



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 four 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, Conjugate Gradient, Delayed Rejection, DR, DRAM, DRM, Gradient Ascent, Laplace Approximation, LaplacesDemon, Laplace’s Demon, Markov chain Monte Carlo, MCMC, Metropolis, Optimization, R, Random Walk, Random-Walk, STATISTICAT.

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 2011) package called **LaplacesDemon** (Hall 2011), 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, discusses large data sets and speed, and explains future goals.

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)
```

```
LaplacesDemon: Software for Bayesian Inference
```

```
``Probability theory is nothing but common sense reduced to
calculation'' (Pierre-Simon Laplace, 1814).
```

```
Laplace's Demon is ready for you.
```

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 **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 <- 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 **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 `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`.

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

3. Specifying a Model

Laplace’s Demon is capable with 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:

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

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

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

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 $\boldsymbol{\beta}$ 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 $\boldsymbol{\mu} = 0$ 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) {
+   beta <- parm[1:Data$J]
+   sigma <- exp(parm[Data$J + 1])
+   beta.prior <- dnorm(beta, 0, sqrt(1000), log = TRUE)
+   sigma.prior <- dhalfcauchy(sigma, 25, log = TRUE)
+   mu <- tcrossprod(beta, Data$X)
+   LL <- sum(dnorm(Data$y, mu, sigma, log = TRUE))
+   LP <- LL + sum(beta.prior) + sigma.prior
+   Modelout <- list(LP = LP, Dev = -2 * LL, Monitor = c(LP,
+     sigma), yhat = mu, parm = parm)
+   return(Modelout)
+ }
```

⁴Examples of more than 50 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).

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’.

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:

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

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 <- dnorm(beta, 0, sqrt(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 trans-

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'.

formed 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 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 $\log(\text{sigma})$. Now it is possible for Laplace's Demon to decrease $\log(\text{sigma})$ 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 vector `beta` and the transpose of the design matrix, `Data$X`. Expectation vector `mu`, vector `Data$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(beta, Data$X)
LL <- sum(dnorm(Data$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 conjugate gradient 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 within an interval, say -3 to 3, 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, Adaptive = 900, Covar = NULL,
+   DR = 1, Initial.Values, Iterations = 10000, Periodicity = 10,
+   Status = 1000, Thinning = 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 four 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 Delayed Rejection Adaptive Metropolis (DRAM) algorithm for updating.

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'.

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 **Adaptive=900** argument indicates that a non-adaptive MCMC algorithm will begin, and that it will become adaptive at the 900th iteration. Beginning with the 900th iteration, the MCMC algorithm will estimate the proposal variance or covariance based on the history of the chains.
- 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 **DR=1** argument indicates that delayed rejection will occur, such that when a proposal is rejected, an additional proposal will be attempted, thus potentially delaying rejection of proposals.
- The **Initial.Values** argument requires a vector of initial values for the parameters.
- The **Iterations=10000** argument indicates that the **LaplacesDemon** function will update 10,000 times before completion.
- The **Periodicity=10** argument indicates that once adaptation begins, the algorithm will adapt every 10 iterations.
- The **Status=1000** argument indicates that a status message will be printed to the R console every 1,000 iterations.
- Finally, the **Thinning=10** argument indicates that only every *n*th iteration will be retained in the output, and in this case, every 10th iteration will be retained.

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

```
> Fit <- LaplacesDemon(Model, Data = MyData, Adaptive = 900, Covar = NULL,
+   DR = 1, Initial.Values, Iterations = 10000, Periodicity = 10,
+   Status = 1000, Thinning = 10)
```

```
Laplace's Demon was called on Sun Oct  2 13:05:39 2011
```

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

```
Algorithm: Delayed Rejection Adaptive Metropolis
```

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

```
Iteration: 1000, Proposal: Multivariate
```

```
Iteration: 2000, Proposal: Multivariate
```

```
Iteration: 3000, Proposal: Multivariate
```

```
Iteration: 4000, Proposal: Multivariate
```

```
Iteration: 5000, Proposal: Multivariate
```

```
Iteration: 6000, Proposal: Multivariate
```

```
Iteration: 7000, Proposal: Multivariate
```

```
Iteration: 8000, Proposal: Multivariate
Iteration: 9000, Proposal: Multivariate
```

```
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. At each status of 1000 iterations, the proposal was multivariate, so it did not have to resort to single-component proposals. 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] 0.3098
```

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, Adaptive = 900, Covar = NULL,
  DR = 1, Initial.Values = Initial.Values, Iterations = 10000,
  Periodicity = 10, Status = 1000, Thinning = 10)
```

```
Acceptance Rate: 0.3098
```

```
Adaptive: 900
```

```
Algorithm: Delayed Rejection Adaptive Metropolis
```

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

[1] 0.08455135 0.06651354 0.44405792 0.26473941 0.13025956 0.05335009
 [7] 0.13046451 0.08367160 0.13312048 0.28243715 0.02832220

Covariance (Diagonal) History: (NOT SHOWN HERE)

Deviance Information Criterion (DIC):

All Stationary

Dbar 94.714 86.939
 pD 8623.997 123.577
 DIC 8718.710 210.516

Delayed Rejection (DR): 1

Initial Values:

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

Iterations: 10000

Log(Marginal Likelihood): NA

Minutes of run-time: 0.29

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: 301

Recommended Burn-In of Un-thinned Samples: 3010

Recommended Thinning: 280

Status is displayed every 1000 iterations

Summary1: (SHOWN BELOW)

Summary2: (SHOWN BELOW)

Thinned Samples: 1000

Thinning: 10

Summary of All Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	4.9883587	0.4198169	0.06106423	47.26569	4.6496986
beta[2]	-0.4586643	0.3568350	0.02598660	188.55415	-1.1569673
beta[3]	-0.4699650	0.9385360	0.08183581	131.52707	-2.4648023
beta[4]	-0.0243241	0.7272354	0.06253795	135.22681	-1.4920584
beta[5]	-0.3976742	0.5040661	0.03957248	162.25139	-1.3702860
beta[6]	-0.4945990	0.3227069	0.01851768	303.69898	-1.0864305
beta[7]	2.2995157	0.5069884	0.03691111	188.66090	1.3055595
beta[8]	0.6365851	0.4069304	0.02834086	206.16477	-0.1674040
beta[9]	-0.2606857	0.5109132	0.04311193	140.44271	-1.2542463
beta[10]	1.6293691	0.7484949	0.07986923	87.82512	0.1807337
log.sigma	-0.3537363	0.2311960	0.03081725	56.28244	-0.6138484
Deviance	94.7137295	131.3316158	7.71728098	289.60760	73.3766875

LP	-94.7754631	65.6664497	3.85878628	289.59137	-128.6801676
sigma	0.7421042	0.3996368	0.05179764	59.52661	0.5058286
	Median		UB		
beta[1]	5.036027059	5.269213918			
beta[2]	-0.460612366	0.275103279			
beta[3]	-0.474619477	1.341343058			
beta[4]	-0.008943934	1.366729547			
beta[5]	-0.399333645	0.617275600			
beta[6]	-0.500263950	0.079923228			
beta[7]	2.317191232	3.330379444			
beta[8]	0.636632082	1.476951176			
beta[9]	-0.260298995	0.674791650			
beta[10]	1.622834531	3.084998541			
log.sigma	-0.381872538	-0.005642328			
Deviance	82.325599382	162.525733351			
LP	-88.581414468	-84.104397038			
sigma	0.680273835	1.138040074			

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.0452164	0.1209224	0.004886365	612.4092	4.8183569	5.04429637
beta[2]	-0.4671705	0.3626808	0.028336287	163.8187	-1.1968524	-0.46441583
beta[3]	-0.3692309	0.9048893	0.075864466	142.2703	-2.0777775	-0.41493101
beta[4]	-0.0938737	0.7137761	0.056098580	161.8901	-1.5789661	-0.06584369
beta[5]	-0.3307827	0.5285388	0.045642283	134.0970	-1.3343019	-0.32294631
beta[6]	-0.4764229	0.2892152	0.020386227	201.2652	-1.0340627	-0.47390391
beta[7]	2.3047324	0.5034679	0.039984174	158.5504	1.3984978	2.30820182
beta[8]	0.6231165	0.4194543	0.033060651	160.9706	-0.2304077	0.62648569
beta[9]	-0.2492006	0.5146298	0.040653599	160.2477	-1.2330966	-0.27251875
beta[10]	1.5464835	0.7636212	0.062942080	147.1885	0.1393432	1.54110602
log.sigma	-0.3705254	0.1321553	0.009543010	191.7778	-0.6107904	-0.37449456
Deviance	86.9389279	15.7211098	0.594202098	700.0000	74.2382550	82.80023276
LP	-90.8877241	7.8613667	0.297131732	700.0000	-111.8142447	-88.81699766
sigma	0.7014016	0.1253897	0.008305580	227.9206	0.5067230	0.68266718
	UB					
beta[1]	5.2715971					
beta[2]	0.2425978					
beta[3]	1.3843414					
beta[4]	1.2143440					
beta[5]	0.7439408					
beta[6]	0.0765831					
beta[7]	3.3820399					
beta[8]	1.4883015					
beta[9]	0.6859437					
beta[10]	3.0234399					
log.sigma	-0.1063204					

```
Deviance 128.7923598
LP       -84.5347748
sigma    0.9929525
```

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 10000 iterations of 11 variables in 0.29 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, Adaptive = 900, Covar = NULL,
  DR = 1, Initial.Values = Initial.Values, Iterations = 10000,
  Periodicity = 10, Status = 1000, Thinning = 10)

Acceptance Rate: 0.3098
Adaptive: 900
Algorithm: Delayed Rejection Adaptive Metropolis
Covariance Matrix: (NOT SHOWN HERE; diagonal shown instead)
 [1] 0.08455135 0.06651354 0.44405792 0.26473941 0.13025956 0.05335009
 [7] 0.13046451 0.08367160 0.13312048 0.28243715 0.02832220

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

All Stationary
 Dbar 94.714 86.939
 pD 8623.997 123.577
 DIC 8718.710 210.516

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

Iterations: 10000
 Log(Marginal Likelihood): NA
 Minutes of run-time: 0.29
 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: 301
 Recommended Burn-In of Un-thinned Samples: 3010
 Recommended Thinning: 280
 Status is displayed every 1000 iterations
 Summary1: (SHOWN BELOW)
 Summary2: (SHOWN BELOW)
 Thinned Samples: 1000
 Thinning: 10

Summary of All Samples

	Mean	SD	MCSE	ESS	LB
beta[1]	4.9883587	0.4198169	0.06106423	47.26569	4.6496986
beta[2]	-0.4586643	0.3568350	0.02598660	188.55415	-1.1569673
beta[3]	-0.4699650	0.9385360	0.08183581	131.52707	-2.4648023
beta[4]	-0.0243241	0.7272354	0.06253795	135.22681	-1.4920584
beta[5]	-0.3976742	0.5040661	0.03957248	162.25139	-1.3702860
beta[6]	-0.4945990	0.3227069	0.01851768	303.69898	-1.0864305
beta[7]	2.2995157	0.5069884	0.03691111	188.66090	1.3055595
beta[8]	0.6365851	0.4069304	0.02834086	206.16477	-0.1674040
beta[9]	-0.2606857	0.5109132	0.04311193	140.44271	-1.2542463
beta[10]	1.6293691	0.7484949	0.07986923	87.82512	0.1807337
log.sigma	-0.3537363	0.2311960	0.03081725	56.28244	-0.6138484
Deviance	94.7137295	131.3316158	7.71728098	289.60760	73.3766875
LP	-94.7754631	65.6664497	3.85878628	289.59137	-128.6801676
sigma	0.7421042	0.3996368	0.05179764	59.52661	0.5058286
	Median	UB			
beta[1]	5.036027059	5.269213918			
beta[2]	-0.460612366	0.275103279			

beta[3]	-0.474619477	1.341343058
beta[4]	-0.008943934	1.366729547
beta[5]	-0.399333645	0.617275600
beta[6]	-0.500263950	0.079923228
beta[7]	2.317191232	3.330379444
beta[8]	0.636632082	1.476951176
beta[9]	-0.260298995	0.674791650
beta[10]	1.622834531	3.084998541
log.sigma	-0.381872538	-0.005642328
Deviance	82.325599382	162.525733351
LP	-88.581414468	-84.104397038
sigma	0.680273835	1.138040074

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.0452164	0.1209224	0.004886365	612.4092	4.8183569	5.04429637
beta[2]	-0.4671705	0.3626808	0.028336287	163.8187	-1.1968524	-0.46441583
beta[3]	-0.3692309	0.9048893	0.075864466	142.2703	-2.0777775	-0.41493101
beta[4]	-0.0938737	0.7137761	0.056098580	161.8901	-1.5789661	-0.06584369
beta[5]	-0.3307827	0.5285388	0.045642283	134.0970	-1.3343019	-0.32294631
beta[6]	-0.4764229	0.2892152	0.020386227	201.2652	-1.0340627	-0.47390391
beta[7]	2.3047324	0.5034679	0.039984174	158.5504	1.3984978	2.30820182
beta[8]	0.6231165	0.4194543	0.033060651	160.9706	-0.2304077	0.62648569
beta[9]	-0.2492006	0.5146298	0.040653599	160.2477	-1.2330966	-0.27251875
beta[10]	1.5464835	0.7636212	0.062942080	147.1885	0.1393432	1.54110602
log.sigma	-0.3705254	0.1321553	0.009543010	191.7778	-0.6107904	-0.37449456
Deviance	86.9389279	15.7211098	0.594202098	700.0000	74.2382550	82.80023276
LP	-90.8877241	7.8613667	0.297131732	700.0000	-111.8142447	-88.81699766
sigma	0.7014016	0.1253897	0.008305580	227.9206	0.5067230	0.68266718
	UB					
beta[1]	5.2715971					
beta[2]	0.2425978					
beta[3]	1.3843414					
beta[4]	1.2143440					
beta[5]	0.7439408					
beta[6]	0.0765831					
beta[7]	3.3820399					
beta[8]	1.4883015					
beta[9]	0.6859437					
beta[10]	3.0234399					
log.sigma	-0.1063204					
Deviance	128.7923598					
LP	-84.5347748					
sigma	0.9929525					

Demonic Suggestion

Due to the combination of the following conditions,

1. Delayed Rejection Adaptive Metropolis
2. The acceptance rate (0.3098) 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 301 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 <- Fit$Posterior1[Fit$Thinned.Samples,]
Fit <- LaplacesDemon(Model, Data=MyData, Adaptive=36,
  Covar=Fit$Covar, DR=0, Initial.Values, Iterations=280000,
  Periodicity=904, Status=34482, Thinning=280)
```

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. There are 48 combinations of these five conditions, though many combinations lead to the same conclusions. 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 hard to estimate for adaptive algorithms. 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 considered 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

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%].

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.

Notice that since stationarity has been estimated beginning with the 301st iteration, the suggested R code changes from `Adaptive=900` to `Adaptive=0`. The suggestion is to abandon the adaptive MCMC algorithm in favor of a non-adaptive algorithm, specifically a Random-Walk Metropolis (RWM). It is also replacing the initial values with the latest values of the parameter chains, and is suggesting to begin with the latest covariance matrix. Some of the arguments in the suggested R code seem excessive, such as `Iterations=280000` and `Thinning=280`. For the sake of the example and saving the reader from a few pages of output, the suggested R code will not be run and the following will be run instead:

```
> Initial.Values <- Fit$Posterior1[Fit$Thinned.Samples, ]
> Fit <- LaplacesDemon(Model, Data = MyData, Adaptive = 0, Covar = Fit$Covar,
+   DR = 0, Initial.Values, Iterations = 290000, Periodicity = 0,
+   Status = 10000, Thinning = 290)
```

Laplace's Demon was called on Sun Oct 2 13:05:56 2011

Performing initial checks...

Adaptation will not occur due to the Adaptive argument.

Adaptation will not occur due to the Periodicity argument.

Algorithm: Random-Walk Metropolis

Laplace's Demon is beginning to update...

```
Iteration: 10000, Proposal: Single-Component
Iteration: 20000, Proposal: Multivariate
Iteration: 30000, Proposal: Multivariate
Iteration: 40000, Proposal: Multivariate
Iteration: 50000, Proposal: Multivariate
Iteration: 60000, Proposal: Multivariate
Iteration: 70000, Proposal: Multivariate
Iteration: 80000, Proposal: Multivariate
Iteration: 90000, Proposal: Multivariate
Iteration: 100000, Proposal: Multivariate
Iteration: 110000, Proposal: Multivariate
Iteration: 120000, Proposal: Multivariate
Iteration: 130000, Proposal: Multivariate
Iteration: 140000, Proposal: Multivariate
Iteration: 150000, Proposal: Multivariate
Iteration: 160000, Proposal: Multivariate
Iteration: 170000, Proposal: Multivariate
Iteration: 180000, Proposal: Multivariate
Iteration: 190000, Proposal: Multivariate
Iteration: 200000, Proposal: Multivariate
Iteration: 210000, Proposal: Multivariate
Iteration: 220000, Proposal: Multivariate
```

```
Iteration: 230000, Proposal: Multivariate
Iteration: 240000, Proposal: Multivariate
Iteration: 250000, Proposal: Multivariate
Iteration: 260000, Proposal: Multivariate
Iteration: 270000, Proposal: Multivariate
Iteration: 280000, Proposal: Multivariate
```

```
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 #
#####
```

Call:

```
LaplacesDemon(Model = Model, Data = MyData, Adaptive = 0, Covar = Fit$Covar,
  DR = 0, Initial.Values = Initial.Values, Iterations = 290000,
  Periodicity = 0, Status = 10000, Thinning = 290)
```

Acceptance Rate: 0.1806

Adaptive: 290001

Algorithm: Random-Walk Metropolis

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

```
[1] 0.08455135 0.06651354 0.44405792 0.26473941 0.13025956 0.05335009
[7] 0.13046451 0.08367160 0.13312048 0.28243715 0.02832220
```

Covariance (Diagonal) History: (NOT SHOWN HERE)

Deviance Information Criterion (DIC):

```
      All Stationary
Dbar 82.055      82.128
pD   15.054     14.740
DIC  97.109     96.869
```

Delayed Rejection (DR): 0

Initial Values:

```
      beta[1]      beta[2]      beta[3]      beta[4]      beta[5]      beta[6]
5.035727314 -0.001927554 -0.661212565 0.150148260 -0.579904003 -0.423509217
      beta[7]      beta[8]      beta[9]      beta[10]      log.sigma
1.124764509 0.838525731 0.257052949 1.837345368 -0.540429186
```

Iterations: 290000
 Log(Marginal Likelihood): -86.60886
 Minutes of run-time: 2.18
 Model: (NOT SHOWN HERE)
 Monitor: (NOT SHOWN HERE)
 Parameters (Number of): 11
 Periodicity: 290001
 Posterior1: (NOT SHOWN HERE)
 Posterior2: (NOT SHOWN HERE)
 Recommended Burn-In of Thinned Samples: 201
 Recommended Burn-In of Un-thinned Samples: 58290
 Recommended Thinning: 870
 Status is displayed every 10000 iterations
 Summary1: (SHOWN BELOW)
 Summary2: (SHOWN BELOW)
 Thinned Samples: 1000
 Thinning: 290

Summary of All Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.0476190	0.10984847	0.003473714	1000.0000	4.8270261	5.04823623
beta[2]	-0.4662030	0.36526652	0.018377095	395.0627	-1.2321903	-0.45438351
beta[3]	-0.4915055	0.91307553	0.036206756	635.9664	-2.2930568	-0.46420443
beta[4]	-0.0579340	0.67729967	0.021418096	1000.0000	-1.3839292	-0.06781109
beta[5]	-0.4058416	0.50043946	0.017553949	812.7434	-1.4253946	-0.40363787
beta[6]	-0.5071423	0.30318141	0.009587438	1000.0000	-1.1221300	-0.49588295
beta[7]	2.3018995	0.54584038	0.022036180	613.5624	1.2251696	2.31786996
beta[8]	0.6399804	0.42184745	0.016071664	688.9523	-0.1619606	0.65994613
beta[9]	-0.2065339	0.57541339	0.021552593	712.7888	-1.3357013	-0.21631982
beta[10]	1.6537835	0.75428713	0.026310818	821.8730	0.1840727	1.63794373
log.sigma	-0.3683079	0.13378099	0.004230526	1000.0000	-0.6010445	-0.36843782
Deviance	82.0549970	5.48702667	0.173515019	1000.0000	73.5924230	81.20998107
LP	-88.4458976	2.74451043	0.086789040	1000.0000	-94.6280743	-88.02410585
sigma	0.6993738	0.09596913	0.003034810	1000.0000	0.5486774	0.69294099

UB

beta[1]	5.26180805
beta[2]	0.22075298
beta[3]	1.29063870
beta[4]	1.28778198
beta[5]	0.57286988
beta[6]	0.08262073
beta[7]	3.29668580
beta[8]	1.40908385
beta[9]	0.89019461
beta[10]	3.08867997

```

log.sigma -0.09797562
Deviance  94.41662369
LP        -84.21281351
sigma     0.91937490

```

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.0487614	0.11098695	0.003923981	800.0000	4.8276524	5.0478311
beta[2]	-0.4462911	0.36804981	0.013012526	800.0000	-1.1984881	-0.4311838
beta[3]	-0.4127919	0.89952083	0.031802864	800.0000	-2.2471119	-0.3759962
beta[4]	-0.1005763	0.66522778	0.023519354	800.0000	-1.3871856	-0.1070156
beta[5]	-0.3881095	0.49963146	0.017664640	800.0000	-1.4203551	-0.3832188
beta[6]	-0.5086246	0.31026265	0.010969441	800.0000	-1.1329532	-0.4950606
beta[7]	2.2831323	0.53905984	0.022160873	591.6982	1.1967685	2.3114036
beta[8]	0.6057946	0.42616788	0.015067310	800.0000	-0.2380697	0.6115900
beta[9]	-0.1951843	0.56778306	0.020074163	800.0000	-1.3339505	-0.2052019
beta[10]	1.6169915	0.74430681	0.026315220	800.0000	0.1840727	1.5960083
log.sigma	-0.3642619	0.13404312	0.004739140	800.0000	-0.6073474	-0.3633139
Deviance	82.1283571	5.42958388	0.191964779	800.0000	73.9554532	81.2635926
LP	-88.4823987	2.71566019	0.096013087	800.0000	-94.6280743	-88.0484537
sigma	0.7008271	0.09527294	0.003368407	800.0000	0.5474287	0.6950558
UB						
beta[1]	5.26483478					
beta[2]	0.26423908					
beta[3]	1.30052180					
beta[4]	1.20910226					
beta[5]	0.63976566					
beta[6]	0.08456518					
beta[7]	3.28063670					
beta[8]	1.39014195					
beta[9]	0.86841712					
beta[10]	3.02095238					
log.sigma	-0.09041084					
Deviance	94.41662369					
LP	-84.39394118					
sigma	0.92059392					

Demonic Suggestion

Due to the combination of the following conditions,

1. Random-Walk Metropolis
2. The acceptance rate (0.1806) is within the interval [0.15,0.5].
3. Each target MCSE is < 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 201 iterations.

Laplace's Demon has been appeased, and suggests the marginal posterior samples should be plotted and subjected to any other MCMC diagnostic deemed fit before using these samples for inference.

Laplace's Demon is finished consorting.

In 2.18 minutes, Laplace's Demon updated 290000 iterations, retaining every 290th iteration due to thinning, and reported an acceptance rate of 0.181. Notice that all criteria have been met: MCSE's are sufficiently small, ESS's are sufficiently large, and stationarity was estimated beginning with the first iteration. Since the algorithm was RWM, the Markov property holds, so let's look at some plots.

7. Plotting Output

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 1: Plots of Marginal Posterior Samples




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, 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 \mathbf{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: 302.861
```

```
L-criterion: 22.177, S.L: 0.393
```

```
Records:
```

	y	Mean	SD	LB	Median	UB	PQ	Discrep
1	4.174387	4.176	0.195	3.803	4.174	4.550	0.498	0.000
2	5.361292	5.296	0.392	4.540	5.300	6.057	0.438	0.028
3	6.089045	5.249	0.533	4.223	5.262	6.345	0.060	2.483
4	5.298317	5.154	0.323	4.534	5.165	5.818	0.311	0.200
5	4.406719	4.071	0.252	3.573	4.081	4.562	0.084	1.777
6	2.197225	3.802	0.196	3.437	3.803	4.192	1.000	66.671
7	5.010635	4.546	0.177	4.196	4.546	4.916	0.008	6.891

8	1.609438	3.860	0.197	3.487	3.860	4.247	1.000	130.549
9	4.343805	4.235	0.227	3.810	4.237	4.652	0.312	0.227
10	4.812184	4.722	0.222	4.290	4.722	5.171	0.332	0.165
11	4.189655	4.409	0.202	4.010	4.413	4.809	0.876	1.177
12	4.919981	4.532	0.175	4.177	4.529	4.874	0.015	4.926
13	4.753590	4.377	0.176	4.037	4.367	4.725	0.016	4.609
14	4.127134	4.166	0.166	3.836	4.168	4.475	0.595	0.055
15	3.713572	4.094	0.193	3.733	4.098	4.480	0.980	3.886
16	4.672829	4.406	0.223	3.981	4.408	4.842	0.105	1.434
17	6.930495	7.202	0.512	6.184	7.213	8.221	0.704	0.282
18	5.068904	4.809	0.239	4.328	4.814	5.296	0.129	1.191
19	6.775366	6.316	0.490	5.333	6.322	7.302	0.169	0.879
20	6.553933	7.220	0.479	6.240	7.234	8.128	0.910	1.930
21	4.890349	5.382	0.350	4.669	5.371	6.048	0.921	1.971
22	4.442651	4.265	0.267	3.719	4.275	4.778	0.234	0.444
23	2.833213	3.067	0.497	2.103	3.082	3.998	0.684	0.222
24	4.787492	4.935	0.255	4.390	4.940	5.461	0.735	0.335
25	6.933423	7.247	0.605	6.065	7.251	8.437	0.695	0.269
26	6.180017	6.044	0.619	4.835	6.034	7.215	0.415	0.048
27	5.652489	5.334	0.305	4.771	5.328	5.950	0.149	1.091
28	5.429346	4.470	0.206	4.040	4.479	4.888	0.000	21.626
29	5.634790	5.561	0.676	4.170	5.596	6.921	0.475	0.012
30	4.262680	4.064	0.204	3.681	4.062	4.454	0.156	0.952
31	3.891820	4.070	0.244	3.590	4.073	4.541	0.774	0.534
32	6.613384	6.607	0.400	5.792	6.607	7.389	0.492	0.000
33	4.919981	4.406	0.177	4.061	4.395	4.745	0.002	8.405
34	6.541030	6.465	0.528	5.426	6.470	7.511	0.442	0.021
35	6.345636	6.452	0.501	5.505	6.447	7.462	0.592	0.045
36	3.737670	4.054	0.270	3.517	4.051	4.564	0.874	1.376
37	7.356280	7.926	0.600	6.719	7.915	9.069	0.832	0.901
38	5.739793	4.765	0.169	4.414	4.764	5.087	0.000	33.114
39	5.517453	5.149	0.252	4.651	5.148	5.647	0.074	2.135

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` that 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.




Figure 7: Posterior Predictive Plots

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 \mathbf{y} and \mathbf{y}^{rep} , providing a criterion to be minimized (Laud and Ibrahim 1995).
- The last part of the summarized output reports \mathbf{y} , information about the distribution of \mathbf{yhat} , and the predictive quantile (PQ). The mean prediction of $\mathbf{y}[1]$, or \mathbf{y}_1^{rep} , given the model and data, is 4.176. Most importantly, $PQ[1]$ is 0.498, indicating that 49.8% of the time, $\mathbf{yhat}[1,]$ was greater than $\mathbf{y}[1]$, or that $\mathbf{y}[1]$ is close to the mean of $\mathbf{yhat}[1,]$. Contrast this with the 6th record, where $\mathbf{y}[6]=2.197$ and $PQ[6]=1$. Therefore, $\mathbf{yhat}[6,]$ was not a good replication of $\mathbf{y}[6]$, because the distribution of $\mathbf{yhat}[6,]$ is always greater than $\mathbf{y}[6]$. While $\mathbf{y}[1]$ is within the 95% probability interval of $\mathbf{yhat}[1,]$, the 95% probability interval of $\mathbf{yhat}[6,]$ is above $\mathbf{y}[6]$ 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)
```

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)`.
- Once a model has been specified in the `Model` function, it may be tempting to specify a large number of iterations and thinning in the `LaplacesDemon` function, and simply let the model update a long time, hoping for convergence. Instead, it is wise to begin with few iterations such as `Iterations=20`, set `Adaptive=0` (preventing adaptation), and set `Thinning=1`. User-error in specifying the `Model` function will be frustrating otherwise.
- As model complexity increases, the number of parameters increases, and as initial values are further from high-probability regions, the initial acceptance rate may be very low. If the previous general suggestion was successful, but the acceptance rate was zero, then update the model again, but for more iterations. The goal here is to verify that proposals are accepted without problems before attempting an “actual” model update.
- 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 was allowed to begin at the 900th iteration (`Adaptive=900`), but also occurred with `Periodicity=10`, so every 10th iteration, adaptation occurred. It is also wise to use delayed rejection to assist with the acceptance rate when the algorithm may begin far from its solution, so set `DR=1`.
- 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 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, and one possible sign of convergence will be when the deviance and the joint posterior seem to become stationary or no longer show a trend.
- 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 the AM or DRAM algorithm 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, continue Laplace's Demon with `Adaptive=0` so that adaptation will not occur. 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).

- Disagreement exists in MCMC literature as to whether to update one, long chain (Geyer 1992), 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, each beginning with different initial values, until multiple output objects of class `demonoid` exist with stationary samples, if time allows.

10. Independence and Observability

For the user, one set of advantages of Laplace’s Demon compared to many other available methods is that it 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 is performed completely within base R (though of course the **LaplacesDemon** package is required). A goal is to provide a complete, Bayesian environment. From personal experience, I’ve used multiple packages to achieve goals before, and have been trapped when one of those packages failed to keep pace with other changes.

Common Bayesian probability distributions (such as Dirichlet, multivariate normal, Wishart, and others, as well as truncated forms of distributions) have been included in **LaplacesDemon** so the user does not have to load numerous R packages. All functions in Laplace’s Demon are written entirely in R, so the user can easily observe or manipulate the algorithm or functions. For example, to print the code for **LaplacesDemon** to the R console, simply enter:

```
> LaplacesDemon
```

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 are 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). In the **LaplacesDemon** package, Laplace Approximation may iterate faster or slower than MCMC, so this is not the fastest possible implementation of Laplace Approximation. 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, though it currently works best with few parameters.

The `LaplaceApproximation` function may be called by the user before using `LaplacesDemon`, or `LaplacesDemon` may call this function if all initial values are zero. Chasing convergence with `LaplaceApproximation` may be time-consuming and unimportant. The goal, instead, is to improve the logarithm of the unnormalized joint posterior density so that it is easier for the `LaplacesDemon` function to begin updating the parameters in search of the target distributions. This can be difficult when the initial values are in low-probability regions, and can cause unreasonably low acceptance rates.

`LaplaceApproximation` seeks a global maximum of the logarithm of the unnormalized joint posterior density by taking steps proportional to an adaptive scale of the approximate gradient. This portion of the `LaplaceApproximation` function uses a conjugate gradient or gradient ascent algorithm, where gradient ascent is called a gradient descent or steepest descent algorithm elsewhere for minimization problems. Laplace’s Demon uses the `LaplaceApproximation` algorithm to optimize initial values, estimate covariance, and save time for the user, though it is used only when sample size is at least five times the number of parameters or initial values.

This algorithm assumes 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.

With adaptive gradient ascent, at 10 evenly-spaced 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. Conjugate gradient should give superior results, except in very large dimensions.

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. 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. 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. 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. 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 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 multi-modal target distributions, which other

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.

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.

Multiple parameters usually exist, and therefore correlations may occur between the parameters. All 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, 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 optimal acceptance rate for an infinite number of distributions that are IID and Gaussian is 0.234 or 23.4%.

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 (see below), 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. DRM is rarely suggested by Laplace's Demon, though the combination of DRM and AM, called DRAM (see below), is suggested frequently.

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.

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 also optimize the acceptance rate (which becomes difficult as dimensionality increases), minimize autocorrelation, or optimize a scale factor. Laplace’s Demon uses a variation of the Adaptive Metropolis (AM) algorithm of Haario *et al.* (2001).

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, and it is hoped that these algorithms will be included in Laplace’s Demon in the future.

Haario *et al.* (2001) tested their algorithm with up to 200 dimensions or parameters. It has been tested in Laplace’s Demon with 2,600 parameters, so it is capable of large-scale Bayesian inference. 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 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

¹⁰Haario, Saksman, and Tamminen (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.

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

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 and 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 "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.

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.

11.4. 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 stationary 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. Here, it is considered important enough to be one of five main criteria to appease Laplace's Demon. It is often recommended that one of several competing batch methods should be used to estimate MCSE, arguing that the simple method ($\text{MCSE} = \sigma/\sqrt{m}$) is biased and reports less error (where m is the ESS). I have calculated both the simple method and non-overlapping batch MCSE's on a wide range of applied models, and noted just as many cases of the simple method producing higher MCSE's as lower MCSE's. As far as Laplace's Demon is concerned, the simple method is used to estimate MCSE, but it is open to debate.

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**, **MCMC-pack**, and **UMACS**) and SAS. Many other R packages use MCMC, but are not intended as general-purpose MCMC software. Hopefully I have not overlooked any general-purpose MCMC packages in R.

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 sampling. 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.

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 textbooks in several fields have been written that are full of WinBUGS examples.

The four 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 website¹². 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 2007), and uses C language for significantly faster sampling. Metropolis-within-Gibbs is not as generalizable as adaptive MCMC. Otherwise, if the user wishes to see the code of the **AMCMC** sampler, then the user must also be familiar with C language.

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 adaptive algorithms.

The **MCMCpack** package (Martin, Quinn, and Park 2011) 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, though the algorithm is evaluated in `optim`, which has not performed well in my testing of Laplace Approximations.

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

¹²AMCMC is available from J. S. Rosenthal's website at <http://www.probability.ca/amcmc/>

avoids ‘for loops’ and `apply` functions wherever possible¹³. 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 `Adaptive=2`, `DR=0`, `Periodicity=10`, and the initial value for each β set to 0.1 to bypass Laplace Approximation. This took 0.98 minutes with Laplace’s Demon, according to a simple, linear regression¹⁴. It was nowhere near convergence, but updating the same model with the same data for 1,000 iterations took 45.55 minutes in JAGS.

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 converge¹⁵.

However, with small data sets, other MCMC software (**AMCMC** is a good example) can be faster than Laplace’s Demon, if it is programmed 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 in **C** as well.

When the non-adaptive algorithm updates in Laplace’s Demon, the expected speed of an iteration should not differ depending on how many iterations it has previously updated. However, the adaptive algorithm will slow as iterations are updated, because each time it adapts, it is adapting to the covariance of the entire history of the chains. As the history increases, the calculations take longer to complete, and the expected speed of an adaptive iteration decreases, compared to earlier adaptive iterations. If time is of the essence and the algorithm needs to be adaptive, then it may be best to make multiple, shorter updates in place of one, longer update.

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 laplacesdemon@statisticat.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

¹³However, when ‘for loops’ or `apply` functions must be used, Laplace’s Demon is typically slower than BUGS.

¹⁴These updates were performed on a 2010 System76 Pangolin Performance laptop with 64-bit Debian Linux and 8GB RAM.

¹⁵To continue this example, JAGS may be *guessed* to take 20,000 iterations or 15.18 hours, and **LaplacesDemon** may take 400,000 iterations or 6.53 hours, and also have less autocorrelation in the chains due to more thinning.

- Inference in Belief Networks with Continuous Variables.” In R̃Mantaras, D̃Poole (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 Chance. 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.
- Crawley M (2007). *The R Book*. John Wiley & Sons Ltd, West Sussex, England.
- Gelfand A (1996). “Model Determination Using Sampling Based Methods.” In W̃Gilks, S̃Richardson, D̃Spiegelhalter (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.
- 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>.
- Gilks W, Roberts G (1996). “Strategies for Improving MCMC.” In W̃Gilks, S̃Richardson, D̃Spiegelhalter (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 (2011). *LaplacesDemon: Software for Bayesian Inference*. R package version 11.10.03, 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.
- 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 (2011). *MCMCpack: Markov chain Monte Carlo (MCMC) Package*. R package version 1.1-1, 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 (2011). *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>.
- Roberts G, Rosenthal J (2001). “Optimal Scaling for Various Metropolis-Hastings Algorithms.” *Statistical Science*, **16**, 351–367.
- Roberts G, Rosenthal J (2007). “Examples of 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, 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.

Affiliation:

Byron Hall

STATISTICAT, LLC

Farmington, CT

E-mail: laplacesdemon@statisticat.com

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