



LaplacesDemon Examples

Byron Hall
STATISTICAT, LLC

Abstract

The **LaplacesDemon** package in R enables Bayesian inference with any Bayesian model, provided the user specifies the likelihood. This vignette is a compendium of examples of how to specify different model forms.

Keywords: ~Bayesian, Bayesian Inference, Laplace’s Demon, LaplacesDemon, R, STATISTICAT.

LaplacesDemon (Hall 2011), usually referred to as Laplace’s Demon, is an R package that is available on CRAN (R Development Core Team 2010). A formal introduction to Laplace’s Demon is provided in an accompanying vignette entitled “**LaplacesDemon Tutorial**”, and an introduction to Bayesian inference is provided in the “Bayesian Inference” vignette.

The purpose of this document is to provide users of the **LaplacesDemon** package with examples of a variety of Bayesian methods. To conserve space, the examples are not worked out in detail, and only the minimum of necessary materials is provided for using the various methodologies. Necessary materials include the form expressed in notation, data (which is often simulated), initial values, and the `Model` function. This vignette will grow over time as examples of more methods become included. Contributed examples are welcome. Please send contributed examples in a similar format in an email to statisticat@gmail.com for review and testing. All accepted contributions are, of course, credited.

Contents

- Autoregression, AR(1) [1](#)
- Binary Logit [2](#)
- Dynamic Linear Model (DLM) [3](#)
- Normal, Multilevel [4](#)
- Linear Regression [5](#)
- Linear Regression, Laplace-Distributed [6](#)

- Poisson Regression [7](#)

1. Autoregression, AR(1)

1.1. Form

$$\begin{aligned}
 y_t &\sim N(\mu_{t-1}, \tau^{-1}), \quad t = 2, \dots, T \\
 \mu_t &= \alpha + \phi y_t, \quad t = 1, \dots, T \\
 \alpha &\sim N(0, 1000) \\
 \phi &\sim N(0, 1000) \\
 \tau &\sim \Gamma(0.001, 0.001)
 \end{aligned}$$

1.2. Data

```

T <- 100
y <- rep(0, T)
y[1] <- 0
for (t in 2:T) {y[t] <- y[t-1] + rnorm(1, 0, 0.1)}
parm.names <- c("alpha", "phi", "log.tau")
MyData <- list(T=T, parm.names=parm.names, y=y)

```

1.3. Initial Values

```
Initial.Values <- c(rep(0, 2), log(1))
```

1.4. Model

```

Model <- function(parm, Data)
{
  ### Prior Parameters
  alpha.mu <- 0; alpha.tau <- 1.0E-3
  phi.mu <- 0; phi.tau <- 1.0E-3
  tau.alpha <- 1.0E-3; tau.beta <- 1.0E-3
  ### Parameters
  alpha <- parm[1]; phi <- parm[2]; tau <- exp(parm[3])
  ### Log(Prior Densities)
  alpha.prior <- dnorm(alpha, alpha.mu, 1/sqrt(alpha.tau), log=TRUE)
  phi.prior <- dnorm(phi, phi.mu, 1/sqrt(phi.tau), log=TRUE)
  tau.prior <- dgamma(tau, tau.alpha, tau.beta, log=TRUE)
  ### Log-Likelihood
  mu <- alpha + phi*y
}

```

```

LL <- sum(dnorm(y[2:T], mu[1:(T-1)], 1/sqrt(tau), log=TRUE))
### Log-Posterior
LP <- LL + alpha.prior + phi.prior + tau.prior
Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(tau), yhat=mu,
    parm=parm)
return(Modelout)
}

```

2. Binary Logit

2.1. Form

$$y \sim Bern(\eta)$$

$$\eta = \log[1 + \exp(\mu)]$$

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

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

2.2. Data

```

data(demonsnacks)
N <- NROW(demonsnacks)
J <- 3
y <- ifelse(demonsnacks$Calories <= 137, 0, 1)
X <- cbind(1, as.matrix(demonsnacks[,c(7,8)]))
for (j in 2:J) {X[,j] <- (X[,j] - mean(X[,j])) / (2*sd(X[,j]))}
parm.names <- rep(NA, J)
for (j in 1:J) {parm.names[j] <- paste("beta[", j, "]", sep="")}
MyData <- list(J=J, X=X, parm.names=parm.names, y=y)

```

2.3. Initial Values

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

2.4. Model

```

Model <- function(parm, MyData)
{
  ### Prior Parameters
  beta.mu <- rep(0, J)
  beta.tau <- rep(1.0E-3, J)
  ### Parameters

```

```

beta <- parm[1:J]
### Log(Prior Densities)
beta.prior <- rep(0,j)
for (j in 1:J)
  {beta.prior[j] <- dnorm(beta[j], beta.mu[j],
    1/sqrt(beta.tau[j]), log=TRUE)}
### Log-Posterior
mu <- beta %*% t(X)
eta <- log(1 + exp(mu))
eta[mu>700] <- mu[mu>700] # Overflow trick
### Log-Likelihood
LL <- sum(y*mu - eta)
### Log-Posterior
LP <- LL + sum(beta.prior)
Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(eta[1],mu[1]),
  yhat=eta, parm=parm)
return(Modelout)
}

```

3. Dynamic Linear Model (DLM)

The data is presented so that the time-series is subdivided into three sections: modeled ($t = 1, \dots, T_m$), one-step ahead forecast ($t = T_m + 1$), and future forecast [$t = (T_m + 2), \dots, T$].

3.1. Form

$$\begin{aligned}
y_t &\sim N(\mu_t, \tau_V^{-1}), \quad t = 1, \dots, T_m \\
y_t^{new} &\sim N(\mu_t, \tau_V^{-1}), \quad t = (T_m + 1), \dots, T \\
\mu_t &= \alpha + x_t \beta_t, \quad t = 1, \dots, T \\
\alpha &\sim N(0, 1000) \\
\beta_1 &\sim N(0, 1000) \\
\beta_t &\sim N(\beta_{t-1}, \tau_W^{-1}), \quad t = 2, \dots, T \\
\tau_V &\sim \Gamma(0.001, 0.001) \\
\tau_W &\sim \Gamma(0.001, 0.001)
\end{aligned}$$

3.2. Data

```

T <- 20
T.m <- 14
beta.orig <- x <- rep(0,T)
for (t in 2:T) {

```

```

beta.orig[t] <- beta.orig[t-1] + rnorm(1,0,0.1)
x[t] <- x[t-1] + rnorm(1,0,0.1)}
y <- 10 + beta.orig*x + rnorm(T,0,0.1)
y[(T.m+2):T] <- NA
parm.names <- rep(NA, T+3)
parm.names[1] <- "alpha"
for (i in 1:T) {parm.names[i+1] <- paste("beta[", i, "]", sep="")}
parm.names[(T+2):(T+3)] <- c("log.beta.w.tau","log.v.tau")
MyData <- list(T=T, T.m=T.m, parm.names=parm.names, x=x, y=y)

```

3.3. Initial Values

```
Initial.Values <- c(rep(0,T+3))
```

3.4. Model

```

Model <- function(parm, Data)
{
  ### Parameters
  alpha <- parm[1]
  beta <- rep(0,T); beta <- parm[2:(T+1)]
  beta.w.tau <- exp(parm[T+2])
  v.tau <- exp(parm[T+3])
  ### Log(Prior Densities)
  alpha.prior <- dnorm(alpha, 0, 1/sqrt(1.0E-3), log=TRUE)
  beta.prior <- rep(0,T)
  beta.prior[1] <- dnorm(beta[1], 0, 1/sqrt(1.0E-3), log=TRUE)
  for (t in 2:T) {
    beta.prior[t] <- dnorm(beta[t], beta[t-1], 1/sqrt(beta.w.tau),
      log=TRUE)}
  beta.w.tau.prior <- dgamma(beta.w.tau, 0.001, 0.001, log=TRUE)
  v.tau.prior <- dgamma(v.tau, 1.0E-3, 1.0E-3, log=TRUE)
  ### Log-Likelihood
  mu <- alpha + beta*x
  LL <- sum(dnorm(y[1:T.m], mu[1:T.m], 1/sqrt(v.tau), log=TRUE))
  ### Log-Posterior
  LP <- LL + alpha.prior + sum(beta.prior) + beta.w.tau.prior +
    v.tau.prior
  Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(mu[(T.m+1):T]),
    yhat=mu, parm=parm)
  return(Modelout)
}

```

4. Normal, Multilevel

This is Gelman's school example ([Gelman, Carlin, Stern, and Rubin 2004](#)). Note that

Laplace's Demon is much slower to converge compared to this example that uses the **R2WinBUGS** package (Gelman 2009), an R package on CRAN. However, also note that Laplace's Demon (eventually) provides a better answer (higher ESS, lower DIC, etc.).

4.1. Form

$$\begin{aligned}y_j &\sim N(\theta_j, \tau_j^{-1}) \\ \theta_j &\sim N(\theta_\mu, \theta_\tau^{-1}) \\ \theta_\mu &\sim N(0, 1000) \\ \theta_\tau &\sim \Gamma(0.001, 0.001) \\ \tau_j &= sd^{-2}\end{aligned}$$

4.2. Data

```
J <- 8
y <- c(28.4, 7.9, -2.8, 6.8, -0.6, 0.6, 18.0, 12.2)
sd <- c(14.9, 10.2, 16.3, 11.0, 9.4, 11.4, 10.4, 17.6)
parm.names <- 2*J+2
for (j in 1:J) {parm.names[j] <- paste("theta[", j, "]", sep="")}
parm.names[J+1] <- paste("theta.mu[", J, "]", sep="")
parm.names[J+2] <- paste("log.theta.sigma[", J, "]", sep="")
MyData <- list(J=J, parm.names=parm.names, sd=sd, y=y)
```

4.3. Initial Values

```
Initial.Values <- rep(0, J+2)
```

4.4. Model

```
Model <- function(parm, MyData)
{
  ### Hyperprior Parameters
  theta.mu.mu <- 0
  theta.mu.tau <- 1.0E-3
  ### Prior Parameters
  theta.mu <- parm[J+1]
  theta.sigma <- exp(parm[J+2])
  tau.alpha <- 1.0E-3
  tau.beta <- 1.0E-3
  ### Parameters
  tau <- theta <- rep(0, J)
  theta <- parm[1:J]; tau <- 1/(sd*sd)
  ### Log(Prior Densities)
```

```

theta.mu.prior <- dnorm(theta.mu, theta.mu.mu,
    1/sqrt(theta.mu.tau), log=TRUE)
tau.prior <- theta.prior <- rep(0,J)
for (j in 1:J) {
    tau.prior[j] <- dgamma(tau[j], tau.alpha, tau.beta, log=TRUE)
    theta.prior[j] <- dnorm(theta[j], theta.mu,
        theta.sigma, log=TRUE)}
#### Log-Likelihood
LL <- sum(dnorm(y, theta, 1/sqrt(tau), log=TRUE))
#### Log-Posterior
LP <- LL + theta.mu.prior + sum(theta.prior) + sum(tau.prior)
Modelout <- list(LP=LP, Dev=-2*LL, Monitor=theta.sigma,
    yhat=theta, parm=parm)
return(Modelout)
}

```

5. Linear Regression

5.1. Form

$$\begin{aligned}
y &\sim N(\mu, \tau^{-1}) \\
\mu &= \mathbf{X}\beta \\
\beta_j &\sim N(0, 1000), \quad j = 1, \dots, J \\
\tau &\sim \Gamma(0.001, 0.001)
\end{aligned}$$

5.2. Data

```

N <- 10000
J <- 5
X <- matrix(1,N,J)
for (j in 2:J) {X[,j] <- rnorm(N, runif(1,-3,3), runif(1,0.1,1))}
beta <- runif(J,-3,3)
e <- rnorm(N,0,0.1)
y <- beta %*% t(X) + e
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=t(y))

```

5.3. Initial Values

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

5.4. Model

```
Model <- function(parm, Data)
{
  ### Prior Parameters
  beta.mu <- rep(0,J)
  beta.tau <- rep(1.0E-3,J)
  tau.alpha <- 1.0E-3
  tau.beta <- 1.0E-3
  ### Parameters
  beta <- rep(0,J)
  beta <- parm[1:J]
  tau <- exp(parm[J+1])
  ### Log(Prior Densities)
  beta.prior <- rep(0,J)
  for (j in 1:J) {
    beta.prior[j] <- dnorm(beta[j], beta.mu[j],
      1/sqrt(beta.tau[j]), log=TRUE)}
  tau.prior <- dgamma(tau, tau.alpha, tau.beta, log=TRUE)
  ### Log-Likelihood
  mu <- beta %*% t(X)
  LL <- sum(dnorm(y, mu, 1/sqrt(tau), log=TRUE))
  ### Log-Posterior
  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)
}
```

6. Linear Regression, Laplace-Distributed

This linear regression specifies that y is Laplace-distributed, where it is usually Gaussian or normally-distributed. It has been claimed that it should be surprising that the normal distribution became the standard, when the Laplace distribution usually fits better and has wider tails (Kotz, Kozubowski, and Podgorski 2001). Another popular alternative is to use the t-distribution, though it is more computationally expensive to estimate, because it has three parameters. The Laplace distribution has only two parameters, location and scale like the normal distribution, and is computationally easier to fit. This example could be taken one step further, and the parameter vector β could be Laplace-distributed. Laplace's Demon recommends that users experiment with replacing the normal distribution with the Laplace distribution.

6.1. Form

$$\begin{aligned} y &\sim L(\mu, \tau^{-1}) \\ \mu &= \mathbf{X}\beta \end{aligned}$$

$$\begin{aligned}\beta_j &\sim N(0, 1000), \quad j = 1, \dots, J \\ \tau &\sim \Gamma(0.001, 0.001)\end{aligned}$$

6.2. Data

```
N <- 10000
J <- 5
X <- matrix(1,N,J)
for (j in 2:J) {X[,j] <- rnorm(N,runif(1,-3,3),runif(1,0.1,1))}
beta <- runif(J,-3,3)
e <- rnorm(N,0,0.1)
y <- beta %*% t(X) + e
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=t(y))
```

6.3. Initial Values

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

6.4. Model

```
Model <- function(parm, Data)
{
  ### Prior Parameters
  beta.mu <- rep(0,J)
  beta.tau <- rep(1.0E-3,J)
  tau.alpha <- 1.0E-3
  tau.beta <- 1.0E-3
  ### Parameters
  beta <- rep(0,J)
  beta <- parm[1:J]
  tau <- exp(parm[J+1])
  ### Log(Prior Densities)
  beta.prior <- rep(0,J)
  for (j in 1:J) {
    beta.prior[j] <- dnorm(beta[j], beta.mu[j],
      1/sqrt(beta.tau[j]), log=TRUE)}
  tau.prior <- dgamma(tau, tau.alpha, tau.beta, log=TRUE)
  ### Log-Likelihood
  mu <- beta %*% t(X)
  LL <- sum(log(1 / (2*(1/sqrt(tau)))) * exp(-(abs(y - mu) /
    (1/sqrt(tau))))))
  ### Log-Posterior
  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)
}

```

7. Poisson Regression

7.1. Form

$$y \sim Pois(\lambda)$$

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

$$\beta \sim N(0, 1000)$$

7.2. Data

```

N <- 10000
J <- 5
X <- matrix(1,N,J)
for (j in 2:J) {X[,j] <- rnorm(N,runif(1,-3,3),runif(1,0.1,1))}
beta <- runif(J,-3,3)
e <- rnorm(N,0,0.1)
y <- exp(beta %*% t(X)) + e
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=t(y))

```

7.3. Initial Values

```
Initial.Values <- rep(0,J)
```

7.4. Model

```

Model <- function(parm, MyData)
{
  ### Prior Parameters
  beta.mu <- rep(0,J)
  beta.tau <- rep(1.0E-3,J)
  ### Parameters
  beta <- rep(0,J)
  beta <- parm[1:J]
  ### Log(Prior Densities)
  beta.prior <- rep(0,j)
}

```

```

for (j in 1:J) {
  beta.prior[j] <- dnorm(beta[j], beta.mu[j],
    1/sqrt(beta.tau[j]), log=TRUE)}
### Log-Likelihood
lambda <- exp(beta %*% t(X))
LL <- sum(dpois(y, lambda, log=TRUE))
### Log-Posterior
LP <- LL + sum(beta.prior)
Modelout <- list(LP=LP, Dev=-2*LL, Monitor=c(lambda[1:2]),
  yhat=lambda, parm=parm)
return(Modelout)
}

```

References

- Gelman A (2009). *R2WinBUGS: Running WinBUGS and OpenBUGS from R / S-PLUS*. R package version 2.1-16, URL <http://www.R-project.org/package=R2WinBUGS>.
- Gelman A, Carlin J, Stern H, Rubin D (2004). *Bayesian Data Analysis*. 2nd edition. Chapman & Hall, Boca Raton, FL.
- Hall B (2011). *LaplaceDemon: Software for Bayesian Inference*. R package version 11.01.18, URL <http://www.R-project.org/package=LaplaceDemon>.
- Kotz S, Kozubowski T, Podgorski K (2001). *The Laplace Distribution and Generalizations: A Revisit with Applications to Communications, Economics, Engineering, and Finance*. Birkhauser, Boston.
- 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>.

Affiliation:

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