



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 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 the likelihood. Laplace’s Demon also attempts to assist the user by creating and offering R code, based on a previous model update, that can be copy/pasted and executed. Posterior predictive checks and many other features are included as well. Laplace’s Demon seeks to be generalizable and user-friendly to Bayesians, especially Laplacians.

Keywords: ~Adaptive, AM, Bayesian, Delayed Rejection, DR, DRAM, DRM, LaplacesDemon, Laplace’s Demon, Markov chain Monte Carlo, MCMC, Metropolis, 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.

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

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.

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 2010) 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) inference, 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:

```
> library(LaplacesDemon)
```

```
LaplacesDemon: Software for Bayesian Inference
Thank you for using LaplacesDemon.
```

2. Data

Laplace’s Demon requires data that is specified in a list. As an example, there is a data set called `demonsnacks` that is provided with the **LaplacesDemon** package. For no good reason, other than to provide an example, the log of `Calories` will be fit as an additive, linear function of the remaining variables. Since an intercept will be included, a vector of 1’s is inserted into design matrix **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] <- (X[, j] - mean(X[, j]))/(2 * sd(X[,
+   j]))
+ }
> parm.names <- rep(NA, J + 1)
> for (j in 1:J) {
+   parm.names[j] <- paste("beta[", j, "]", sep = "")
+ }
> parm.names[J + 1] <- "log.tau"
> MyData <- list(J = J, X = X, 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 precision, tau, must be included as well. The reason why it is called `log.tau` 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. Also, note that each predictor has been centered and scaled, as per Gelman (2008). The way in which centering and scaling is performed in this example is crude, where more elegant methods exist, but it is used here because it is easy for the reader to see how it is done.

3. Specifying a Model

Laplace’s Demon is capable with any Bayesian model for which the likelihood is specified². To use Laplace’s Demon, the user must specify a model. Let’s consider a linear regression model, which is often denoted as:

$$y \sim N(\mu, \sigma^2)$$

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

²Examples of numerous Bayesian models may be found in the “Examples” vignette that comes with the **LaplacesDemon** package.

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

For a Bayesian model, the notation for the residual variance, σ^2 , is often replaced with the residual precision, τ^{-1} . Prior probabilities are specified for β and τ :

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

$$\tau \sim \Gamma(0.001, 0.001)$$

Each of the J β parameters is assigned an uninformative prior probability distribution that is normally-distributed according to $\mu = 0$ and $\sigma^2 = 1000$, where the precision is $\tau = 0.001$. The large variance or small precision indicates a lot of uncertainty about each β , and is hence a uninformative distribution. The residual precision τ is gamma-distributed according to two parameters of its distribution: $\alpha = 0.001$ and $\beta = 0.001$.

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.mu <- rep(0, Data$J)
+   beta.tau <- rep(0.001, Data$J)
+   tau.alpha <- 0.001
+   tau.beta <- 0.001
+   beta <- parm[1:Data$J]
+   tau <- exp(parm[Data$J + 1])
+   beta.prior <- dnorm(beta, beta.mu, 1/sqrt(beta.tau),
+     log = TRUE)
+   tau.prior <- dgamma(tau, tau.alpha, tau.beta,
+     log = TRUE)
+   mu <- beta %*% t(Data$X)
+   LL <- sum(dnorm(Data$y, mu, 1/sqrt(tau), log = TRUE))
+   LP <- LL + sum(beta.prior) + tau.prior
+   Modelout <- list(LP = LP, Dev = -2 * LL, Monitor = c(tau,
+     mu[1]), yhat = mu, parm = parm)
+   return(Modelout)
+ }
```

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

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

where Θ is a set of parameters, y is the data, \propto means ‘proportional to’³, $p(\Theta|y)$ is the joint posterior density, $p(y|\Theta)$ is the likelihood, and $p(\Theta)$ is the set of prior densities.

³For 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’.

During each iteration in which Laplace's Demon is maximizing the log of the 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 log of the 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, y^{rep} (`yhat`) or replicates of 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) + tau.prior
Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(tau,mu[1]),
  yhat=mu, parm=parm)
return(Modelout)
```

The rest of the function specifies the prior parameters, parameters, log of the prior densities, and calculates the log-likelihood. The prior parameters specify the parameters for the prior distributions. Since design matrix \mathbf{X} has $J=10$ column vectors (including the intercept), there are 10 `beta` parameters and a `tau` parameter for residual precision, the inverse of the variance. Each of the J `beta` parameters will be distributed normally according to mean `beta.mu` and precision `beta.tau`, and the additional `tau` parameter will be gamma-distributed according to `tau.alpha` and `tau.beta`. Here are the specifications for the prior parameters:

```
beta.mu <- rep(0,Data$J)
beta.tau <- rep(1.0E-3,Data$J)
tau.alpha <- 1.0E-3
tau.beta <- 1.0E-3
```

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 `tau` is exponentiated will, again, be explained later.

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

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

```
beta.prior <- dnorm(beta, beta.mu, 1/sqrt(beta.tau), log=TRUE)
tau.prior <- dgamma(tau, tau.alpha, tau.beta, log=TRUE)
```

It is important to reparameterize all parameters to be real-valued. For example, a positive-only parameter such as variance should be allowed to range from $-\infty$ to ∞ , and be transformed in the `Model` function with the `exp` function, which will force it to positive values. A parameter θ that needs to be bounded in the model, such as in the interval $[1,5]$, can be transformed to that range with a logistic function, such as $1 + 4[\exp(\theta)/(\exp(\theta) + 1)]$. Alternatively, each parameter may be constrained in the `Model` function. 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 precision `tau` receives a gamma-distributed prior of the form:

$$\tau \sim \Gamma(0.001, 0.001)$$

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

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

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

Finally, everything is put together to calculate `LP`, the log of the joint posterior density. The expectation vector `mu` is the inner product (`%*%`) of the vector `beta` and the transposed design matrix, `t(Data$X)`. Expectation vector `mu`, vector `Data$y`, and scalar `tau` are used to estimate the sum of the log-likelihoods, where:

$$y \sim N(\mu, \tau^{-1})$$

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

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

```
mu <- beta %*% t(Data$X)
LL <- sum(dnorm(Data$y, mu, 1/sqrt(tau), log=TRUE))
LP <- LL + sum(beta.prior) + tau.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 can be estimated in Laplace's Demon, but each missing value must be specified as a parameter in the `Model` function so that an initial value is assigned.

4. Initial Values

Laplace's Demon requires a vector of initial values for the parameters. Each initial value is a user-specified starting point for a parameter. In this example, there are 11 parameters, and with no prior knowledge, it is a good idea to set them equal to zero or use a random function. The first 10 parameters, the `beta` parameters, have been set equal to zero, and the remaining parameter, `log.tau`, has been set equal to `log(1)`, which is equal to zero. This visually reminds me that I am working with the log of `tau`, rather than `tau`, 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 demon⁴:

```
> 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)
```

Laplace's Demon was called on Sun Jan 30 14:39:24 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.

⁴Demonic 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'.

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 four MCMC algorithms (which are explained in section 11). The above example declares the Delayed Rejection Adaptive Metropolis (DRAM) algorithm.

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 Jan 30 14:39:46 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 independent 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] 0.151
```

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: 0.151
 Adaptive: 900
 Algorithm: Delayed Rejection Adaptive Metropolis
 Covar: (NOT SHOWN HERE)
 DIC of all samples (Dbar): 87.038
 DIC of all samples (pD): 909.02
 DIC of all samples (DIC): 996.06
 DIC of stationary samples (Dbar): 85.098
 DIC of stationary samples (pD): 296.32
 DIC of stationary samples (DIC): 381.42
 DR: 1
 Iterations: 10000
 Minutes of run-time: 0.38
 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: 101
 Recommended Burn-In of Un-thinned Samples: 1010
 Recommended Thinning: 140
 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]	5.023845	0.33840	0.028311	142.869	4.8098768
beta[2]	-0.423014	0.40478	0.036956	119.970	-1.0665893
beta[3]	-0.468149	0.88428	0.088965	98.796	-2.2457468
beta[4]	-0.051459	0.68067	0.056668	144.277	-1.4197627
beta[5]	-0.405247	0.52220	0.048073	117.997	-1.3001923
beta[6]	-0.512253	0.31105	0.021180	215.681	-1.0988330
beta[7]	2.236277	0.60102	0.049644	146.568	1.1135842
beta[8]	0.641497	0.42790	0.030244	200.178	-0.2695962
beta[9]	-0.232757	0.56686	0.054812	106.955	-1.3177832
beta[10]	1.627831	0.77833	0.070468	121.995	-0.0048381
log.tau	0.645996	0.40456	0.044354	83.199	-0.0319729
Deviance	87.037593	42.63843	2.311360	340.304	73.5399438
Monitor	2.037325	0.67806	0.053828	158.678	0.8952463
Monitor	4.180418	0.34765	0.016009	471.591	3.7464816
	Median	UB			

beta[1]	5.053093	5.271651
beta[2]	-0.442104	0.330382
beta[3]	-0.503749	1.444549
beta[4]	-0.057801	1.323524
beta[5]	-0.442618	0.732194
beta[6]	-0.498716	0.038452
beta[7]	2.250847	3.369559
beta[8]	0.640800	1.494984
beta[9]	-0.207839	0.798964
beta[10]	1.667156	3.034656
log.tau	0.676993	1.241983
Deviance	81.700079	124.733676
Monitor	1.969744	3.484118
Monitor	4.179164	4.696071

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.047606	0.11401	0.0052038	480.021	4.82297	5.051445
beta[2]	-0.464420	0.33531	0.0201334	277.374	-1.07451	-0.463978
beta[3]	-0.484691	0.90504	0.0944059	91.905	-2.27525	-0.529352
beta[4]	-0.041958	0.69862	0.0595537	137.616	-1.43164	-0.032474
beta[5]	-0.419372	0.49431	0.0447920	121.788	-1.30078	-0.438338
beta[6]	-0.507743	0.30025	0.0217726	190.168	-1.06780	-0.495138
beta[7]	2.278773	0.53204	0.0373720	202.674	1.32130	2.258343
beta[8]	0.650700	0.42245	0.0312737	182.473	-0.25847	0.664909
beta[9]	-0.243283	0.57955	0.0583164	98.765	-1.40380	-0.227560
beta[10]	1.657062	0.74681	0.0646415	133.475	0.16624	1.701984
log.tau	0.673556	0.27459	0.0266237	106.372	0.10367	0.674788
Deviance	85.097775	24.34440	0.9043630	724.624	73.91767	81.799162
Monitor	2.051439	0.65477	0.0473069	191.572	1.04396	1.963429
Monitor	4.180156	0.22584	0.0104361	468.309	3.77354	4.172922

UB

beta[1]	5.273329
beta[2]	0.216787
beta[3]	1.467696
beta[4]	1.350325
beta[5]	0.644929
beta[6]	0.034743
beta[7]	3.418600
beta[8]	1.491993
beta[9]	0.833338
beta[10]	3.036424
log.tau	1.245165
Deviance	110.998763
Monitor	3.503541
Monitor	4.622298

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 0.38 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: 0.151
Adaptive: 900
Algorithm: Delayed Rejection Adaptive Metropolis
Covar: (NOT SHOWN HERE)
DIC of all samples (Dbar): 87.038
DIC of all samples (pD): 909.02
DIC of all samples (DIC): 996.06
DIC of stationary samples (Dbar): 85.098
DIC of stationary samples (pD): 296.32
DIC of stationary samples (DIC): 381.42
DR: 1
Iterations: 10000
Minutes of run-time: 0.38
```

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: 101
 Recommended Burn-In of Un-thinned Samples: 1010
 Recommended Thinning: 140
 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]	5.023845	0.33840	0.028311	142.869	4.8098768
beta[2]	-0.423014	0.40478	0.036956	119.970	-1.0665893
beta[3]	-0.468149	0.88428	0.088965	98.796	-2.2457468
beta[4]	-0.051459	0.68067	0.056668	144.277	-1.4197627
beta[5]	-0.405247	0.52220	0.048073	117.997	-1.3001923
beta[6]	-0.512253	0.31105	0.021180	215.681	-1.0988330
beta[7]	2.236277	0.60102	0.049644	146.568	1.1135842
beta[8]	0.641497	0.42790	0.030244	200.178	-0.2695962
beta[9]	-0.232757	0.56686	0.054812	106.955	-1.3177832
beta[10]	1.627831	0.77833	0.070468	121.995	-0.0048381
log.tau	0.645996	0.40456	0.044354	83.199	-0.0319729
Deviance	87.037593	42.63843	2.311360	340.304	73.5399438
Monitor	2.037325	0.67806	0.053828	158.678	0.8952463
Monitor	4.180418	0.34765	0.016009	471.591	3.7464816

	Median	UB
beta[1]	5.053093	5.271651
beta[2]	-0.442104	0.330382
beta[3]	-0.503749	1.444549
beta[4]	-0.057801	1.323524
beta[5]	-0.442618	0.732194
beta[6]	-0.498716	0.038452
beta[7]	2.250847	3.369559
beta[8]	0.640800	1.494984
beta[9]	-0.207839	0.798964
beta[10]	1.667156	3.034656
log.tau	0.676993	1.241983
Deviance	81.700079	124.733676
Monitor	1.969744	3.484118
Monitor	4.179164	4.696071

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.047606	0.11401	0.0052038	480.021	4.82297	5.051445
beta[2]	-0.464420	0.33531	0.0201334	277.374	-1.07451	-0.463978
beta[3]	-0.484691	0.90504	0.0944059	91.905	-2.27525	-0.529352
beta[4]	-0.041958	0.69862	0.0595537	137.616	-1.43164	-0.032474
beta[5]	-0.419372	0.49431	0.0447920	121.788	-1.30078	-0.438338
beta[6]	-0.507743	0.30025	0.0217726	190.168	-1.06780	-0.495138
beta[7]	2.278773	0.53204	0.0373720	202.674	1.32130	2.258343
beta[8]	0.650700	0.42245	0.0312737	182.473	-0.25847	0.664909
beta[9]	-0.243283	0.57955	0.0583164	98.765	-1.40380	-0.227560
beta[10]	1.657062	0.74681	0.0646415	133.475	0.16624	1.701984
log.tau	0.673556	0.27459	0.0266237	106.372	0.10367	0.674788
Deviance	85.097775	24.34440	0.9043630	724.624	73.91767	81.799162
Monitor	2.051439	0.65477	0.0473069	191.572	1.04396	1.963429
Monitor	4.180156	0.22584	0.0104361	468.309	3.77354	4.172922
UB						
beta[1]	5.273329					
beta[2]	0.216787					
beta[3]	1.467696					
beta[4]	1.350325					
beta[5]	0.644929					
beta[6]	0.034743					
beta[7]	3.418600					
beta[8]	1.491993					
beta[9]	0.833338					
beta[10]	3.036424					
log.tau	1.245165					
Deviance	110.998763					
Monitor	3.503541					
Monitor	4.622298					

Demonic Suggestion

Due to the combination of the following conditions,

1. Delayed Rejection Adaptive Metropolis
2. The acceptance rate (0.151) is within the interval $[0.15, 0.5]$.
3. At least one target MCSE is $\geq 6.27\%$ of its marginal posterior standard deviation.
4. At least one target distribution has an effective sample size (ESS) less than 100.
5. Each target distribution became stationary by 101 iterations.

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

```
Initial.Values <- Fit$Posterior1[Fit$Thinned.Samples,]
Fit <- LaplacesDemon(Model, Data=MyData, Adaptive=0,
  Covar=Fit$Covar, DR=0, Initial.Values, Iterations=140000,
  Periodicity=0, Status=26316, Thinning=140)
```

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

Notice that since stationarity has been estimated beginning with the 101st 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=140000** and **Thinning=140**. For the sake of the example, the suggested R code will be run:

```
> Initial.Values <- Fit$Posterior1[Fit$Thinned.Samples,
+ ]
```

⁵While 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%].

```
> Fit <- LaplacesDemon(Model, Data = MyData, Adaptive = 0,
+   Covar = Fit$Covar, DR = 0, Initial.Values, Iterations = 140000,
+   Periodicity = 0, Status = Rec.Status, Thinning = 140)
```

Laplace's Demon was called on Sun Jan 30 14:40:09 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: 26316, Proposal: Multivariate

Iteration: 52632, Proposal: Multivariate

Iteration: 78948, Proposal: Multivariate

Iteration: 105264, Proposal: Multivariate

Iteration: 131580, Proposal: Multivariate

Assessing Stationarity

Assessing Thinning and ESS

Creating Summaries

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 = 140000,
  Periodicity = 0, Status = Rec.Status, Thinning = 140)
```

Acceptance Rate: 0.104

Adaptive: 140001

Algorithm: Random-Walk Metropolis

Covar: (NOT SHOWN HERE)

DIC of all samples (Dbar): 82.599

DIC of all samples (pD): 16.17

DIC of all samples (DIC): 98.769

DIC of stationary samples (Dbar): 82.599

DIC of stationary samples (pD): 16.17

DIC of stationary samples (DIC): 98.769

DR: 0
 Iterations: 140000
 Minutes of run-time: 1.37
 Model: (NOT SHOWN HERE)
 Monitor: (NOT SHOWN HERE)
 Parameters (Number of): 11
 Periodicity: 140001
 Posterior1: (NOT SHOWN HERE)
 Posterior2: (NOT SHOWN HERE)
 Recommended Burn-In of Thinned Samples: 1
 Recommended Burn-In of Un-thinned Samples: 140
 Recommended Thinning: 140
 Status is displayed every 26316 iterations
 Summary1: (SHOWN BELOW)
 Summary2: (SHOWN BELOW)
 Thinned Samples: 1000
 Thinning: 140

Summary of All Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.043050	0.11735	0.0037110	1000.00	4.82230	5.041210
beta[2]	-0.460943	0.39564	0.0125112	1000.00	-1.24428	-0.456113
beta[3]	-0.425834	0.95345	0.0305721	972.63	-2.20963	-0.463028
beta[4]	-0.080472	0.70178	0.0235502	887.99	-1.47775	-0.074824
beta[5]	-0.400099	0.52748	0.0190397	767.52	-1.47697	-0.378210
beta[6]	-0.486054	0.28826	0.0091157	1000.00	-1.03108	-0.480930
beta[7]	2.234295	0.57834	0.0182888	1000.00	1.15831	2.237885
beta[8]	0.616238	0.44232	0.0139875	1000.00	-0.26938	0.606442
beta[9]	-0.161214	0.60270	0.0190590	1000.00	-1.33288	-0.172599
beta[10]	1.617871	0.76159	0.0256503	881.57	0.10063	1.619918
log.tau	0.667224	0.26619	0.0088640	901.84	0.10291	0.667983
Deviance	82.599193	5.68681	0.2009995	800.47	73.63728	81.862022
Monitor	2.016305	0.52285	0.0178561	857.39	1.14732	1.952132
Monitor	4.173337	0.20585	0.0065096	1000.00	3.74971	4.185542
UB						
beta[1]	5.269494					
beta[2]	0.322310					
beta[3]	1.461353					
beta[4]	1.279225					
beta[5]	0.574134					
beta[6]	0.042421					
beta[7]	3.467258					
beta[8]	1.478068					
beta[9]	0.970311					
beta[10]	3.151475					
log.tau	1.169123					

Deviance 96.426988
 Monitor 3.219167
 Monitor 4.551336

Summary of Stationary Samples

	Mean	SD	MCSE	ESS	LB	Median
beta[1]	5.043050	0.11735	0.0037110	1000.00	4.82230	5.041210
beta[2]	-0.460943	0.39564	0.0125112	1000.00	-1.24428	-0.456113
beta[3]	-0.425834	0.95345	0.0305721	972.63	-2.20963	-0.463028
beta[4]	-0.080472	0.70178	0.0235502	887.99	-1.47775	-0.074824
beta[5]	-0.400099	0.52748	0.0190397	767.52	-1.47697	-0.378210
beta[6]	-0.486054	0.28826	0.0091157	1000.00	-1.03108	-0.480930
beta[7]	2.234295	0.57834	0.0182888	1000.00	1.15831	2.237885
beta[8]	0.616238	0.44232	0.0139875	1000.00	-0.26938	0.606442
beta[9]	-0.161214	0.60270	0.0190590	1000.00	-1.33288	-0.172599
beta[10]	1.617871	0.76159	0.0256503	881.57	0.10063	1.619918
log.tau	0.667224	0.26619	0.0088640	901.84	0.10291	0.667983
Deviance	82.599193	5.68681	0.2009995	800.47	73.63728	81.862022
Monitor	2.016305	0.52285	0.0178561	857.39	1.14732	1.952132
Monitor	4.173337	0.20585	0.0065096	1000.00	3.74971	4.185542

UB

beta[1]	5.269494
beta[2]	0.322310
beta[3]	1.461353
beta[4]	1.279225
beta[5]	0.574134
beta[6]	0.042421
beta[7]	3.467258
beta[8]	1.478068
beta[9]	0.970311
beta[10]	3.151475
log.tau	1.169123
Deviance	96.426988
Monitor	3.219167
Monitor	4.551336

Demonic Suggestion

Due to the combination of the following conditions,

1. Random-Walk Metropolis
2. The acceptance rate (0.104) is below 0.15.
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 1 iterations.

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

```
Initial.Values <- Fit$Posterior1[Fit$Thinned.Samples,]
Fit <- LaplacesDemon(Model, Data=MyData, Adaptive=0,
  Covar=Fit$Covar, DR=1, Initial.Values, Iterations=140000,
  Periodicity=1346, Status=102190, Thinning=140)
```

Laplace's Demon is finished consorting.

In 1.37 minutes, Laplace's Demon updated 140000 iterations, retaining every 140th iteration due to thinning, and reported an acceptance rate of 0.104, which is low enough to trigger a suggestion to continue updating. However, notice that all other criteria have been met: MCSE's are sufficiently small, ESS's are sufficiently large, and stationarity was estimated beginning with the first iteration.

Notice `DR=1` in the suggested R code. Together with `Adaptive=0`, this indicates that a Delayed Rejection Metropolis (DRM) algorithm is suggested. DRM is discussed in section 11, but here it is enough to state that in the case of a rejected proposal, the DRM algorithm will make a second proposal before simply moving on to the next iteration. This serves to increase the acceptance rate, and is obviously why Laplace's Demon suggests DRM.

The low acceptance rate suggests that the proposal distribution may not be optimal, and that the chains may not have mixed well. Even though Laplace's Demon has not been appeased due to the acceptance rate, everything else looks good. 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, Parms = Fit$Parameters)
```

There are three plots for each parameter, the deviance, and each monitored variable (which in this example are `tau` 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

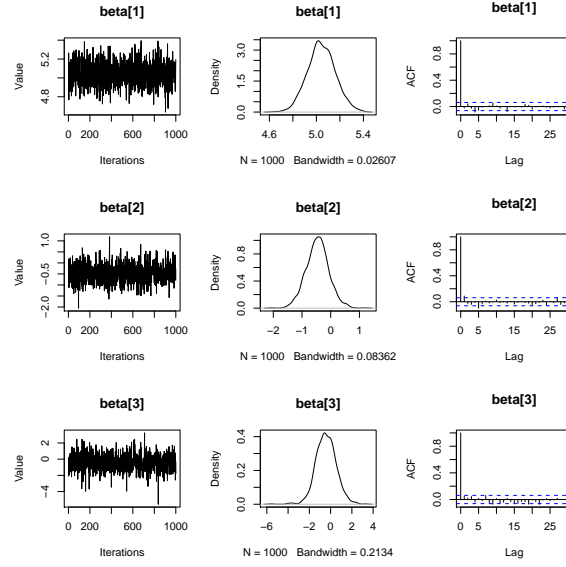


Figure 1: Plots of Marginal Posterior Samples

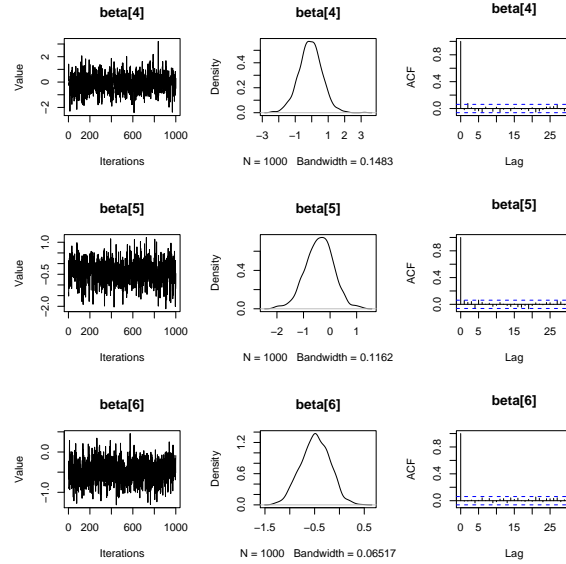


Figure 2: Plots of Marginal Posterior Samples

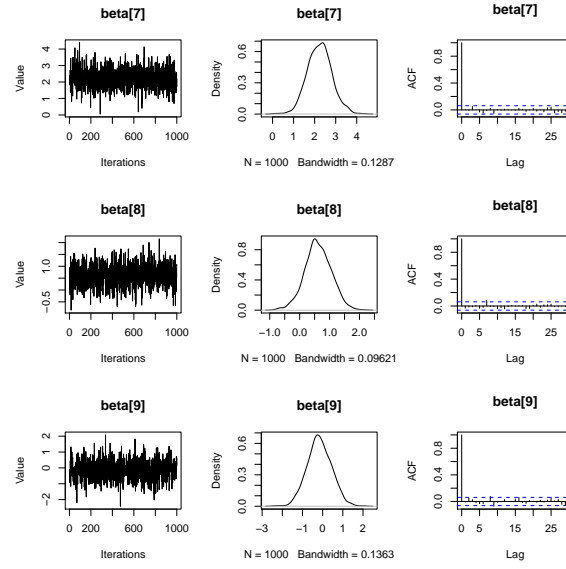


Figure 3: Plots of Marginal Posterior Samples

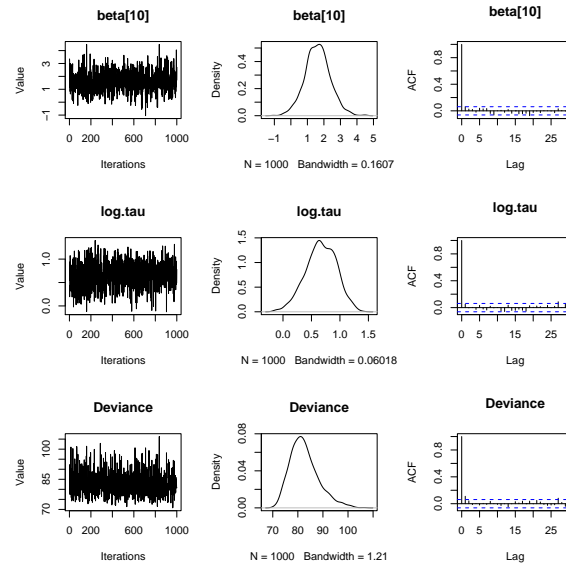


Figure 4: Plots of Marginal Posterior Samples

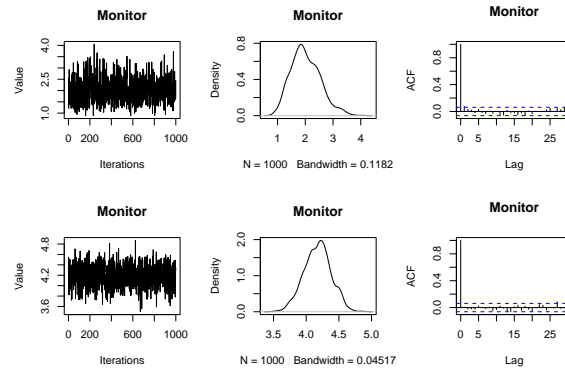


Figure 5: Plots of Marginal Posterior Samples

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.

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, marginal posterior samples may be used for inference.

When predicting the logarithm of y (Calories) with the `demonSnacks` data, 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.

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 y^{rep}) are estimated. If a new data set is supplied in `predict`, then new, unobserved instances of y (called y^{new}) are estimated. Note that with new data, a `y` vector must still be supplied, and if unknown, can be set to something sensible such as the mean of the `y` vector in the model.

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

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

```
> summary(Pred)
```

```
Concordance: 0.76923
```

```
Records:
```

	y	Mean	SD	LB	Median	UB	PQ
1	4.1744	4.1726	0.21037	3.7512	4.1839	4.5672	0.516
2	5.3613	5.2842	0.42346	4.4636	5.2753	6.0913	0.423
3	6.0890	5.2601	0.54282	4.2511	5.2740	6.3467	0.063
4	5.2983	5.1334	0.34303	4.4995	5.1269	5.7947	0.313
5	4.4067	4.0852	0.23872	3.6153	4.0839	4.5474	0.087
6	2.1972	3.8058	0.20690	3.4028	3.8042	4.2354	1.000
7	5.0106	4.5456	0.19153	4.1668	4.5442	4.9176	0.007
8	1.6094	3.8655	0.20400	3.4730	3.8638	4.2901	1.000
9	4.3438	4.2267	0.25582	3.7354	4.2200	4.7431	0.327
10	4.8122	4.7227	0.23010	4.2708	4.7184	5.1976	0.341
11	4.1897	4.4151	0.20038	4.0277	4.4152	4.8209	0.880
12	4.9200	4.5311	0.18160	4.1635	4.5303	4.9088	0.023
13	4.7536	4.3754	0.18619	4.0112	4.3762	4.7498	0.024
14	4.1271	4.1647	0.18529	3.8093	4.1597	4.5443	0.574
15	3.7136	4.0985	0.20079	3.7153	4.0972	4.5214	0.977
16	4.6728	4.3979	0.24908	3.9084	4.3914	4.8838	0.135
17	6.9305	7.1909	0.55017	6.0630	7.1960	8.2511	0.694
18	5.0689	4.8040	0.25662	4.3087	4.8049	5.3068	0.140
19	6.7754	6.3193	0.47844	5.3625	6.3187	7.2636	0.169
20	6.5539	7.2110	0.47059	6.2740	7.2252	8.1290	0.925
21	4.8903	5.3757	0.37621	4.5990	5.3771	6.0821	0.914
22	4.4427	4.2568	0.29632	3.6601	4.2583	4.8365	0.269
23	2.8332	3.0955	0.50240	2.1002	3.1132	4.0705	0.715
24	4.7875	4.9359	0.26554	4.3790	4.9382	5.4548	0.741
25	6.9334	7.2029	0.62457	6.0094	7.2076	8.4322	0.666
26	6.1800	6.0359	0.61201	4.7555	6.0338	7.1975	0.395
27	5.6525	5.3244	0.30417	4.7397	5.3197	5.9207	0.139
28	5.4293	4.4739	0.21926	4.0240	4.4680	4.9182	0.000
29	5.6348	5.5200	0.71131	4.0465	5.5185	6.8674	0.433
30	4.2627	4.0687	0.21434	3.6561	4.0705	4.5069	0.171
31	3.8918	4.0639	0.27065	3.5396	4.0697	4.5707	0.738
32	6.6134	6.6059	0.38030	5.8264	6.6151	7.3139	0.503
33	4.9200	4.4037	0.19003	4.0329	4.4047	4.7860	0.009
34	6.5410	6.4245	0.48658	5.4747	6.4296	7.4401	0.393
35	6.3456	6.3978	0.49137	5.4727	6.3882	7.4039	0.532
36	3.7377	4.0645	0.24543	3.5970	4.0628	4.5511	0.907
37	7.3563	7.9146	0.63457	6.6721	7.8958	9.2386	0.824
38	5.7398	4.7681	0.17696	4.4250	4.7674	5.1379	0.000
39	5.5175	5.1366	0.27009	4.5760	5.1404	5.6466	0.078

The `summary.demonoid.ppc` function returns a list with 2 components:

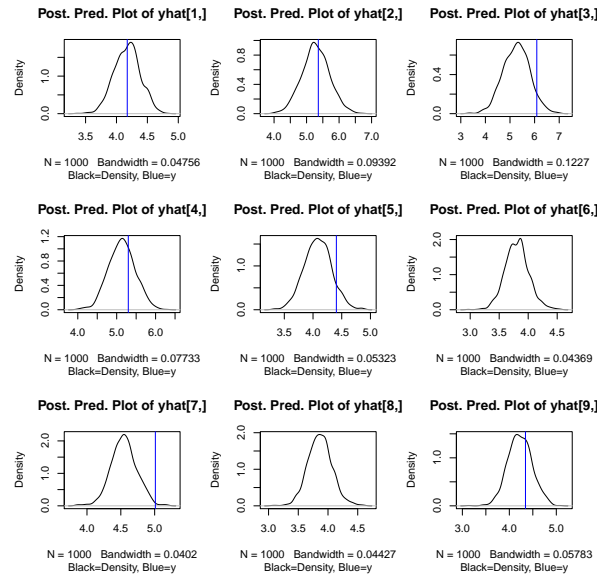


Figure 6: Posterior Predictive Plots

- **Concordance** is the predictive concordance of [Gelfand \(1996\)](#), that indicates the percentage of y that was within the 95% probability interval of \hat{y} . 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 \hat{y} . 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.
- The last part of the summarized output reports y , information about the distribution of \hat{y} , and the predictive quantile (PQ). The mean prediction of $y[1]$, or y_1^{rep} , given the model and data, is 4.173. Most importantly, $PQ[1]$ is 0.516, indicating that 51.6% of the time, $\hat{y}[1,]$ was greater than $y[1]$, or that $y[1]$ is close to the mean of $\hat{y}[1,]$. Contrast this with the 6th record, where $y[6]=2.197$ and $PQ[6]=1$. Therefore, $\hat{y}[6,]$ was not a good replication of $y[6]$, because the distribution of $\hat{y}[6,]$ is always greater than $y[6]$. While $y[1]$ is within the 95% probability interval of $\hat{y}[1,]$, the 95% probability interval of $\hat{y}[6,]$ is above $y[6]$ 100% of the time, indicating a strong discrepancy between the model and data, in this case.

The last component of this summary may be viewed graphically as well. Rather than observing plots for each of 39 records or rows, only the first 9 will be shown here:

```
> plot(Pred, Rows = c(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)).`
- Do not forget to reparameterize any bounded parameters in the `Model` function to be real-valued.
- 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 changing the initial values. 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 algorithm may take longer than usual, and 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. 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-plot of the deviance. If it is decreasing run after run, then the model is continuously fitting better and better, and one possible sign of convergence will be when the deviance seems to become stationary or no longer shows a trend.

- 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.
- 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). 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) 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

Laplace's Demon 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 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. 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.

11.1. 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.

11.2. 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 proposal⁶ 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 log of the 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 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

⁶Laplace'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.

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, then Laplace's Demon temporarily resorts to independent proposals by estimating univariate variances, 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%.

11.3. 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.

11.4. 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 state. 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, Saksman, and Tamminen (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 $2.4/\sqrt{d}$ 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: $2.381204/\sqrt{d}$, even though Haario *et al.* (2001) use a different form: $2.4^2/d$. 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, 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, it defaults to independent adaptive proposals. Although independent adaptive proposals should take longer 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:

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

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

For independent adaptive tuning, the formula across K parameters and t iterations is:

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

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 independent 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, it uses the second suggested trick of shrinking the proposal as long as the acceptance rate is less than 5%. 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 99%. 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.

11.5. 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.

11.6. Afterward

Once the model is updated, 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. A suggestion for BUGS and JAGS would be to attempt Gibbs sampling and abandon it if correlations are too high. 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*. 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. For example, BUGS automatically handles missing values in the dependent variable, where Laplace’s Demon requires specifications for each one in the `Model` function. BUGS also allows the user to specify the model graphically as a directed acyclic graph (DAG) in Doodle BUGS. 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 website⁷. 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 2010) 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 avoids 'for loops' wherever possible. 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`, and `Periodicity=10` in 1.55 minutes by Laplace's Demon, according to a simple, linear regression⁸. 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

⁷**AMCMC** is available from J. S. Rosenthal's website at <http://www.probability.ca/amcmc/>

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

sets better, and it may estimate each iteration faster, but it may also take more iterations to converge.

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**, or **C++**. 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. Laplace’s Demon Predicts Its Own Future

Following are some predictions of the future of Laplace’s Demon:

1. Additional MCMC algorithms will be explored and considered for inclusion.
2. The MCMC algorithms will be attempted to be coded in (C++) for faster sampling, but will remain available in **R** code. The user will have the option to run the algorithm in **R** or **C++**, probably with an additional command such as `C=TRUE` in the **LaplacesDemon** function.
3. An approximate, deterministic method such as Laplace Approximation will be explored to assist in getting initial values reasonably close to their target distributions, hopefully saving the user a lot of time.
4. The “Examples” vignette will grow as numerous examples of methods are included.
5. Posterior predictive checks will be attempted to be expanded to accommodate a wider variety of methods, and more checks and plotting options will be included.

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

References

- 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.

- Crawley M (2007). *The R Book*. John Wiley & Sons Ltd, West Sussex, England.
- Gelfand A (1996). “Model Determination Using Sampling Based Methods.” In WˆGilks, SˆRichardson, DˆSpiegelhalter (eds.), *Markov Chain Monte Carlo in Practice*, pp. 145–161. Chapman & Hall, Boca Raton, FL.
- Gelman A (2008). “Scaling Regression Inputs by Dividing by Two Standard Deviations.” *Statistics in Medicine*, **27**, 2865–2873.
- Gelman A, Carlin J, Stern H, Rubin D (2004). *Bayesian Data Analysis*. 2nd edition. Chapman & Hall, Boca Raton, FL.
- Gelman A, Meng X, Stern H (1996a). “Posterior Predictive Assessment of Model Fitness via Realized Discrepancies.” *Statistica Sinica*, **6**, 773–807.
- Gelman A, Roberts G, Gilks W (1996b). “Efficient Metropolis Jumping Rules.” *Bayesian Statistics*, **5**, 599–608.
- Gelman A, Rubin D (1992). “Inference from Iterative Simulation Using Multiple Sequences.” *Statistical Science*, **7**(4), 457–472.
- Geweke J (1992). “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments.” *Bayesian Statistics*, **4**, 1–31.
- 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://www.R-project.org/package=mcmc>.
- 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.
- Hall B (2011). *LaplacesDemon: Software for Bayesian Inference*. R package version 11.02.01, URL <http://www.R-project.org/package=LaplacesDemon>.
- 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).
- 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 (2010). **MCMCpack**: *Markov chain Monte Carlo (MCMC) Package*. R package version 1.0-8, URL <http://www.R-project.org/package=MCMCpack>.
- 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 (2010). *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>.

Affiliation:

Byron Hall
 STATISTICAT, LLC
 Farmington, CT
 E-mail: statisticat@gmail.com
 URL: <http://www.statisticat.com/laplacesdemon.html>