covariance matrix `Covar` or the stationary posterior samples `Posterior2`. As usual, these can be printed to the screen by appending a dollar sign, followed by the desired component, such as:

```
> Fit$Posterior2
```

Although a lot can be learned from the above output, notice that it completed 2000 iterations of 11 variables in 0.11 minutes. Of course this was fast, since there were only 39 records, and the form of the specified model was simple. As discussed later, Laplace's Demon does better than most other MCMC software with large numbers of records, such as 100,000 (see section 13).

In R, there is usually a `summary` function associated with each class of output object. The `summary` function usually summarizes the output. For example, with frequentist models, the `summary` function usually creates a table of parameter estimates, complete with p-values.

Since this is not a frequentist package, p-values are not part of any table with the `LaplacesDemon` function, and the marginal posterior distributions of the parameters and other variables have already been summarized in `Fit`, there is no point to have an associated `summary` function. Going one more step toward useability, `LaplacesDemon` has a `Consort` function, where the user consorts with Laplace's Demon about the output object.

Consorting with Laplace's Demon produces two kinds of output. The first section is identical to `print(Fit)`, but by consorting with Laplace's Demon, it also produces a second section called `Demonic Suggestion`.

```
> Consort(Fit)
```

```

#####
# Consort with Laplace's Demon                                     #
#####
Call:
LaplacesDemon(Model = Model, Data = MyData, Initial.Values = Initial.Values,
  Covar = NULL, Iterations = 2000, Status = 100, Thinning = 2,
  Algorithm = "AMWG", Specs = list(Periodicity = 10))

Acceptance Rate: 0.44886
Adaptive: 2
Algorithm: Adaptive Metropolis-within-Gibbs

```

Covariance Matrix: (NOT SHOWN HERE; diagonal shown instead)
 [1] 0.07254651 0.15313539 0.52611710 0.32452284 0.26883273 0.08655878
 [7] 0.36971329 0.14092320 0.30733137 0.52198351 0.02791225

Covariance (Diagonal) History: (NOT SHOWN HERE)
 Deviance Information Criterion (DIC):

All Stationary
 Dbar 91.333 91.333
 pD 379.806 379.806
 DIC 471.139 471.139

Delayed Rejection (DR): 0
 Initial Values:
 [1] 0 0 0 0 0 0 0 0 0 0 0

Iterations: 2000
 Log(Marginal Likelihood): NA
 Minutes of run-time: 0.11
 Model: (NOT SHOWN HERE)
 Monitor: (NOT SHOWN HERE)
 Parameters (Number of): 11
 Periodicity: 10
 Posterior1: (NOT SHOWN HERE)
 Posterior2: (NOT SHOWN HERE)
 Recommended Burn-In of Thinned Samples: 1
 Recommended Burn-In of Un-thinned Samples: 2
 Recommended Thinning: 48
 Status is displayed every 100 iterations
 Summary1: (SHOWN BELOW)
 Summary2: (SHOWN BELOW)
 Thinned Samples: 1000
 Thinning: 2

Summary of All Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	5.02226985	0.2884288	0.02650105	215.18847	4.8253944
beta[2]	-0.49813076	0.3966606	0.06861114	82.64569	-1.2073272
beta[3]	-0.59171623	0.7238099	0.16408563	27.31984	-1.9587287
beta[4]	-0.05955862	0.5692229	0.10713934	53.65730	-1.2369838
beta[5]	-0.52257030	0.5205659	0.13948247	18.46695	-1.8337525
beta[6]	-0.50089658	0.2953064	0.03375876	131.99351	-1.0344104
beta[7]	2.33021125	0.6084677	0.12232181	46.55444	1.0841906
beta[8]	0.65251502	0.3752885	0.06200709	82.12562	-0.1364323
beta[9]	-0.17187270	0.5503216	0.12967205	30.34461	-1.1536858
beta[10]	1.81150568	0.7271973	0.20140623	17.30472	0.6038835
log.sigma	-0.35187059	0.1677690	0.01869313	135.77891	-0.5997676

Deviance	91.33298240	27.5610776	1.32724857	105.15859	74.8429399
LP	-93.08538028	13.7811811	0.66363206	105.15050	-120.5291468
sigma	0.72238843	0.1958220	0.02040342	130.13890	0.5278232
	Median	UB			
beta[1]	5.03798788	5.24629592			
beta[2]	-0.52312007	0.29513706			
beta[3]	-0.61718994	0.68827737			
beta[4]	-0.04278835	1.05910022			
beta[5]	-0.43862326	0.33894182			
beta[6]	-0.50079053	0.14795765			
beta[7]	2.38422530	3.40702248			
beta[8]	0.63603579	1.37584623			
beta[9]	-0.20618005	1.11747094			
beta[10]	1.78630651	3.48992133			
log.sigma	-0.36665800	-0.02735787			
Deviance	85.09887675	146.20979961			
LP	-89.97078909	-84.83955081			
sigma	0.69430263	1.08931120			

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	5.02226985	0.2884288	0.02650105	215.18847	4.8253944
beta[2]	-0.49813076	0.3966606	0.06861114	82.64569	-1.2073272
beta[3]	-0.59171623	0.7238099	0.16408563	27.31984	-1.9587287
beta[4]	-0.05955862	0.5692229	0.10713934	53.65730	-1.2369838
beta[5]	-0.52257030	0.5205659	0.13948247	18.46695	-1.8337525
beta[6]	-0.50089658	0.2953064	0.03375876	131.99351	-1.0344104
beta[7]	2.33021125	0.6084677	0.12232181	46.55444	1.0841906
beta[8]	0.65251502	0.3752885	0.06200709	82.12562	-0.1364323
beta[9]	-0.17187270	0.5503216	0.12967205	30.34461	-1.1536858
beta[10]	1.81150568	0.7271973	0.20140623	17.30472	0.6038835
log.sigma	-0.35187059	0.1677690	0.01869313	135.77891	-0.5997676
Deviance	91.33298240	27.5610776	1.32724857	105.15859	74.8429399
LP	-93.08538028	13.7811811	0.66363206	105.15050	-120.5291468
sigma	0.72238843	0.1958220	0.02040342	130.13890	0.5278232
	Median	UB			
beta[1]	5.03798788	5.24629592			
beta[2]	-0.52312007	0.29513706			
beta[3]	-0.61718994	0.68827737			
beta[4]	-0.04278835	1.05910022			
beta[5]	-0.43862326	0.33894182			
beta[6]	-0.50079053	0.14795765			
beta[7]	2.38422530	3.40702248			
beta[8]	0.63603579	1.37584623			
beta[9]	-0.20618005	1.11747094			
beta[10]	1.78630651	3.48992133			

```
log.sigma -0.36665800 -0.02735787
Deviance  85.09887675 146.20979961
LP        -89.97078909 -84.83955081
sigma     0.69430263  1.08931120
```

Demonic Suggestion

Due to the combination of the following conditions,

1. Adaptive Metropolis-within-Gibbs
2. The acceptance rate (0.4488636) is within the interval [0.15,0.5].
3. At least one target MCSE is $\geq 6.27\%$ of its marginal posterior standard deviation.
4. At least one target distribution has an effective sample size (ESS) less than 100.
5. Each target distribution became stationary by 1 iterations.

Laplace's Demon has not been appeased, and suggests copy/pasting the following R code into the R console, and running it.

```
Initial.Values <- as.initial.values(Fit)
Fit <- LaplacesDemon(Model, Data=MyData, Initial.Values,
  Covar=Fit$Covar, Iterations=48000, Status=18181, Thinning=48,
  Algorithm="AMWG", Specs=list(Periodicity=107))
```

Laplace's Demon is finished consorting.

The `Demonic Suggestion` is a very helpful section of output. When Laplace's Demon was developed initially in late 2010, there were not to my knowledge any tools of Bayesian inference that make suggestions to the user.

Before making its `Demonic Suggestion`, Laplace's Demon considers and presents five conditions: the algorithm, acceptance rate, Monte Carlo standard error (MCSE), effective sample size (ESS), and stationarity. In addition to these conditions, there are other suggested values, such as a recommended number of iterations or values for the `Periodicity` and `Status` arguments. The suggested value for `Status` is seeking to print a status message every minute when the expected time is longer than a minute, and is based on the time in minutes it took, the number of iterations, and the recommended number of iterations. This estimate is fairly accurate for non-adaptive algorithms, and is harder to estimate for some adaptive algorithms that calculate the observed covariance matrix for proposals, such as AM and DRAM (explained later). But, back to the really helpful part...

If these five conditions are unsatisfactory, then Laplace's Demon is not appeased, and suggests it should continue updating, and that the user should copy/paste and execute its suggested R code. Here are the criteria it measures against. The final algorithm must be non-adaptive, so that the Markov property holds (this is covered in section 11). The acceptance rate is consid-

ered satisfactory if it is within the interval [15%,50%]⁸. MCSE is considered satisfactory for each target distribution if it is less than 6.27% of the standard deviation of the target distribution. This allows the true mean to be within 5% of the area under a Gaussian distribution around the estimated mean. ESS is considered satisfactory for each target distribution if it is at least 100, which is usually enough to describe 95% probability intervals. And finally, each variable must be estimated as stationary.

In an effort to save several pages in this tutorial due to output, the following will be run instead:

```
> Initial.Values <- as.initial.values(Fit)
> Fit <- LaplacesDemon(Model, Data=MyData, Initial.Values,
+   Covar=NULL, Iterations=70000, Status=10000, Thinning=70,
+   Algorithm="MWG", Specs=NULL)
```

Laplace's Demon was called on Mon Jun 4 11:46:26 2012

Performing initial checks...

Algorithm: Metropolis-within-Gibbs

Laplace's Demon is beginning to update...

Iteration: 10000, Proposal: Componentwise

Iteration: 20000, Proposal: Componentwise

Iteration: 30000, Proposal: Componentwise

Iteration: 40000, Proposal: Componentwise

Iteration: 50000, Proposal: Componentwise

Iteration: 60000, Proposal: Componentwise

Iteration: 70000, Proposal: Componentwise

Assessing Stationarity

Assessing Thinning and ESS

Creating Summaries

Estimating Log of the Marginal Likelihood

Creating Output

Laplace's Demon has finished.

Next, the user consorts with Laplace's Demon:

```
> Consort(Fit)
```

```
#####
# Consort with Laplace's Demon                                     #
#####
```

⁸While Spiegelhalter, Thomas, Best, and Lunn (2003) recommend updating until the acceptance rate is within the interval [20%,40%], and Roberts and Rosenthal (2001) suggest [10%,40%], the interval recommended here is [15%,50%].

Call:

```
LaplacesDemon(Model = Model, Data = MyData, Initial.Values = Initial.Values,
  Covar = NULL, Iterations = 70000, Status = 10000, Thinning = 70,
  Algorithm = "MWG", Specs = NULL)
```

Acceptance Rate: 0.42196

Adaptive: 70001

Algorithm: Metropolis-within-Gibbs

Covariance Matrix: (NOT SHOWN HERE; diagonal shown instead)

```
[1] 0.0468606 0.0468606 0.0468606 0.0468606 0.0468606 0.0468606 0.0468606
[8] 0.0468606 0.0468606 0.0468606 0.0468606
```

Covariance (Diagonal) History: (NOT SHOWN HERE)

Deviance Information Criterion (DIC):

All Stationary

Dbar 91.891 92.095

pD 519.743 563.407

DIC 611.634 655.502

Delayed Rejection (DR): 0

Initial Values:

```
[1] 5.0912999 0.3135047 0.3217194 -0.2216508 0.1680635 -0.3160828
[7] 1.7929112 0.3650947 -0.4251597 0.6036350 -0.3435412
```

Iterations: 70000

Log(Marginal Likelihood): -290.2598

Minutes of run-time: 2.02

Model: (NOT SHOWN HERE)

Monitor: (NOT SHOWN HERE)

Parameters (Number of): 11

Periodicity: 70001

Posterior1: (NOT SHOWN HERE)

Posterior2: (NOT SHOWN HERE)

Recommended Burn-In of Thinned Samples: 101

Recommended Burn-In of Un-thinned Samples: 7070

Recommended Thinning: 280

Status is displayed every 10000 iterations

Summary1: (SHOWN BELOW)

Summary2: (SHOWN BELOW)

Thinned Samples: 1000

Thinning: 70

Summary of All Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	5.04556506	0.1126520	0.003574479	1000.0000	4.8248563
beta[2]	-0.45657625	0.3529318	0.011181752	1000.0000	-1.1544148

beta[3]	-0.41110953	0.8570366	0.053396487	442.0197	-2.0537891
beta[4]	-0.09878181	0.6615657	0.032351946	559.1875	-1.4286129
beta[5]	-0.39979249	0.5095699	0.026621627	543.2583	-1.4236528
beta[6]	-0.49017679	0.2827103	0.009918262	905.3024	-1.0481164
beta[7]	2.30402406	0.5325780	0.018553925	874.7659	1.2732954
beta[8]	0.58125550	0.4046358	0.017332557	704.6629	-0.2229524
beta[9]	-0.21769847	0.5630305	0.021659681	838.9528	-1.3199996
beta[10]	1.62219490	0.7186156	0.042402902	437.9637	0.2430715
log.sigma	-0.36662200	0.1329518	0.004117531	1000.0000	-0.6186175
Deviance	91.89081331	32.2410687	1.022550846	1000.0000	74.7170636
LP	-93.36376670	16.1207443	0.511286243	1000.0000	-118.3623989
sigma	0.71148780	0.1873927	0.005880848	1000.0000	0.4902330
	Median		UB		
beta[1]	5.04467519	5.25617081			
beta[2]	-0.45724306	0.21633686			
beta[3]	-0.42768680	1.27955769			
beta[4]	-0.08114074	1.19724282			
beta[5]	-0.41649999	0.59715631			
beta[6]	-0.48935446	0.05361971			
beta[7]	2.29276189	3.35059885			
beta[8]	0.57567489	1.41475077			
beta[9]	-0.20102797	0.86312655			
beta[10]	1.64694691	3.06170554			
log.sigma	-0.37019986	-0.10863809			
Deviance	85.74399002	141.87940028			
LP	-90.29242365	-84.77418702			
sigma	0.69037135	1.03619122			

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.0456107	0.1129305	0.003803198	900.0000	4.8257812	5.0446752
beta[2]	-0.4570311	0.3525301	0.012268226	900.0000	-1.1416880	-0.4579727
beta[3]	-0.4200113	0.8536847	0.057137178	397.4777	-2.0676380	-0.4312795
beta[4]	-0.1029847	0.6478656	0.032890427	558.0625	-1.4311633	-0.0794948
beta[5]	-0.4176005	0.5111293	0.027878587	483.2898	-1.4285411	-0.4321244
beta[6]	-0.4934814	0.2868459	0.010599474	900.0000	-1.0507664	-0.4911701
beta[7]	2.3056460	0.5384893	0.019765266	796.8549	1.2614365	2.3041014
beta[8]	0.5779617	0.4022343	0.017530273	682.7413	-0.2076138	0.5687412
beta[9]	-0.2088682	0.5602649	0.021948817	773.0448	-1.3195047	-0.1940136
beta[10]	1.6408508	0.7202921	0.045540951	376.8018	0.2510485	1.6628810
log.sigma	-0.3674662	0.1334197	0.004308773	900.0000	-0.6195045	-0.3711961
Deviance	92.0953910	33.5680528	1.123896550	900.0000	74.7117577	85.7117382
LP	-93.4660774	16.7842406	0.561960712	900.0000	-119.0843251	-90.2731535
sigma	0.7106827	0.1929500	0.006332367	900.0000	0.4912335	0.6872738
	UB					
beta[1]	5.25640927					

```

beta[2]    0.21700519
beta[3]    1.27668340
beta[4]    1.16522578
beta[5]    0.59984712
beta[6]    0.05779339
beta[7]    3.35066745
beta[8]    1.37521266
beta[9]    0.86102352
beta[10]   3.06972711
log.sigma -0.10837679
Deviance  143.32300027
LP        -84.77231654
sigma     1.03178805

```

Demonic Suggestion

Due to the combination of the following conditions,

1. Metropolis-within-Gibbs
2. The acceptance rate (0.4219649) is within the interval [0.15,0.5].
3. At least one target MCSE is $\geq 6.27\%$ of its marginal posterior standard deviation.
4. Each target distribution has an effective sample size (ESS) of at least 100.
5. Each target distribution became stationary by 101 iterations.

Laplace's Demon has not been appeased, and suggests copy/pasting the following R code into the R console, and running it.

```

Initial.Values <- as.initial.values(Fit)
Fit <- LaplacesDemon(Model, Data=MyData, Initial.Values,
  Covar=Fit$Covar, Iterations=280000, Status=34653, Thinning=280,
  Algorithm="AMWG", Specs=list(Periodicity=664))

```

Laplace's Demon is finished consorting.

In 2.02 minutes, Laplace's Demon updated 70000 iterations, retaining every 70th iteration due to thinning, and reported an acceptance rate of 0.422. Notice that all criteria have been met: MCSEs are sufficiently small, ESSs are sufficiently large, and all parameters were estimated to be stationary. Since the algorithm was the non-adaptive Metropolis-within-Gibbs (MWG), the Markov property holds, so let's look at some plots.

7. Plotting Output

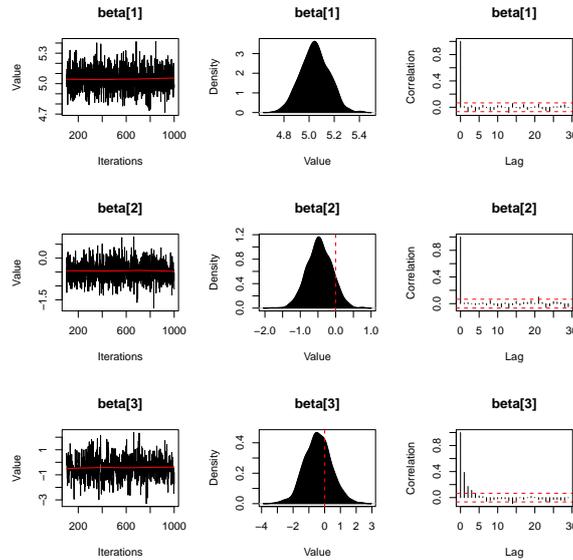


Figure 1: Plots of Marginal Posterior Samples

Laplace's Demon has a `plot.demonoid` function to enable its own customized plots with `demonoid` objects. The variable `BurnIn` (below) may be left as it is so it will show only the stationary samples (samples that are no longer trending), or set equal to one so that all samples can be plotted. In this case, it will already be one, so I will leave it alone. The function also enables the user to specify whether or not the plots should be saved as a `.pdf` file, and allows the user to limit the number of parameters plotted, in case the number is very large and only a quick glance is desired.

```
> BurnIn <- Fit$Rec.BurnIn.Thinned
```

```
> plot(Fit, BurnIn, MyData, PDF=FALSE, Params=Fit$Parameters)
```

There are three plots for each parameter, the deviance, and each monitored variable (which in this example are `sigma` and `mu[1]`). The leftmost plot is a trace-plot, showing the history of the value of the parameter according to the iteration. The middlemost plot is a kernel density plot. The rightmost plot is an ACF or autocorrelation function plot, showing the autocorrelation at different lags. The chains look stationary (do not exhibit a trend), the kernel densities look Gaussian, and the ACF's show low autocorrelation.

Another useful plot is called the caterpillar plot, which plots a horizontal representation of three quantiles (2.5%, 50%, and 97.5%) of each selected parameter from the posterior samples summary. The caterpillar plot will attempt to plot the stationary samples first (`Fit$Summary2`), but if stationary samples do not exist, then it will plot all samples (`Fit$Summary1`). Here, only the first ten parameters are selected for a caterpillar plot:

```
> caterpillar.plot(Fit, Params=1:10)
```

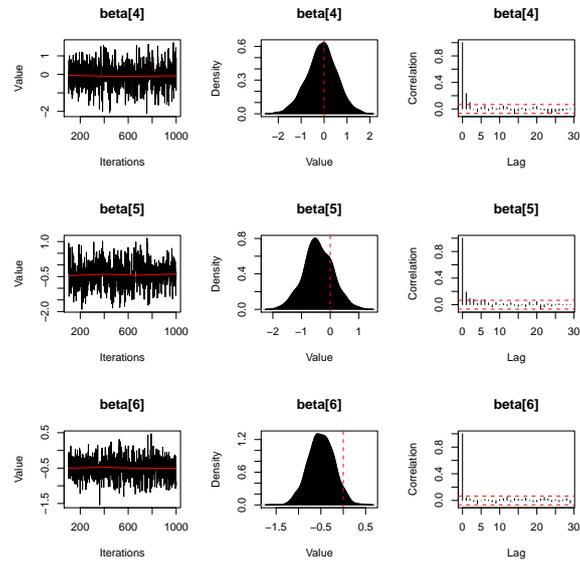


Figure 2: Plots of Marginal Posterior Samples

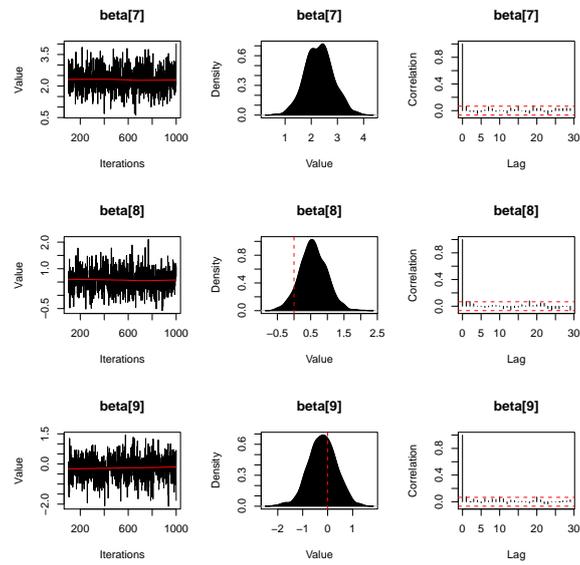


Figure 3: Plots of Marginal Posterior Samples

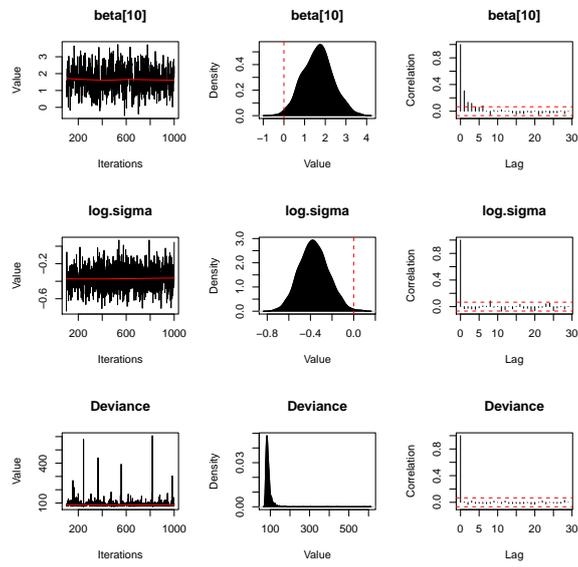


Figure 4: Plots of Marginal Posterior Samples

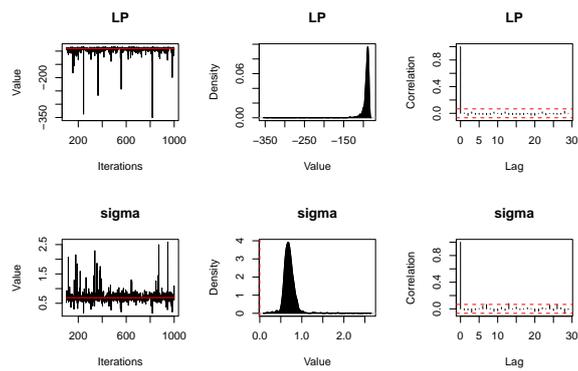


Figure 5: Plots of Marginal Posterior Samples

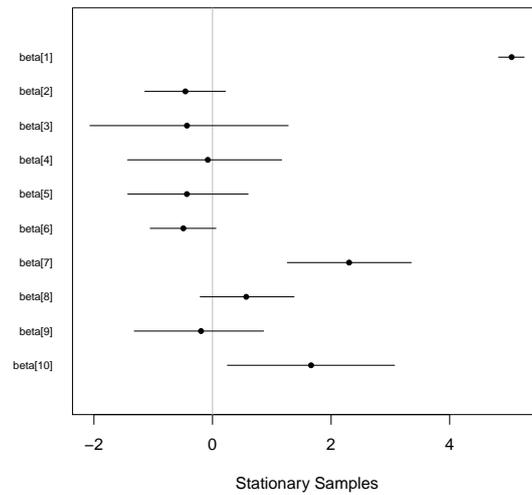


Figure 6: Caterpillar Plot

When predicting the logarithm of y (Calories) with the `demonSnacks` data, the caterpillar plot shows that the best fitting variables are `beta[6]` (Sodium), `beta[7]` (Total.Carbohydrate), and `beta[10]` (Protein). Overall, Laplace's Demon seems to have done well, eating `demonSnacks` for breakfast.

If all is well, then the Markov chains should be studied with MCMC diagnostics (such as visual inspections with the `CSF` or Cumulative Sample Function, introduced in the `LaplacesDemon` package), and finally, further assessments of model fit should be estimated with posterior predictive checks, showing how well (or poorly) the model fits the data. When the user is satisfied, the `BayesFactor` function may be useful in selecting the best model, and the marginal posterior samples may be used for inference.

8. Posterior Predictive Checks

A posterior predictive check is a method to assess discrepancies between the model and the data (Gelman, Meng, and Stern 1996a). To perform posterior predictive checks with Laplace's Demon, simply use the `predict` function:

```
> Pred <- predict(Fit, Model, MyData)
```

This creates `Pred`, which is an object of class `demonoid.ppc` (where `ppc` is short for posterior predictive check) that is a list which contains `y` and `yhat`. If the data set that was used to estimate the model is supplied in `predict`, then replicates of y (also called \mathbf{y}^{rep}) are estimated. If a new data set is supplied in `predict`, then new, unobserved instances of y (called \mathbf{y}^{new}) are estimated. Note that with new data, a `y` vector must still be supplied, and if unknown, can be set to something sensible such as the mean of the `y` vector in the model.

The `predict` function calls the `Model` function once for each set of stationary samples in `Fit$Posterior2`. Each set of samples is used to calculate `mu`, which is the expectation of `y`, and `mu` is reported here as `yhat`. When there are few discrepancies between `y` and y^{rep} , the model is considered to fit well to the data.

Since `Pred$yhat` is a large (39 x 1000) matrix, let's look at the summary of the posterior predictive distribution:

```
> summary(Pred, Discrep="Chi-Square")
```

```
Concordance: 0.7692308
```

```
Discrepancy Statistic: 310.263
```

```
L-criterion: 21.924, S.L: 0.395
```

```
Records:
```

	y	Mean	SD	LB	Median	UB	PQ	Discrep
1	4.174387	4.161	0.189	3.777	4.165	4.522	0.481	0.005
2	5.361292	5.263	0.394	4.483	5.264	6.006	0.410	0.062
3	6.089045	5.222	0.523	4.198	5.216	6.271	0.044	2.748
4	5.298317	5.128	0.319	4.459	5.129	5.761	0.291	0.284
5	4.406719	4.080	0.239	3.577	4.084	4.557	0.081	1.860
6	2.197225	3.802	0.197	3.408	3.800	4.184	1.000	66.569
7	5.010635	4.553	0.172	4.209	4.555	4.907	0.002	7.051
8	1.609438	3.862	0.194	3.490	3.864	4.238	1.000	134.593
9	4.343805	4.225	0.232	3.768	4.223	4.666	0.292	0.264
10	4.812184	4.717	0.204	4.326	4.705	5.133	0.310	0.216
11	4.189655	4.404	0.181	4.060	4.398	4.781	0.874	1.407
12	4.919981	4.533	0.175	4.193	4.534	4.884	0.017	4.895
13	4.753590	4.381	0.180	4.010	4.383	4.726	0.020	4.306
14	4.127134	4.161	0.166	3.837	4.163	4.495	0.584	0.042
15	3.713572	4.097	0.185	3.742	4.096	4.460	0.979	4.305
16	4.672829	4.389	0.223	3.925	4.395	4.817	0.102	1.618
17	6.930495	7.159	0.537	6.086	7.161	8.204	0.672	0.182
18	5.068904	4.814	0.229	4.364	4.806	5.269	0.128	1.232
19	6.775366	6.295	0.472	5.348	6.303	7.260	0.142	1.033
20	6.553933	7.197	0.458	6.257	7.204	8.076	0.920	1.974
21	4.890349	5.378	0.339	4.741	5.374	6.066	0.931	2.067
22	4.442651	4.243	0.265	3.707	4.241	4.738	0.227	0.570
23	2.833213	3.077	0.476	2.134	3.083	4.009	0.703	0.262
24	4.787492	4.934	0.243	4.480	4.930	5.469	0.730	0.363
25	6.933423	7.259	0.609	6.076	7.262	8.422	0.710	0.286
26	6.180017	6.103	0.580	4.965	6.093	7.273	0.429	0.018
27	5.652489	5.361	0.299	4.817	5.352	5.992	0.163	0.949
28	5.429346	4.468	0.196	4.085	4.474	4.864	0.000	23.957
29	5.634790	5.515	0.691	4.171	5.528	6.929	0.429	0.030
30	4.262680	4.066	0.197	3.667	4.072	4.452	0.156	0.998
31	3.891820	4.054	0.244	3.560	4.053	4.522	0.750	0.443
32	6.613384	6.595	0.384	5.853	6.599	7.340	0.483	0.002
33	4.919981	4.411	0.183	4.038	4.416	4.762	0.002	7.719

34	6.541030	6.501	0.476	5.539	6.500	7.450	0.473	0.007
35	6.345636	6.476	0.466	5.568	6.478	7.425	0.614	0.078
36	3.737670	4.060	0.249	3.549	4.060	4.552	0.908	1.671
37	7.356280	7.910	0.609	6.700	7.933	9.126	0.811	0.827
38	5.739793	4.762	0.169	4.421	4.756	5.098	0.000	33.304
39	5.517453	5.162	0.248	4.666	5.159	5.661	0.068	2.066

The `summary.demonoid.ppc` function returns a list with 4 components when `y` is continuous (different output occurs for categorical dependent variables when given the argument `Categorical=TRUE`):

- **Concordance** is the predictive concordance of [Gelfand \(1996\)](#), that indicates the percentage of times that `y` was within the 95% probability interval of `yhat`. A goal is to have 95% predictive concordance. For more information, see the accompanying vignette entitled “Bayesian Inference”. In this case, roughly 1% of the time, `y` is within the 95% probability interval of `yhat`. These results suggest that the model should be attempted again under different conditions, such as using different predictors, or specifying a different form to the model.
- **Discrepancy.Statistic** is a summary of a specified discrepancy measure. There are many options for discrepancy measures that may be specified in the `Discrep` argument. In this example, the specified discrepancy measure was the χ^2 test in [Gelman *et al.* \(2004, p. 175\)](#), and higher values indicate a worse fit.
- **L-criterion** is a posterior predictive check for model and variable selection that measures the distance between `y` and `yrep`, providing a criterion to be minimized ([Laud and Ibrahim 1995](#)).
- The last part of the summarized output reports `y`, information about the distribution of `yhat`, and the predictive quantile (PQ). The mean prediction of `y[1]`, or `y1rep`, given the model and data, is 4.161. Most importantly, `PQ[1]` is 0.481, indicating that 48.1% of the time, `yhat[1,]` was greater than `y[1]`, or that `y[1]` is close to the mean of `yhat[1,]`. Contrast this with the 6th record, where `y[6]=2.197` and `PQ[6]=1`. Therefore, `yhat[6,]` was not a good replication of `y[6]`, because the distribution of `yhat[6,]` is always greater than `y[6]`. While `y[1]` is within the 95% probability interval of `yhat[1,]`, the 95% probability interval of `yhat[6,]` is above `y[6]` 100% of the time, indicating a strong discrepancy between the model and data, in this case.

There are also a variety of plots for posterior predictive checks, and the type of plot is controlled with the `Style` argument. Many styles exist, such as producing plots of covariates and residuals. The last component of this summary may be viewed graphically as posterior densities. Rather than observing plots for each of 39 records or rows, only the first 9 densities will be shown here:

```
> plot(Pred, Style="Density", Rows=1:9)
```

The **Importance** function is not presented here in detail, but is a useful way to assess variable importance, which is defined here as the impact of each variable on `yrep`, when the variable is

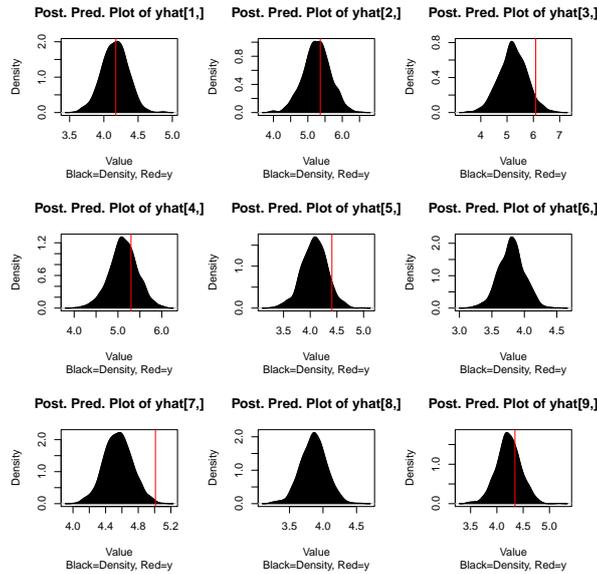


Figure 7: Posterior Predictive Plots

removed (or set to zero). Variable importance consists of differences in discrepancy statistics, showing how well the model fits the data with each variable removed. This information may be used for model revision, or presenting the relative importance of variables.

These posterior predictive checks indicate that there is plenty of room to improve this model.

9. General Suggestions

Following are general suggestions on how best to use Laplace's Demon:

- As suggested by [Gelman \(2008\)](#), continuous predictors should be centered and scaled. Here is an explicit example in R of how to center and scale a single predictor called `x`: `x.cs <- (x - mean(x)) / (2*sd(x))`. However, it is instead easier to use the `CenterScale` function provided in **LaplacesDemon**.
- Do not forget to reparameterize any bounded parameters in the `Model` function to be real-valued in the `parm` vector.
- MCMC is a stochastic method of numerical approximation, and as such, results may differ with each run due to the use of pseudo-random number generation. It is good practice to set a seed so that each update of the model may be reproduced. Here is an example in R: `set.seed(666)`.
- Rather than specify the final, intended model in the `Model` function, start by specifying the simplest possible form. Rather than beginning with actual data, start by simulating data given specified parameters. Update the simple model on simulated data and verify that the algorithm converges to the correct target distributions. One by one, add components to the model specification, simulate more complicated data, update, verify, and

progress toward the intended model. When confident the model is specified correctly, finally use actual data, but with few iterations, such as `Iterations=20`.

- After studying updates with few iterations, the first “actual” update should be long enough that proposals are accepted (the acceptance rate is not zero), adaptation begins to occur, and that enough iterations occur after the first adaptation to allow the user to study the adaptation. In the supplied example, adaptation began at the 2nd iteration (which is implicit with Adaptive Metropolis-within-Gibbs), but also occurred with `Periodicity=10`, so every 10th iteration, adaptation occurred.
- If the algorithm does not mix as expected, try a different algorithm. DRAM may be useful in increasing the acceptance rate when initial values begin in low probability regions.
- If adaptation does not seem to improve estimation or the initial movement in the chains is worse than expected, then consider optimizing the initial values with the `LaplaceApproximation` function, changing the initial values, or setting all initial values equal to zero so the `LaplacesDemon` function will use the `LaplaceApproximation` function. In MCMC, initial values are most effective when the starting points are close to the target distributions (though, if the target distributions were known *a priori*, then there would be little point in much of this). When initial values are far enough away from the target distributions to be in low-probability regions, the algorithms (both Laplace Approximation and MCMC) may take longer than usual. The MCMC algorithms herein will struggle more as the proposal covariance matrix approaches near-singularity. In extreme examples, it is possible for the proposal covariance matrix to become singular, which will stop Laplace’s Demon. If there is no information available to make a better selection, then randomize the initial values with the `GIV` function and use `LaplaceApproximation`. Centered and scaled predictors also help by essentially standardizing the possible range of the target distributions.
- If Laplace’s Demon exhibits an unreasonably low acceptance rate (say, arbitrarily, lower than 15%, but greater than 0%) and is having a hard time exploring (but is still able to explore) after significant iterations, then investigate the latest proposal covariance matrix by entering `Fit$Covar`. Chances are that the elements of the diagonal, the variances, are large. In this case, it may be best to set `Covar=NULL` for the next time it continues to update, which will begin by default with a scaled identity matrix that should get more movement in the chains. As is usual practice, the latest sampled values should also replace the initial values, so it begins from the last update, but with larger proposal variances. The chains will mix better the closer they get to their target distributions. The user can confirm that Laplace’s Demon is making progress and moving overall in the right direction by observing the trace-plots of the deviance, or better yet, the logarithm of the unnormalized joint posterior density. If the deviance is decreasing and the joint posterior is increasing run after run, then the model is continuously fitting better and better.
- When speed is a concern, such as with complex models, there may be things in the `Model` function that can be commented out, such as sometimes calculating `yhat`. The model can be updated without some features, that can be un-commented and used for posterior predictive checks. By commenting out things that are strictly unnecessary to updating, the model will update more quickly.

- If Laplace’s Demon is exploring areas of the state space that the user knows *a priori* should not be explored, then the parameters may be constrained in the `Model` function before being passed back to the `LaplacesDemon` function. Simply change the parameter of interest as appropriate and place the constrained value back in the `parm` vector.
- `Demonic Suggestion` is intended as an aid, not an infallible replacement for critical thinking. As with anything else, its suggestions are based on assumptions, and it is the responsibility of the user to check those assumptions. For example, the `Geweke.Diagnostic` may indicate stationarity (lack of a trend) when it does not exist, and this most likely occurs when too few thinned samples remain. Or, the `Demonic Suggestion` may indicate that the next update may need to run for a million iterations in a complex model, requiring weeks to complete. Is this really best for the user?
- Use a two-phase approach with Laplace’s Demon, where the first phase consists of using an adaptive algorithm (usually AMWG, AMM, AM, DRAM, RAM, or SAMWG) to achieve stationary samples that seem to have converged to the target distributions (convergence can never be determined with MCMC, but some instances of non-convergence can be observed). Once it is believed that convergence has occurred, use a non-adaptive algorithm, such as DRM, MWG, RWM, or SMWG. The final samples should again be checked for signs of non-convergence and, if satisfactory, used for inference.
- The desirable number of final, thinned samples for inference depends on the required precision of the inferential goal. A good, general goal is to end up with 1,000 thinned samples (Gelman *et al.* 2004, p. 295), where the ESS is at least 100 (and more is desirable). See the `ESS` function for more information.
- Disagreement exists in MCMC literature as to whether to update one, long chain (Geyer 1992, 2011), or multiple, long chains with different, randomized initial values (Gelman and Rubin 1992). Laplace’s Demon is not designed to simultaneously update multiple chains. Nonetheless, if multiple chains are desired, then Laplace’s Demon can be updated a series of times (simultaneously, with different sessions), each beginning with different initial values, until multiple output objects of class `demonoid` exist with stationary samples, if time allows, and the `Gelman.Diagnostic` function may be used to compare multiple chains. Samples from multiple chains may be put together with the `Combine` function.

10. Independence and Observability

Laplace’s Demon was designed with independence and observability in mind. By independence, it is meant that a goal was to minimize dependence on other software. Laplace’s Demon requires only base R. The variety of packages makes R extremely attractive. However, depending on multiple packages can be problematic when a change is made in one package, but other packages do not keep pace, and the user is dependent on packages being in sync. By avoiding dependencies on other packages, Laplace’s Demon is attempting to be consistent and dependable for the user.

For example, common MCMC diagnostics and probability distributions (such as Dirichlet, multivariate normal, Wishart, and many others, as well as truncated forms of distributions)

in Bayesian inference have been included in **LaplacesDemon** so the user does not have to load numerous R packages.

By observability, it is meant that Laplace’s Demon is written entirely in R. Certain functions could be sped up in another language, but this may prevent some R users from understanding the code. Laplace’s Demon is intended to be open and accessible. If a user desires speed and is familiar with a faster language, then the user is encouraged to program the model specification function in the faster language. See the documentation for the `Model.Spec.Time` function for more information. Moreover, it is demonstrated in section 13 that Laplace’s Demon is often significantly faster than other MCMC software that was programmed in faster languages, and users are encouraged to time comparisons, especially with large samples.

Observability also enables users to investigate or customize functions in Laplace’s Demon. To access any function, simply enter the function name and press enter. For example, to print the code for **LaplacesDemon** to the R console, simply enter:

```
> LaplacesDemon
```

To access undocumented, internal-only functions, use the `:::` operator, such as:

```
> LaplacesDemon:::RWM
```

Laplace’s Demon seeks to provide a complete, Bayesian environment within R. Independence from other software facilitates dependability, and its open code makes it easier for a user to investigate and customize.

11. Details

The **LaplacesDemon** package uses two broad types of numerical approximation algorithms: Laplace Approximation and Markov chain Monte Carlo (MCMC), and Approximate Bayesian Computation (ABC) may be estimated within each. Each is described below, but MCMC is emphasized.

11.1. Approximate Bayesian Computation

Approximate Bayesian Computation (ABC), also called likelihood-free estimation, is a family of numerical approximation techniques in Bayesian inference. ABC is especially useful when evaluation of the likelihood, $p(\mathbf{y}|\Theta)$ is computationally prohibitive, or when suitable likelihoods are unavailable. As such, ABC algorithms estimate likelihood-free approximations. ABC is usually faster than a similar likelihood-based numerical approximation technique, because the likelihood is not evaluated directly, but replaced with an approximation that is usually easier to calculate. The approximation of a likelihood is usually estimated with a measure of distance between the observed sample, \mathbf{y} , and its replicate given the model, \mathbf{y}^{rep} , or with summary statistics of the observed and replicated samples. See the accompanying vignette entitled “Examples” for an example.

11.2. Laplace Approximation

The Laplace Approximation or Laplace Method is a family of asymptotic techniques used to

approximate integrals. Laplace’s method seems to accurately approximate uni-modal posterior moments and marginal posterior distributions in many cases. Since it is not applicable in all cases, it is recommended here that Laplace Approximation is used cautiously in its own right, or preferably, it is used before MCMC.

After introducing the Laplace Approximation (Laplace 1774, p. 366–367), a proof was published later (Laplace 1814) as part of a mathematical system of inductive reasoning based on probability. Laplace used this method to approximate posterior moments.

Since its introduction, the Laplace Approximation has been applied successfully in many disciplines. In the 1980s, the Laplace Approximation experienced renewed interest, especially in statistics, and some improvements in its implementation were introduced (Tierney and Kadane 1986; Tierney, Kass, and Kadane 1989). Only since the 1980s has the Laplace Approximation been seriously considered by statisticians in practical applications.

There are many variations of Laplace Approximation, with an effort toward replacing Markov chain Monte Carlo (MCMC) algorithms as the dominant form of numerical approximation in Bayesian inference. The run-time of Laplace Approximation is a little longer than Maximum Likelihood Estimation (MLE), and much shorter than MCMC (Azevedo-Filho and Shachter 1994).

The speed of Laplace Approximation depends on the optimization algorithm selected, and typically involves many evaluations of the objective function per iteration (where the AMM MCMC algorithm evaluates once per iteration), making most of the MCMC algorithms faster per iteration. The attractiveness of Laplace Approximation is that it typically improves the objective function better than MCMC when the parameters are in low-probability regions (in which MCMC algorithms may suffer unreasonably low acceptance rates) until an adaptive MCMC has “learned” how to move better. Laplace Approximation is also typically faster because it is seeking point-estimates, rather than attempting to represent the target distribution with enough simulation draws. Laplace Approximation extends MLE, but shares similar limitations, such as its asymptotic nature with respect to sample size. Bernardo and Smith (2000) note that Laplace Approximation is an attractive numerical approximation algorithm, and will continue to develop.

`LaplaceApproximation` seeks a global maximum of the logarithm of the unnormalized joint posterior density. The approach differs by `Method`. The `LaplacesDemon` function uses the `LaplaceApproximation` algorithm to optimize initial values, estimate covariance, and save time for the user.

Most optimization algorithms assume that the logarithm of the unnormalized joint posterior density is defined and differentiable. An approximate gradient is taken for each initial value as the difference in the logarithm of the unnormalized joint posterior density due to a slight increase versus decrease in the parameter.

Adaptive Gradient Ascent

With adaptive gradient ascent, at 10 evenly-space times, `LaplaceApproximation` attempts several step sizes, which are also called rate parameters in other literature, and selects the best step size from a set of 10 fixed options. Thereafter, each iteration in which an improvement does not occur, the step size shrinks, being multiplied by 0.999.

Gradient ascent is criticized for sometimes being relatively slow when close to the maximum, and its asymptotic rate of convergence is inferior to other methods. However, compared to

other popular optimization algorithms such as Newton-Raphson, an advantage of the gradient ascent is that it works in infinite dimensions, requiring only sufficient computer memory. Although Newton-Raphson converges in fewer iterations, calculating the inverse of the negative Hessian matrix of second-derivatives is more computationally expensive and subject to singularities. Therefore, gradient ascent takes longer to converge, but is more generalizable.

Limited-Memory BFGS

The limited-memory BFGS (Broyden-Fletcher-Goldfarb-Shanno) algorithm is a quasi-Newton optimization algorithm that compactly approximates the Hessian matrix. Rather than storing the dense Hessian matrix, L-BFGS stores only a few vectors that represent the approximation. This algorithm is better suited for large-scale models than the BFGS algorithm. This is the default algorithm (`method="LBFGS"`) for `LaplaceApproximation`, which calls `method="L-BFGS-B"` in the `optim` function of base R.

Resilient Backpropagation

“Rprop” stands for resilient backpropagation. In Rprop, the approximate gradient is taken for each parameter in each iteration, and its sign is compared to the approximate gradient in the previous iteration. A weight element in a weight vector is associated with each approximate gradient. A weight element is multiplied by 1.2 when the sign does not change, or by 0.5 if the sign changes. The weight vector is the step size, and is constrained to the interval $[0.001, 50]$, and initial weights are 0.0125. This is the resilient backpropagation algorithm, which is often denoted as the “Rprop-” algorithm of [Riedmiller \(1994\)](#).

Afterward

After `LaplaceApproximation` finishes, due either to early convergence or completing the number of specified iterations, it approximates the Hessian matrix of second derivatives, and attempts to calculate the covariance matrix by taking the inverse of the negative of this matrix. If successful, then this covariance matrix may be passed to `LaplacesDemon`, and the diagonal of this matrix is the variance of the parameters. If unsuccessful, then a scaled identity matrix is returned, and each parameter’s variance will be 1.

11.3. Markov Chain Monte Carlo

Although the `LaplacesDemon` function may be assisted by Laplace Approximation, Laplace’s Demon mainly accomplishes numerical approximation with Markov chain Monte Carlo (MCMC) algorithms, also called samplers. There are a large number of MCMC algorithms, too many to review here. Popular families (which are often non-distinct) include Gibbs sampling, Metropolis-Hastings, Random-Walk Metropolis (RWM), slice sampling, and many others, including hybrid algorithms such as Metropolis-within-Gibbs ([Tierney 1994](#)), and algorithms for specific methods, such Updating Adaptive Metropolis-within-Gibbs for state-space models (SSMs). RWM was developed first ([Metropolis, Rosenbluth, M.N., and Teller 1953](#)), and Metropolis-Hastings was a generalization of RWM ([Hastings 1970](#)). All MCMC algorithms are known as special cases of the Metropolis-Hastings algorithm. Regardless of the algorithm, the goal in Bayesian inference is to maximize the unnormalized joint posterior distribution and collect samples of the target distributions, which are marginal posterior distributions,

later to be used for inference.

While designing Laplace’s Demon, the primary goal in numerical approximation was generalization. The most generalizable MCMC algorithm is the Metropolis-Hastings (MH) generalization of the RWM algorithm. The MH algorithm extended RWM to include asymmetric proposal distributions. Having no need of asymmetric proposals, Laplace’s Demon uses variations of the original RWM algorithm, which use symmetric proposal distributions, specifically Gaussian proposals (and sometimes others, such as in the RAM algorithm of [Vihola \(2011\)](#)). For years, the main disadvantage of the RWM and MH algorithms was that the proposal variance (see below) had to be tuned manually, and therefore other MCMC algorithms have become popular because they do not need to be tuned.

Gibbs sampling became popular for Bayesian inference, though it requires conditional sampling of conjugate distributions, so it is precluded from non-conjugate sampling in its purest form. Gibbs sampling also suffers under high correlations ([Gilks and Roberts 1996](#)). Due to these limitations, Gibbs sampling is less generalizable than RWM. Slice sampling samples a distribution by sampling uniformly from the region under the plot of its density function, and is more appropriate with bounded distributions that cannot approach infinity.

There are valid ways to tune the RWM algorithm as it updates. This is known by many names, including adaptive Metropolis and adaptive MCMC, among others. A brief discussion follows of RWM and its adaptive variants.

Block Updating

Usually, there is more than one target distribution, in which case it must be determined whether it is best to sample from target distributions individually, in groups, or all at once. Block updating refers to splitting a multivariate vector into groups called blocks, so each block may be treated differently. A block may contain one or more variables. Advantages of block updating are that a different MCMC algorithm may be used for each block (or variable, for that matter), creating a more specialized approach, and the acceptance of a newly proposed state is likely to be higher than sampling from all target distributions at once in high dimensions. Disadvantages of block updating are that correlations probably exist between variables between blocks, and each block is updated while holding the other blocks constant, ignoring these correlations of variables between blocks. Without simultaneously taking everything into account, the algorithm may converge slowly or never arrive at the proper solution. However, there are instances when it may be best when everything is not taken into account at once, such as in state-space models. Also, as the number of blocks increases, more computation is required, which slows the algorithm. In general, block updating allows a more specialized approach at the expense of accuracy, generalization, and speed. Laplace’s Demon generally avoids block updating, though this increases the importance that the initial values are not in low-probability regions, and may cause Laplace’s Demon to have chains that are slow to begin moving.

Random-Walk Metropolis

In MCMC algorithms, each iterative estimate of a parameter is part of a changing state. The succession of states or iterations constitutes a Markov chain when the current state is influenced only by the previous state. In random-walk Metropolis (RWM), a proposed

future estimate, called a proposal⁹ or candidate, of the joint posterior density is calculated, and a ratio of the proposed to the current joint posterior density, called α , is compared to a random number drawn uniformly from the interval (0,1). In practice, the logarithm of the unnormalized joint posterior density is used, so $\log(\alpha)$ is the proposal density minus the current density. The proposed state is accepted, replacing the current state with probability 1 when the proposed state is an improvement over the current state, and may still be accepted if the logarithm of a random draw from a uniform distribution is less than $\log(\alpha)$. Otherwise, the proposed state is rejected, and the current state is repeated so that another proposal may be estimated at the next iteration. By comparing $\log(\alpha)$ to the log of a random number when $\log(\alpha)$ is not an improvement, random-walk behavior is included in the algorithm, and it is possible for the algorithm to backtrack while it explores.

Random-walk behavior is desirable because it allows the algorithm to explore, and hopefully avoid getting trapped in undesirable regions. On the other hand, random-walk behavior is undesirable because it takes longer to converge to the target distribution while the algorithm explores. The algorithm generally progresses in the right direction, but may periodically wander away. Such exploration may uncover multimodal target distributions, which other algorithms may fail to recognize, and then converge incorrectly. With enough iterations, RWM is guaranteed theoretically to converge to the correct target distribution, regardless of the starting point of each parameter, provided the proposal variance for each proposal of a target distribution is sensible. Nonetheless, multimodal target distributions are often problematic.

Multiple parameters usually exist, and therefore correlations may occur between the parameters. Most MCMC algorithms in Laplace's Demon are modified to attempt to estimate multivariate proposals, thereby taking correlations into account through a covariance matrix. If a failure is experienced in attempting to estimate multivariate proposals in the Adaptive Metropolis (AM) of Haario, Saksman, and Tamminen (2001) here, or if the acceptance rate is less than 5%, then Laplace's Demon temporarily resorts to single-component proposals by updating one randomly-selected parameter, and will continue to attempt to return to multivariate proposals at each iteration.

Throughout the RWM algorithm, the proposal covariance or variance remains fixed. The user may enter a vector of proposal variances or a proposal covariance matrix, and if neither is supplied, then Laplace's Demon estimates both before it begins, based on the number of variables.

The acceptance or rejection of each proposal should be observed at the completion of the RWM algorithm as the acceptance rate, which is the number of acceptances divided by the total number of iterations. If the acceptance rate is too high, then the proposal variance or covariance is too small. In this case, the algorithm will take longer than necessary to find the target distribution and the samples will be highly autocorrelated. If the acceptance rate is too low, then the proposal variance or covariance is too large, and the algorithm is ineffective at exploration. In the worst case scenario, no proposals are accepted and the algorithm fails to move. Under theoretical conditions, the optimal acceptance rate for a sole, independent and identically distributed (IID), Gaussian, marginal posterior distribution is 0.44 or 44%. The

⁹Laplace's Demon allows the user to constrain proposals in the `Model` function. Laplace's Demon generates a proposal vector, which is passed to the `Model` function in the `parm` vector. In the `Model` function, the user may constrain the proposal to prevent the sampler from exploring certain areas of the state space by altering the proposed values and placing them back into the `parm` vector, which will be passed back to Laplace's Demon.

optimal acceptance rate for an infinite number of distributions that are IID and Gaussian is 0.234 or 23.4%.

Markov Chain Properties

This tutorial introduces only briefly the basics of Markov chain properties. A Markov chain is Markovian when the current iteration depends only on the previous iteration. Many (but not all) adaptive algorithms are merely chains but not Markov chains when the adaptation is based on the history of the chains, not just the previous iteration. A Markov chain is said to be aperiodic when it is not repeating a cycle. A Markov chain is considered irreducible when it is possible to go from any state to any other state, though not necessarily in one iteration. A Markov chain is said to be recurrent if it will eventually return to a given state with probability 1, and it is positive recurrent if the expected return time is finite, and null recurrent otherwise. The ergodic theorem states that a Markov chain is ergodic when it is aperiodic, irreducible, and positive recurrent.

The non-Markovian chains of an adaptive algorithm that adapt based on the history of the chains should have two conditions: containment and diminishing adaptation. Containment is difficult to implement and is not currently programmed into Laplace’s Demon. The condition of diminishing adaptation is fulfilled when the amount of adaptation diminishes with the length of the chain. Diminishing adaptation can be achieved when the proposal variances become smaller or by decreasing the probability of performing adaptations with more iterations (Roberts and Rosenthal 2007). Trace-plots of the output of the `LaplacesDemon` function automatically include plots of the absolute differences in proposal variance with each adaptation for adaptive algorithms, and the `Consort` function will try to suggest a different adaptive algorithm when these absolute differences are not trending downward.

The remaining MCMC algorithms in the **LaplacesDemon** package are now presented alphabetically.

Adaptive Metropolis

In traditional, non-adaptive RWM, the Markov property is satisfied, creating valid Markov chains, but it is difficult to manually optimize the proposal variance or covariance, and it is crucial that it is optimized for good mixing of the Markov chains. Adaptive MCMC may be used to automatically optimize the proposal variance or covariance based on the history of the chains, though this violates the Markov property, which declares the proposed state is influenced only by the current state¹⁰. To retain the Markov property, and therefore valid Markov chains, a two-phase approach may be used, in which adaptive MCMC is used in the first phase to arrive at the target distributions while violating the Markov property, and non-adaptive DRM or RWM is used in the second phase to sample from the target distributions for inference, while possessing the Markov property.

There are too many adaptive MCMC algorithms to review here. All of them adapt the proposal variance to improve mixing. Some adapt the proposal variance to optimize the acceptance rate, minimize autocorrelation, or optimize a scale factor. Laplace’s Demon uses a variation of the Adaptive Metropolis (AM) algorithm of Haario *et al.* (2001).

¹⁰Haario *et al.* (2001) assert that the chains remain ergodic in the limit as the amount of change in the adaptations should decrease to zero as the chains approach the target distributions, now referred to as the diminishing adaptation condition of Roberts and Rosenthal (2007).

Given the number of dimensions (d) or parameters, the optimal scale of the proposal variance, also called the jumping kernel, has been reported as $2.4^2/d^{11}$ based on the asymptotic limit of infinite-dimensional Gaussian target distributions that are independent and identically-distributed (Gelman, Roberts, and Gilks 1996b). In applied settings, each problem is different, so the amount of correlation varies between variables, target distributions may be non-Gaussian, the target distributions may be non-IID, and the scale should be optimized. Laplace's Demon uses a scale that is accurate to more decimals: $2.381204^2/d$. There are algorithms in statistical literature that attempt to optimize this scale, such as the RAM algorithm. Haario *et al.* (2001) tested their algorithm with up to 200 dimensions or parameters. It has been tested in Laplace's Demon with as many as 2,600 parameters, so it is capable of large-scale Bayesian inference, but you would not want to depend on it in a timely manner. The version in Laplace's Demon should be capable of more dimensions than the AM algorithm as it was presented, because when Laplace's Demon experiences an error in multivariate AM, or when the acceptance rate is less than 5%, it defaults to random-scan single-component adaptive proposals (Haario, Saksman, and Tamminen 2005). Although single-component adaptive proposals should take more iterations to converge, the algorithm is limited in dimension only by the RAM of the computer.

For multivariate adaptive tuning, the formula across K parameters and T iterations is:

$$\Sigma^* = [\phi_K \text{cov}(\Theta_{1:t,1:K})] + (\phi_K C \mathbf{I}_K)$$

where ϕ_K is the scale according to K parameters, C is a small (1.0E-5) constant to ensure the proposal covariance matrix is positive definite (does not have zero or negative variance on the diagonal), and \mathbf{I}_K is a $K \times K$ identity matrix. The initial proposal covariance matrix, when none is provided, defaults to the scaling component multiplied by its identity matrix: $\phi_K \mathbf{I}_K$. For single-component adaptive tuning, the formula across K parameters at t iterations is:

$$\sigma_k^{*2} = \phi_k \text{var}(\Theta_{1:t,k}) + \phi_k C$$

Each element in the initial vector of proposal variances is set equal to the asymptotic scale according to its dimensions: ϕ_k .

In both the multivariate and single-component cases, the AM algorithm begins with a fixed proposal variance or covariance that is either estimated internally or supplied by the user. Next, the algorithm begins, and it does not adapt until the iteration is reached that is specified by the user in the `Adaptive` argument of the `LaplacesDemon` function. Then, the algorithm will adapt with every `n` iterations according to the `Periodicity` argument. Therefore, the user has control over when the AM algorithm begins to adapt, and how often it adapts. The value of the `Adaptive` argument in Laplace's Demon is chosen subjectively by the user according to their confidence in the accuracy of the initial proposal covariance or variance. The value of the `Periodicity` argument is chosen by the user according to their patience: when the value is 1, the algorithm will adapt continuously, which will be slower to calculate. The AM algorithm adapts the proposal covariance or variance according to the observed covariance or variance in the entire history of all parameter chains, as well as the scale factor.

As recommended by Haario *et al.* (2001), there are two tricks that may be used to assist the AM algorithm in the beginning. Although Laplace's Demon does not use the suggested

¹¹The optimal proposal standard deviation in this case is approximately $2.4/\sqrt{d}$.

“greedy start” method (and will instead use Laplace Approximation when sample size permits), it uses the second suggested trick of shrinking the proposal as long as the acceptance rate is less than 5%, and there have been at least five acceptances. Haario *et al.* (2001) suggest loosely that if “it has not moved enough during some number of iterations, the proposal could be shrunk by a constant factor”. For each iteration that the acceptance rate is less than 5% and that the AM algorithm is used but the current iteration is prior to adaptation, Laplace’s Demon multiplies the proposal covariance or variance by $(1 - 1/\text{Iterations})$. Over pre-adaptive time, this encourages a smaller proposal covariance or variance to increase the acceptance rate so that when adaptation begins, the observed covariance or variance of the chains will not be constant, and then shrinkage will cease and adaptation will take it from there.

The AM algorithm performs very well in practice, though each adaptation is time-consuming after numerous iterations. The Adaptive-Mixture Metropolis (AMM) of Roberts and Rosenthal (2009) and Robust Adaptive Metropolis (Vihola 2011) are extensions of the AM algorithm.

Adaptive Metropolis-within-Gibbs

The Adaptive Metropolis-within-Gibbs (AMWG) algorithm is presented in (Roberts and Rosenthal 2009; Rosenthal 2007). The standard deviation of the proposal of each parameter is manipulated to optimize the associated acceptance rate toward 0.44. This is much simpler than other adaptive methods that adapt based on sample covariance in large dimensions. Large covariance matrices require a large number of elements to adapt, which takes exponentially longer to adapt as the dimension increases. Regardless of dimension, the AMWG optimizes each parameter to a univariate acceptance rate, and a sample covariance matrix does not need to be estimated for adaptation, which consumes time and memory. The order of the parameters for updating is randomized each iteration (random-scan AMWG), as opposed to sequential updating (deterministic-scan AMWG).

Compared to other adaptive algorithms in **LaplacesDemon**, a disadvantage is the time to complete each iteration increases as a function of parameters and model complexity, as noted in MWG. For example, in a 100-parameter model, AMWG completes its first iteration as the AMM algorithm completes its 100th. However, to adapt accurately, the AMM algorithm must correctly estimate 5,050 elements of a sample covariance matrix, while AMWG must correctly estimate only 100 proposal standard deviations. Roberts and Rosenthal (2009) have shown an example model with 500 parameters that had a burn-in of around 25,000 iterations.

The advantages of AMWG over AMM are that AMWG does not require a burn-in period before it can begin to adapt, and that AMWG does not need to estimate a covariance matrix to adapt properly. The disadvantage of AMWG compared to AMM is that AMWG solves the model function once per parameter per iteration, which can be unacceptably slow with large or complicated models. The advantage of AMWG over RAM is that AMWG does not need to estimate a covariance matrix to adapt properly. The disadvantages of AMWG compared to RAM are AMWG is less likely to handle multimodal or heavy-tailed targets, and AMWG solves the model function once per parameter per iteration, which can be unacceptably slow with large or complicated models. If AMWG is used for adaptation, then the final, non-adaptive algorithm should be MWG rather than RWM.

Adaptive-Mixture Metropolis

The Adaptive-Mixture Metropolis (AMM) algorithm is an extension by [Roberts and Rosenthal \(2009\)](#) of the AM algorithm of [Haario *et al.* \(2001\)](#). AMM differs from the AM algorithm in two respects. First, AMM updates a scatter matrix based on the cumulative current parameters and the cumulative associated outer-products, and these are used to generate a multivariate normal proposal. This is more efficient with large numbers of parameters adapting over many iterations, especially with frequent adaptations, and results in a much faster algorithm. The second (and main) difference, is that the proposal is a mixture. The two mixture components are adaptive multivariate and static/symmetric univariate proposals. The mixture is determined at each iteration with a mixture weight. The mixture weight must be in the interval $(0,1]$, and it defaults to 0.05, as in [Roberts and Rosenthal \(2009\)](#). A higher value of the mixture weight is associated with more static/symmetric univariate proposals, and a lower weight is associated with more adaptive multivariate proposals. The algorithm will be unable to include the multivariate mixture component until it has accumulated some history, and models with more parameters will take longer to be able to use adaptive multivariate proposals until later.

The advantage of AMM over AMWG is that it is much faster to update each iteration. The disadvantages are that AMWG does not require a burn-in period before it can begin to adapt, and more information must be learned in the covariance matrix to adapt properly (see AMWG for more). Disadvantages of AMM compared to RAM are that RAM does not require a burn-in period before it can begin to adapt, RAM is more likely to better handle multimodal or heavy-tailed targets, and RAM also adapts to the shape of the target distributions and coerces the acceptance rate. If AMM is used for adaptation, then the final, non-adaptive algorithm should be RWM.

Delayed Rejection Metropolis

The Delayed Rejection Metropolis (DRM or DR) algorithm is a RWM with one, small twist. Whenever a proposal is rejected, the DRM algorithm will try one or more alternate proposals, and correct for the probability of this conditional acceptance. By delaying rejection, autocorrelation in the chains may be decreased, and the algorithm is encouraged to move. Currently, Laplace's Demon will attempt one alternate proposal when using the DRAM (see below) or DRM algorithm. The additional calculations may slow each iteration of the algorithm in which the first set of proposals is rejected, but it may also converge faster. For more information on DRM, see [Mira \(2001\)](#).

DRM may be considered to be an adaptive MCMC algorithm, because it adapts the proposal based on a rejection. However, DRM does not violate the Markov property, because the proposal is based on the current state. For the purposes of Laplace's Demon, DRM is not considered to be an adaptive MCMC algorithm, because it is not adapting to the target distribution by considering previous states in the Markov chain, but merely makes more attempts from the current state. Considered as a non-adaptive algorithm, it is acceptable to conclude model updates with this algorithm, rather than following up with RWM.

Laplace's Demon also temporarily shrinks the proposal covariance arbitrarily by 50% for delayed rejection. A smaller proposal covariance is more likely to be accepted, and the goal of delayed rejection is to increase acceptance. In the long-term, a proposal covariance that is too small is undesirable, and so it is only used in this case to assist acceptance.

Each problem is different, and this can be a useful algorithm. In general, however, it is more likely that other algorithms are used.

Delayed Rejection Adaptive Metropolis

The Delayed Rejection Adaptive Metropolis (DRAM) algorithm is merely the combination of both DRM (or DR) and AM (Haario, Laine, Mira, and Saksman 2006). DRAM has been demonstrated as robust in extreme situations where DRM or AM fail separately. Haario *et al.* (2006) present an example involving ordinary differential equations in which least squares could not find a stable solution, and DRAM did well.

The DRAM algorithm is useful to assist the AM algorithm when the acceptance rate is low. As an alternative, the Adaptive-Mixture Metropolis (AMM) is an extension of the AM algorithm that includes a mixture of proposals, and one mixture component has a small proposal standard deviation to assist in overcoming initially low acceptance rates. If DRAM is used for adaptation, then the final, non-adaptive algorithm should be RWM.

Metropolis-within-Gibbs

Metropolis-within-Gibbs (MWG) is a hybrid algorithm, combining Metropolis-Hastings and Gibbs sampling, and was suggested in Tierney (1994). Also referred to as Metropolis within Gibbs or Metropolis-in-Gibbs, it is a componentwise algorithm in which the model specification function is evaluated a number of times equal to the number of parameters, per iteration. The order of the parameters for updating is randomized each iteration (random-scan MWG), as opposed to sequential updating (deterministic-scan MWG). MWG often uses blocks, but in **LaplacesDemon**, all blocks have dimension 1, meaning that each parameter is updated in turn. If parameters were grouped into blocks, then they would undesirably share a proposal standard deviation. MWG runs most efficiently when the acceptance rate of each parameter is 0.44, which is the optimal acceptance rate of a target distribution that is univariate and Gaussian.

The advantage of MWG over RWM is that it is more efficient with information per iteration, so convergence is faster in iterations. The disadvantage of MWG is that it is more time-consuming due to the evaluation of the model specification function for each parameter per iteration. As the number of parameters increases, and especially as model complexity increases, the run-time per iteration decreases. Since fewer iterations are completed in a given time-interval, the possible amount of thinning is also at a disadvantage. MWG may be used for simple models with few parameters, but is not recommended here for large and complex models.

Robust Adaptive Metropolis

The AM and AMM algorithms adapt the scale of the proposal distribution to attain a theoretical acceptance rate. However, these algorithms are unable to adapt to the shape of the target distribution. The Robust Adaptive Metropolis (RAM) algorithm estimates the shape of the target distribution and simultaneously coerces the acceptance rate (Vihola 2011). If the acceptance probability, α , is less (or greater) than an acceptance rate target, α^* , then the proposal distribution is shrunk (or expanded). Matrix \mathbf{S} is computed as a rank one Cholesky update. Therefore, the algorithm is computationally efficient up to a relatively high dimension. The AM and AMM algorithms require a burn-in period prior to adaptation, so that

these algorithms can adapt to the sample covariance. The RAM algorithm does not require a burn-in period prior to adaptation. The RAM algorithm allows the user the option of using the traditional normally-distributed proposals, or t-distributed proposals for heavier-tailed target densities. Unlike AM and AMM, RAM can cope with targets having arbitrarily heavy tails, and handles multimodal targets better than AM. The user is still assumed to know and specify the target acceptance rate.

This version of RAM does not force positive-definiteness of the variance-covariance matrix, and adapts only when it is positive-definite. Alternative versions exist elsewhere that force positive-definiteness, but in testing here, it seems better to allow it to adapt only when it is positive-definite without coercion.

RAM is slow when `Periodicity=1`, and does not seem to perform well when `Periodicity` is ≥ 100 . A general recommendation is `Periodicity=10`.

In testing the RAM algorithm, it has not been observed to obtain its acceptance rate goal and some wild fluctuations have been observed in the proposal variance after many iterations in some cases.

The advantages of RAM over AMM are that RAM does not require a burn-in period before it can begin to adapt, RAM is more likely to better handle multimodal or heavy-tailed targets, RAM also adapts to the shape of the target distributions, and coerces the acceptance rate. The advantage of RAM over AMWG is that RAM is much faster to update each iteration. The disadvantage of RAM compared to AMWG is that more information must be learned in the covariance matrix to adapt properly (see AMWG for more). If RAM is used for adaptation, then the final, non-adaptive algorithm should be RWM.

Sequential Adaptive Metropolis-within-Gibbs

The Sequential Adaptive Metropolis-within-Gibbs (SAMWG) algorithm is for state-space models (SSMs), including dynamic linear models (DLMs). It is identical to the AMWG algorithm, except with regard to the order of updating parameters (and here, sequential does not refer to deterministic-scan). Parameters are grouped into two blocks: static and dynamic. At each iteration, static parameters are updated first, followed by dynamic parameters, which are updated sequentially through the time-periods of the model. The order of the static parameters is randomly selected at each iteration, and if there are multiple dynamic parameters for each time-period, then the order of the dynamic parameters is also randomly selected. The argument `Dyn` receives a $T \times K$ matrix of T time-periods and K dynamic parameters. The SAMWG algorithm is adapted from [Geweke and Tanizaki \(2001\)](#) for `LaplacesDemon`. The SAMWG is a single-site update algorithm that is more efficient in terms of iterations, though convergence can be slow with high intercorrelations in the state vector ([Fearnhead 2011](#)). If SAMWG is used for adaptation, then the final, non-adaptive algorithm should be SMWG.

Sequential Metropolis-within-Gibbs

The Sequential Metropolis-within-Gibbs (SMWG) algorithm is the non-adaptive version of the SAMWG algorithm, and is used for final sampling of state-space models (SSMs).

t-walk

The t-walk algorithm of [Christen and Fox \(2010\)](#) is a general-purpose algorithm that requires no tuning, is scale-invariant, is technically non-adaptive (but self-adjusting), and can sample from target distributions with arbitrary scale and correlation structures. A random subset of two sets of points is moved around the state-space to influence two chains. In this implementation, the user specifies initial values for one chain, while the other vector of initial values is generated with the `GIV` function. Here, the second chain is used to help the first, and is not reported.

The authors have provided the t-walk algorithm in R code as well as other languages. It is called the “t-walk” for “traverse” or “thoughtful” walk, as opposed to RWM. Where adaptive algorithms are designed to adapt to the scale and correlation structure of target distributions, the t-walk is invariant to this structure. The step-size and direction continuously “adjust” to the local structure. Since it is technically non-adaptive, it may also be used as a final algorithm. The t-walk uses four proposal distributions or ‘moves’: walk, traverse, hop, and blow.

The t-walk has three specification arguments or tuning parameters, though the authors recommend using the default values. The first specification argument is n_1 , which affects the size of the subset of each set of points to adjust, and relates to the number of parameters. For example, if $n_1 = 4$ and a model has $J = 100$ parameters, then there is a $p(0.04) = 4/100$ probability that a point is moved that affects each parameter. Put another way, there is a 40% chance that each parameter changes. The traverse specification argument, a_t , affects the traverse move, which helps when some parameters are highly correlated, and the correlation structure may change through the state-space. The walk specification argument, a_w , affects the walk move. The authors recommend keeping these specification arguments in $n_1 \in [2, 20]$, $a_t \in [2, 10]$, and $a_w \in [0.3, 2]$. The hop and blow moves do not have specifications, but help with multimodality and ensure irreducibility.

Testing in **LaplacesDemon** with the default specifications suggests the t-walk is very promising, but due to the subset of proposals, it is important to note that the reported acceptance rate indicates the proportion of iterations in which moves were accepted, but that only a subset of parameters changed. Therefore, a user who updates a high-dimensional model should find that parameter values change much less frequently, and this requires more iterations.

The main advantage of t-walk, like the MWG family, over multivariate adaptive algorithms such as AMM and RAM is that t-walk does not adapt to a covariance structure, which can be limiting in RAM and other respects in large dimensions, making t-walk suitable for truly high-dimensional exploration. Other advantages are that t-walk is invariant to correlation structures, and does not need to burn-in before adapting, since it technically is non-adaptive (though it ‘adjusts’ continuously), and continuous adjustment is an advantage, so `Periodicity` does not need to be specified. The advantage of t-walk over single-component algorithms such as the MWG family, is that the model specification does not have to be evaluated a number of times equal to the number of parameters in each iteration, allowing the t-walk algorithm to iterate significantly faster in high dimension. The disadvantage of t-walk, compared to these algorithms, is that more iterations are required because only a subset of parameters can change at each iteration. A disadvantage compared to the RAM algorithm is that the t-walk does not adapt to the shape of the target distribution.

The t-walk algorithm seems best suited for high-dimensional problems, especially initial ex-

ploration. With enough iterations and thinning, the t-walk has produced excellent results in testing, and it has been subjected to an extreme test here on a model with 8,000 parameters. Due to limitations in computer memory, RAM, the model was updated several times for 30,000 iterations, and the thinned results were appended together. This model is large enough that, when including a covariance matrix, other algorithms such as AMM and RAM would crash the same computer. Convergence was not pursued. Though this is one, extreme and uncompleted test, the question remains as to which is better suited for high-dimensional problems: AMWG or t-walk. AMWG is a more efficient algorithm, but slows significantly as the number of parameters increases, while t-walk remains relatively fast per iteration for large problems, but the required number of iterations increases substantially.

Updating Sequential Adaptive Metropolis-within-Gibbs

The Updating Sequential Adaptive Metropolis-within-Gibbs (USAMWG) is for state-space models (SSMs), including dynamic linear models (DLMs). After a model is fit with SAMWG and SMWG, and information is later obtained regarding the first future state predicted by the model, the USAMWG algorithm may be applied to update the model given the new information. In SSM terminology, updating is filtering and predicting. The `Begin` argument tells the sampler to begin updating at a specified time-period. This is more efficient than re-estimating the entire model each time new information is obtained.

Updating Sequential Metropolis-within-Gibbs

The Updating Sequential Metropolis-within-Gibbs (USMWG) algorithm is the non-adaptive version of the USAMWG algorithm, and is used for final sampling when updating state-space models (SSMs).

Sampler Selection

When the user is ready to select a general-purpose sampler, the best place to begin, in general, is with the AMM algorithm. In testing, each example in the “Examples” vignette was estimated with AMM, AMWG, and RAM, running AMWG for 2,000 iterations, while AMM and RAM ran for 10,000 iterations. Performance was compared subjectively with a visual inspection of the chains and observing the absolute change in proposal variances with time, among other things. There are plenty of examples where each of these algorithms performs best, and the mixture is pretty even. AMM may have an increasing absolute change in proposal variance over time in some cases, but is otherwise quite stable. AMWG is prone to flat-lining some chains, increasing absolute changes in proposal variance with time, and outright failed to run in one example. RAM is prone to increase absolute changes in proposal variance with time, or more troubling, failed to run on a handful of occasions. After the initial update, the Demonic Suggestion in the `Consort` function attempts to further help with sampler selection. The t-walk will be compared here in the future. In short, experimentation is required, and sampler selection remains more of an art than science.

Afterward

Once the model is updated with the `LaplacesDemon` function, the `Geweke.Diagnostic` function of Geweke (1992) is iteratively applied to successively smaller tail-sections of the thinned samples to assess stationarity (or lack of trend). When all parameters are estimated as sta-

tionary beyond a given iteration, the previous iterations are suggested to be considered as burn-in and discarded. The number of thinned samples is divided into cumulative 10% groups, and the `Geweke.Diagnostic` function is applied by beginning with each cumulative group.

The importance of Monte Carlo Standard Error (MCSE) is debated (Gelman *et al.* 2004; Jones, Haran, Caffo, and Neath 2006). It is included in posterior summaries of `LaplacesDemon`, and is one of five main criteria as a stopping rule to appease Laplace’s Demon. MCSE has been shown to be a better stopping rule than MCMC diagnostics (Jones *et al.* 2006). Laplace’s Demon provides a MCSE function that allows three methods of estimation: sample variance, batch means (Jones *et al.* 2006), and Geyer’s method (Geyer 1992).

The user is encouraged to explore MCMC diagnostics (also called convergence diagnostics). The `LaplacesDemon` package offers a Cumulative Sample Function (CSF), Effective Sample Size (ESS), `Gelman.Diagnostic`, `Geweke.Diagnostic`, Integrated Autocorrelation Time (IAT), the Kolmogorov-Smirnov test (`KS.Diagnostic`), Monte Carlo Standard Error (MCSE), and both the `plot` and `PosteriorChecks` functions include multiple diagnostics.

12. Software Comparisons

There is now a wide variety of software to perform MCMC for Bayesian inference. Perhaps the most common is BUGS, which is an acronym for Bayesian Using Gibbs Sampling (Lunn, Spiegelhalter, Thomas, and Best 2009). BUGS has several versions. A popular variant is JAGS, which is an acronym for Just Another Gibbs Sampler (Plummer 2003). The only other comparisons made here are with some R packages (`AMCMC`, `mcmc`, `MCMCpack`, and `UMACS`) and SAS. Many other R packages use MCMC, but are not intended as general-purpose MCMC software. Hopefully, there are not any general-purpose MCMC packages in R have not been overlooked here.

WinBUGS has been the most common version of BUGS, though it is no longer developed. BUGS is an intelligent MCMC engine that is capable of numerous MCMC algorithms, but prefers Gibbs sampling. According to its user manual (Spiegelhalter *et al.* 2003), WinBUGS 1.4 uses Gibbs sampling with full conditionals that are continuous, conjugate, and standard. For full conditionals that are log-concave and non-standard, derivative-free Adaptive Rejection Sampling (ARS) is used. Slice sampling is selected for non-log-concave densities on a restricted range, and tunes itself adaptively for 500 iterations. Seemingly as a last resort, an adaptive MCMC algorithm is used for non-conjugate, continuous, full conditionals with an unrestricted range. The standard deviation of the Gaussian proposal distribution is tuned over the first 4,000 iterations to obtain an acceptance rate between 20% and 40%. Samples from the tuning phases of both Slice sampling and adaptive MCMC are ignored in the calculation of all summary statistics, although they appear in trace-plots.

The current version of BUGS, OpenBUGS, allows the user to specify an MCMC algorithm from a long list for each parameter (Lunn *et al.* 2009). This is a step forward, overcoming what is perceived here as an over-reliance on Gibbs sampling¹². However, if the user does not customize the selection of the MCMC sampler, then Gibbs sampling will be selected for full conditionals that are continuous, conjugate, and standard, just as with WinBUGS.

¹²To quote Geyer (2011), “many naive users still have a preference for Gibbs updates that is entirely unwarranted. If I had a nickel for every time someone had asked for help with slowly converging MCMC and the answer had been to stop using Gibbs, I would be rich”.

Based on years of almost daily experience with WinBUGS and JAGS, which are excellent software packages for Bayesian inference, Gibbs sampling is selected too often in these automatic, MCMC engines. An advantage of Gibbs sampling is that the proposals are accepted with probability 1, so convergence may be faster, whereas the RWM algorithm backtracks due to its random-walk behavior. Unfortunately, Gibbs sampling is not as generalizable, because it can function only when certain conjugate distributional forms are known *a priori* (Gilks and Roberts 1996). Moreover, Gibbs sampling was avoided for Laplace’s Demon because it doesn’t perform well with correlated variables or parameters, which usually exist, and I have been bitten by that *bug* many times.

The BUGS and JAGS families of MCMC software are excellent. BUGS is capable of several things that Laplace’s Demon is not. BUGS allows the user to specify the model graphically as a directed acyclic graph (DAG) in Doodle BUGS. Laplace’s Demon limits the user to one chain per parameter per update, where BUGS can update multiple chains per parameter simultaneously. Lastly, many journal articles and textbooks in several fields have been published that use BUGS, and many include example code¹³.

Advantages of **LaplacesDemon** over JAGS and WinBUGS (not much experience with **OpenBUGS**) are: Bayes factors, confidence in results (correlations do not cause trouble like in Gibbs), elicitation, environment is part of R for data manipulation and posterior analysis, examples in documentation are more plentiful, faster with large data sets (when model specification avoids loops), Importance (Variable and Parameter), Laplace Approximation, log-posterior is available, likelihood-free estimation, marginal likelihood calculated automatically, modes (functions for multimodality), missing values do not require initial values (unless predicting predictors with missing predictors), model specification gives the user complete control on how everything is calculated (including the log-likelihood, posterior, etc., and “tricks” do not have to be used), more MCMC algorithms, posterior predictive checks and discrepancy statistics, `predict` function for posterior predictive checks or scoring new data sets, suggested code is provided at the end of each run, trap errors do not exist or occur, and weights can be applied easily (such as weighting records in the likelihood).

The MCMC algorithms in Laplace’s Demon are generalizable, and generally robust to correlation between variables or parameters. The disadvantages are that convergence is slower and RWM may get stuck in regions of low probability. The advantages, however, are faster convergence when correlations are high, and more confidence in the results.

At the time this article was written, the **AMCMC** package in R is unavailable on CRAN, but may be downloaded from the author’s website¹⁴. This download is best suited for a Linux, Mac, or UNIX operating system, because it requires the `gcc` C compiler, which is unavailable in Windows. It performs adaptive Metropolis-within-Gibbs (Roberts and Rosenthal 2009; Rosenthal 2007), and uses C language, which results in significantly faster sampling, but only when the model specification function is also programmed in C. This algorithm is included in **LaplacesDemon**, where it is referred to as AMWG, for Adaptive Metropolis-within-Gibbs. The algorithm is excellent, except it is associated with long run-times for large and complex models.

Also in R, the **mcmc** package (Geyer 2010) offers RWM with multivariate Gaussian proposals and allows batching, as well as a simulated tempering algorithm, but it does not have any

¹³The first published journal article to use **LaplacesDemon** is Gallo, Miller, and Fender (2012).

¹⁴**AMCMC** is available from J. S. Rosenthal’s website at <http://www.probability.ca/amcmc/>

adaptive algorithms.

The **MCMCpack** package (Martin, Quinn, and Park 2012) in R takes a canned-function approach to RWM, which is convenient if the user needs the specific form provided, but is otherwise not generalizable. General-purpose RWM is included, but adaptive algorithms are not. It also offers the option of Laplace Approximation to optimize initial values.

At the time this article was written, the **UMACS** package (Kerman 2007) has been removed from CRAN. It became outdated due to lack of interest, but did include an adaptive MCMC algorithm as well as Gibbs sampling.

In SAS 9.2 (SAS Institute Inc. 2008), an experimental procedure called PROC MCMC has been introduced. It is undeniably a rip-off of BUGS (including its syntax), though OpenBUGS is much more powerful, tested, and generalizable. Since SAS is proprietary, the user cannot see or manipulate the source code, and should expect much more from it than OpenBUGS or any open-source software, given the absurd price.

13. Large Data Sets and Speed

An advantage of Laplace’s Demon compared to other MCMC software is that the model is specified in a way that takes advantage of R’s vectorization. BUGS and JAGS, for example, require models to be specified so that each record of data is processed one by one inside a ‘for loop’, which significantly slows updating with larger data sets. In contrast, Laplace’s Demon avoids ‘for loops’ and `apply` functions wherever possible¹⁵. For example, a data set of 100,000 rows and 16 columns (the dependent variable, a column vector of 1’s for the intercept, and 14 predictors) was updated 1,000 times with the AMM algorithm (`Adaptive=500` and `Periodicity=100`), and the initial value for each β set to 0.1 to bypass Laplace Approximation. This took 0.35 minutes with Laplace’s Demon, according to a simple, linear regression¹⁶. It was nowhere near convergence, but updating the same model with the same data for 1,000 iterations took 15.02 minutes in JAGS 3.1.0.

However, the speed with which an iteration is estimated is not a good, overall criterion of performance. For example, a Gibbs sampling algorithm with uncorrelated target distributions should converge in fewer iterations than a random-walk algorithm, such as those used in Laplace’s Demon. Depending on circumstances, Laplace’s Demon may handle larger data sets better, and it may estimate each iteration faster, but it may also take more iterations to converge¹⁷.

However, with small data sets, other MCMC software (**AMCMC** is a good example) can be faster than Laplace’s Demon, if it is programmed entirely (meaning, the model specification function as well) in a faster language such as Component Pascal, C, C++, or FORTRAN. I have not studied all MCMC algorithms in R, but most are probably programmed in C and called from R. And Laplace’s Demon could be much faster if programmed (entirely) in C as

¹⁵However, when ‘for loops’ or `apply` functions must be used, Laplace’s Demon is typically slower than BUGS.

¹⁶These updates were performed on a 2010 System76 Pangolin Performance laptop with 64-bit Debian Linux and 8GB RAM.

¹⁷To continue this example, JAGS may be *guessed* to take 20,000 iterations or 5.01 hours, and `LaplacesDemon` may take 400,000 iterations or 2.33 hours, and also have less autocorrelation in the chains due to more thinning. The slowest adaptive algorithm in Laplace’s Demon is AMWG, which updated in 5.04 minutes, and should finish around 50,000 iterations or 4.2 hours.

well.

14. Conclusion

The **LaplacesDemon** package is a significant contribution toward Bayesian inference in R. In turn, contributions toward the development of Laplace’s Demon are welcome. Please send an email to statisticat@gmail.com with constructive criticism, reports of software bugs, or offers to contribute to Laplace’s Demon.

References

- Azevedo-Filho A, Shachter R (1994). “Laplace’s Method Approximations for Probabilistic Inference in Belief Networks with Continuous Variables.” In RˆMantaras, DˆPoole (eds.), *Uncertainty in Artificial Intelligence*, pp. 28–36. Morgan Kaufman, San Francisco, CA.
- Bayes T, Price R (1763). “An Essay Towards Solving a Problem in the Doctrine of Chances. By the late Rev. Mr. Bayes, communicated by Mr. Price, in a letter to John Canton, MA. and F.R.S.” *Philosophical Transactions of the Royal Society of London*, **53**, 370–418.
- Bernardo J, Smith A (2000). *Bayesian Theory*. John Wiley & Sons, West Sussex, England.
- Christen J, Fox C (2010). “A General Purpose Sampling Algorithm for Continuous Distributions (the t-walk).” *Bayesian Analysis*, **5(2)**, 263–282.
- Crawley M (2007). *The R Book*. John Wiley & Sons Ltd, West Sussex, England.
- Fearnhead P (2011). “MCMC for State-Space Models.” In SˆBrooks, AˆGelman, GˆJones, MˆXiao-Li (eds.), *Handbook of Markov Chain Monte Carlo*, pp. 513–530. Chapman & Hall, Boca Raton, FL.
- Gallo E, Miller B, Fender R (2012). “Assessing luminosity correlations via cluster analysis: Evidence for dual tracks in the radio/X-ray domain of black hole X-ray binaries.” **1203.4263**.
- Gelfand A (1996). “Model Determination Using Sampling Based Methods.” In WˆGilks, SˆRichardson, DˆSpiegelhalter (eds.), *Markov Chain Monte Carlo in Practice*, pp. 145–161. Chapman & Hall, Boca Raton, FL.
- Gelman A (2008). “Scaling Regression Inputs by Dividing by Two Standard Deviations.” *Statistics in Medicine*, **27**, 2865–2873.
- Gelman A, Carlin J, Stern H, Rubin D (2004). *Bayesian Data Analysis*. 2nd edition. Chapman & Hall, Boca Raton, FL.
- Gelman A, Meng X, Stern H (1996a). “Posterior Predictive Assessment of Model Fitness via Realized Discrepancies.” *Statistica Sinica*, **6**, 773–807.
- Gelman A, Roberts G, Gilks W (1996b). “Efficient Metropolis Jumping Rules.” *Bayesian Statistics*, **5**, 599–608.

- Gelman A, Rubin D (1992). “Inference from Iterative Simulation Using Multiple Sequences.” *Statistical Science*, **7**(4), 457–472.
- Geweke J (1992). “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments.” *Bayesian Statistics*, **4**, 1–31.
- Geweke J, Tanizaki H (2001). “Bayesian Estimation of State-Space Models Using the Metropolis-Hastings Algorithm within Gibbs Sampling.” *Computational Statistics and Data Analysis*, **37**, 151–170.
- Geyer C (1992). “Practical Markov Chain Monte Carlo (with Discussion).” *Statistical Science*, **7**(4), 473–511.
- Geyer C (2010). *mcmc: Markov Chain Monte Carlo*. R package version 0.8, URL <http://cran.r-project.org/web/packages/mcmc/index.html>.
- Geyer C (2011). “Introduction to Markov Chain Monte Carlo.” In S~Brooks, A~Gelman, G~Jones, M~Xiao-Li (eds.), *Handbook of Markov Chain Monte Carlo*, pp. 3–48. Chapman & Hall, Boca Raton, FL.
- Gilks W, Roberts G (1996). “Strategies for Improving MCMC.” In W~Gilks, S~Richardson, D~Spiegelhalter (eds.), *Markov Chain Monte Carlo in Practice*, pp. 89–114. Chapman & Hall, Boca Raton, FL.
- Haario H, Laine M, Mira A, Saksman E (2006). “DRAM: Efficient Adaptive MCMC.” *Statistical Computing*, **16**, 339–354.
- Haario H, Saksman E, Tamminen J (2001). “An Adaptive Metropolis Algorithm.” *Bernoulli*, **7**(2), 223–242.
- Haario H, Saksman E, Tamminen J (2005). “Componentwise Adaptation for High Dimensional MCMC.” *Computational Statistics*, **20**(2), 265–274.
- Hall B (2012). *LaplacesDemon: Software for Bayesian Inference*. R package version 12.06.04, URL <http://cran.r-project.org/web/packages/LaplacesDemon/index.html>.
- Hastings W (1970). “Monte Carlo Sampling Methods Using Markov Chains and Their Applications.” *Biometrika*, **57**(1), 97–109.
- Jones G, Haran M, Caffo B, Neath R (2006). “Fixed-Width Output Analysis for Markov chain Monte Carlo.” *Journal of the American Statistical Association*, **100**(1), 1537–1547.
- Kerman J (2007). *UMACS: Universal Markov Chain Sampler*. R package version 0.924, URL <http://www.R-project.org/package=UMACS>.
- Laplace P (1774). “Memoire sur la Probabilite des Causes par les Evenements.” *l’Academie Royale des Sciences*, **6**, 621–656. English translation by S.M. Stigler in 1986 as “Memoir on the Probability of the Causes of Events” in *Statistical Science*, **1**(3), 359–378.
- Laplace P (1812). *Theorie Analytique des Probabilites*. Courcier, Paris. Reprinted as “Oeuvres Completes de Laplace”, **7**, 1878–1912. Paris: Gauthier-Villars.

- Laplace P (1814). “Essai Philosophique sur les Probabilites.” English translation in Truscott, F.W. and Emory, F.L. (2007) from (1902) as “A Philosophical Essay on Probabilities”. ISBN 1602063281, translated from the French 6th ed. (1840).
- Laud P, Ibrahim J (1995). “Posterior Model Selection.” *Journal of the Royal Statistical Society*, **B 57**, 247–262.
- Lunn D, Spiegelhalter D, Thomas A, Best N (2009). “The BUGS Project: Evolution, Critique, and Future Directions.” *Statistics in Medicine*, **28**, 3049–3067.
- Martin A, Quinn K, Park J (2012). *MCMCpack: Markov chain Monte Carlo (MCMC) Package*. R package version 1.2-2, URL <http://cran.r-project.org/web/packages/MCMCpack/index.html>.
- Metropolis N, Rosenbluth A, MN R, Teller E (1953). “Equation of State Calculations by Fast Computing Machines.” *Journal of Chemical Physics*, **21**, 1087–1092.
- Mira A (2001). “On Metropolis-Hastings Algorithms with Delayed Rejection.” *Metron*, **LIX(3–4)**, 231–241.
- Plummer M (2003). “JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling.” In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*. March 20-22, Vienna, Austria. ISBN 1609–395X.
- R Development Core Team (2012). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org>.
- Riedmiller M (1994). “Advanced Supervised Learning in Multi-Layer Perceptrons - From Backpropagation to Adaptive Learning Algorithms.” *Computer Standards and Interfaces*, **16**, 265–278.
- Roberts G, Rosenthal J (2001). “Optimal Scaling for Various Metropolis-Hastings Algorithms.” *Statistical Science*, **16**, 351–367.
- Roberts G, Rosenthal J (2007). “Coupling and Ergodicity of Adaptive Markov Chain Monte Carlo Algorithms.” *Journal of Applied Probability*, **44**, 458–475.
- Roberts G, Rosenthal J (2009). “Examples of Adaptive MCMC.” *Computational Statistics and Data Analysis*, **18**, 349–367.
- Rosenthal J (2007). “AMCMC: An R Interface for Adaptive MCMC.” *Computational Statistics and Data Analysis*, **51**, 5467–5470.
- SAS Institute Inc (2008). *SAS/STAT 9.2 User’s Guide*. Cary, NC: SAS Institute Inc.
- Spiegelhalter D, Thomas A, Best N, Lunn D (2003). *WinBUGS User Manual, Version 1.4*. MRC Biostatistics Unit, Institute of Public Health and Department of Epidemiology and Public Health, Imperial College School of Medicine, UK. <http://www.mrc-bsu.cam.ac.uk/bugs>.
- Tierney L (1994). “Markov Chains for Exploring Posterior Distributions.” *The Annals of Statistics*, **22(4)**, 1701–1762. With discussion and a rejoinder by the author.

- Tierney L, Kadane J (1986). “Accurate Approximations for Posterior Moments and Marginal Densities.” *Journal of the American Statistical Association*, **81**(393), 82–86.
- Tierney L, Kass R, Kadane J (1989). “Fully Exponential Laplace Approximations to Expectations and Variances of Nonpositive Functions.” *Journal of the American Statistical Association*, **84**(407), 710–716.
- Vihola M (2011). “Robust Adaptive Metropolis Algorithm with Coerced Acceptance Rate.” In Forthcoming (ed.), *Statistics and Computing*, pp. 1–12. Springer, Netherlands.

Affiliation:

Byron Hall

STATISTICAT, LLC

Farmington, CT

E-mail: statisticat@gmail.com

URL: <http://www.statisticat.com/laplacesdemon.html>