

MCMC(Metropolis-Hastings & Gibbs) & JAGS

Justin Kao



Course Tools

RStudio: Your Data Science Playground

I code the data, and I tell the story.

1. Install R
2. Install Rstudio

Markdown & Quarto: Your Storytelling Superpowers

- 🧠 **RMarkdown or Quarto:** Make sure it *knits, renders, and saves* without RStudio crying.
(If it knits, you're officially a data scientist.)
- 📄 **Markdown Basics** — Learn how to make text bold, italic, draw diagrams, or show off anything you want to demonstrate.
- 🚀 **Hello, Quarto** by Dr. Mine Çetinkaya-Rundel (Duke University) — From “*What’s a chunk?*” to “*Look, I just published my work on GitHub!*”

Markov Chain Monte Carlo(MCMC)

MCMC - Metropolis-Hastings Algorithm

Let's recap of Key Concepts

- **Step 1:** Propose a candidate sample from the proposal (transition) distribution

$$g(x_{t+1} | x_t)$$

For example, if the proposal is Normal,

$$g(x_{t+1} | x_t) = \mathcal{N}(x_t, \sigma^2)$$

- **Step 2:** Accept the proposed value with probability

$$A(x_t \rightarrow x_{t+1})$$

Detail Balance Condition to reach stationary: $p(a)T(a \rightarrow b) = p(b)T(b \rightarrow a), \quad \forall a, b$

$$\frac{f(a)}{NC} g(b | a) A(a \rightarrow b) = \frac{f(b)}{NC} g(a | b) A(b \rightarrow a)$$

MCMC - Metropolis-Hastings Algorithm

$$\frac{f(a)}{NC} g(b | a) A(a \rightarrow b) = \frac{f(b)}{NC} g(a | b) A(b \rightarrow a)$$

Rewrite the balance condition:

$$\frac{A(a \rightarrow b)}{A(b \rightarrow a)} = \underbrace{\frac{f(b)}{f(a)}}_{r_f} \times \underbrace{\frac{g(a | b)}{g(b | a)}}_{r_g}$$

- **Scenario 1:** $r_f r_g < 1$, $A(a \rightarrow b) = r_f r_g$, $A(b \rightarrow a) = 1$
- **Scenario 2:** $r_f r_g > 1$, $A(a \rightarrow b) = 1$, $A(b \rightarrow a) = \frac{1}{r_f r_g}$

So $A(a \rightarrow b) = \min(r_f r_g, 1)$

This matches the expression shown in Dr. Marina Vannucci's slides:

$$A(\theta_{t-1} \rightarrow \theta^* = \theta_t) = \rho(\theta^{(t-1)}, \theta^*) = \min \left\{ \frac{p(\theta^* | x)}{p(\theta^{(t-1)} | x)} \frac{q(\theta^{(t-1)} | \theta^*)}{q(\theta^* | \theta^{(t-1)})}, 1 \right\}$$

$$\theta^{(t)} = \begin{cases} \theta^* & \text{with probability } \rho(\theta^{(t-1)}, \theta^*) \\ \theta^{(t-1)} & \text{with probability } 1 - \rho(\theta^{(t-1)}, \theta^*) \end{cases}$$

MCMC - Metropolis-Hastings Algorithm - Intuition

Let's assume the distribution is symmetric

$$\frac{A(a \rightarrow b)}{A(b \rightarrow a)} = \underbrace{\frac{f(b)}{f(a)}}_{r_f} \times \underbrace{\frac{g(a | b)}{g(b | a)}}_{r_g} = \underbrace{\frac{f(b)}{f(a)}}_{r_f} = \frac{p(b)}{p(a)} = \frac{p(\theta^* | y)}{p(\theta^{(t-1)} | y)} = r$$

$$A(a \rightarrow b) = \min\left(1, \frac{f(b)}{f(a)}\right) = \min\left(1, \frac{p(b)}{p(a)}\right) = \min\left(1, \frac{p(\theta^* | y)}{p(\theta^{(t-1)} | y)}\right) = r$$

• If $r > 1$:

- Intuition: Since $\theta^{(t-1)}$ is already in our set, we should include θ^* as it has a higher probability than $\theta^{(t-1)}$.
- Procedure: Accept θ^* into our set, i.e. set $\theta^{(s)} = \theta^*$.

• If $r < 1$:

- Intuition: The relative frequency of θ -values in our set equal to θ^* compared to those equal to $\theta^{(t-1)}$ should be $p(\theta^* | y) / p(\theta^{(t-1)} | y) = r$. This means that for every instance of $\theta^{(t-1)}$, we should have only a “fraction” of an instance of a θ^* value.
- Procedure: Set $\theta^{(t)}$ equal to either θ^* or $\theta^{(t-1)}$, with probability r and $1 - r$ respectively.

MCMC - Metropolis-Hastings Algorithm(HW problem 1)

Consider the file `school1.dat` (from the Book website) which contains data on the amount of times students at a high school spent on studying or homework during an exam period. Suppose we fit the following model:

$$Y_i \sim \mathcal{N}(\mu, \sigma^2), \quad i = 1, \dots, n$$

$$\mu \sim \mathcal{N}(\mu_0, \tau_0^2)$$

with $\mu_0 = 5, \tau_0^2 = 0.5$. Let us fix $\sigma^2 = 1$. Implement a Metropolis-Hastings:

```

1 library(coda)
2
3 # Read data
4 y <- scan(url("https://www2.stat.duke.edu/~pdh10/FCBS/Exercises/school1.dat"))
5 n <- length(y)
6
7 # Known hyperparameters
8 mu0 <- 5
9 tau0_sq <- 0.5
10 sigma_sq <- 1
11
12 # Log-posterior: log p(mu | y) ∝ log p(y | mu) + log p(mu)
13 log_post <- function(mu, y, mu0, tau0_sq, sigma_sq) {
14   log_lik <- sum(dnorm(y, mean = mu, sd = sqrt(sigma_sq), log = TRUE))
15   log_prior <- dnorm(mu, mean = mu0, sd = sqrt(tau0_sq), log = TRUE)
16   log_lik + log_prior
17 }

```

MCMC - Metropolis-Hastings Algorithm(HW problem 1, cont'd)

- A proposal function for μ centered at the previously sampled value.
- My proposal $\mu^* \sim q(\mu^* | \mu^{(0)}) = N(\mu^{(0)}, \sigma_p^2)$, where we can set $\mu^{(0)} = \bar{y}$

```
1 mean(y)
```

```
[1] 9.464
```

- $\mu^* = \mu^{(t-1)} + \varepsilon$, $\varepsilon \sim N(0, \sigma_p^2)$ Hence, $q(\mu^* | \mu^{(t-1)}) = N(\mu^{(t-1)}, \sigma_p^2)$

```
1 # Metropolis-Hastings sampler (Random-Walk proposal)
2 mh_sampler <- function(sigma_p, n_iter = 5000, burn_in = 1000) {
3   mu <- numeric(n_iter)
4   mu[1] <- mean(y)
5
6   for (t in 2:n_iter) {
7     mu_prop <- rnorm(1, mean = mu[t - 1], sd = sigma_p)
8     log_r <- log_post(mu_prop, y, mu0, tau0_sq, sigma_sq) -
9               log_post(mu[t - 1], y, mu0, tau0_sq, sigma_sq)
10    if (log(runif(1)) < log_r) {
11      mu[t] <- mu_prop
12    } else {
13      mu[t] <- mu[t - 1]
14    }
15  }
16
17  # Drop burn-in and return MCMC object
18  chain <- mcmc(mu[(burn_in + 1):n_iter])
19  return(chain)
20 }
```

MCMC - Metropolis-Hastings Algorithm(HW problem 1, contd)

```

1 # Run chain with sigma_p = 0.5
2 set.seed(425)
3 chain_small <- mh_sampler(sigma_p = 0.5)
4
5 # Summaries and plots
6 summary(chain_small)

```

Iterations = 1:4000
 Thinning interval = 1
 Number of chains = 1
 Sample size per chain = 4000

1. Empirical mean and standard deviation for each variable, plus standard error of the mean:

Mean	SD	Naive SE	Time-series SE
9.132470	0.193392	0.003058	0.006434


2. Quantiles for each variable:

2.5%	25%	50%	75%	97.5%
8.768	9.002	9.133	9.264	9.505

MCMC - Metropolis-Hastings Algorithm(HW problem 1, contd)

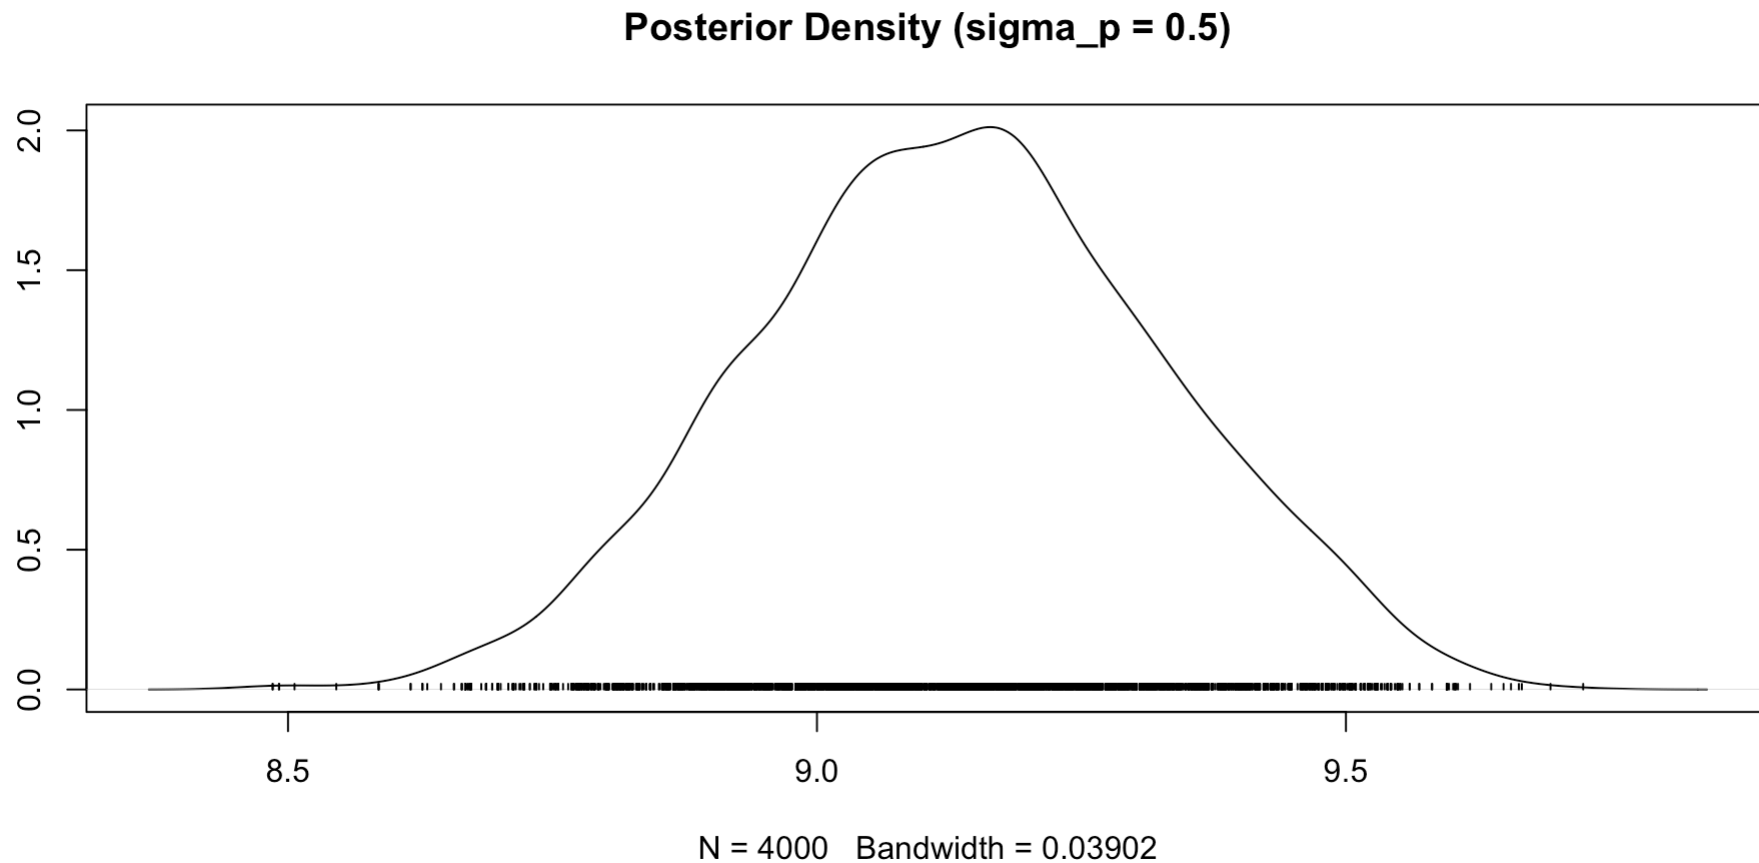
```
1 traceplot(chain_small, main = "Trace Plot (sigma_p = 0.5)")
```

Trace Plot (sigma_p = 0.5)



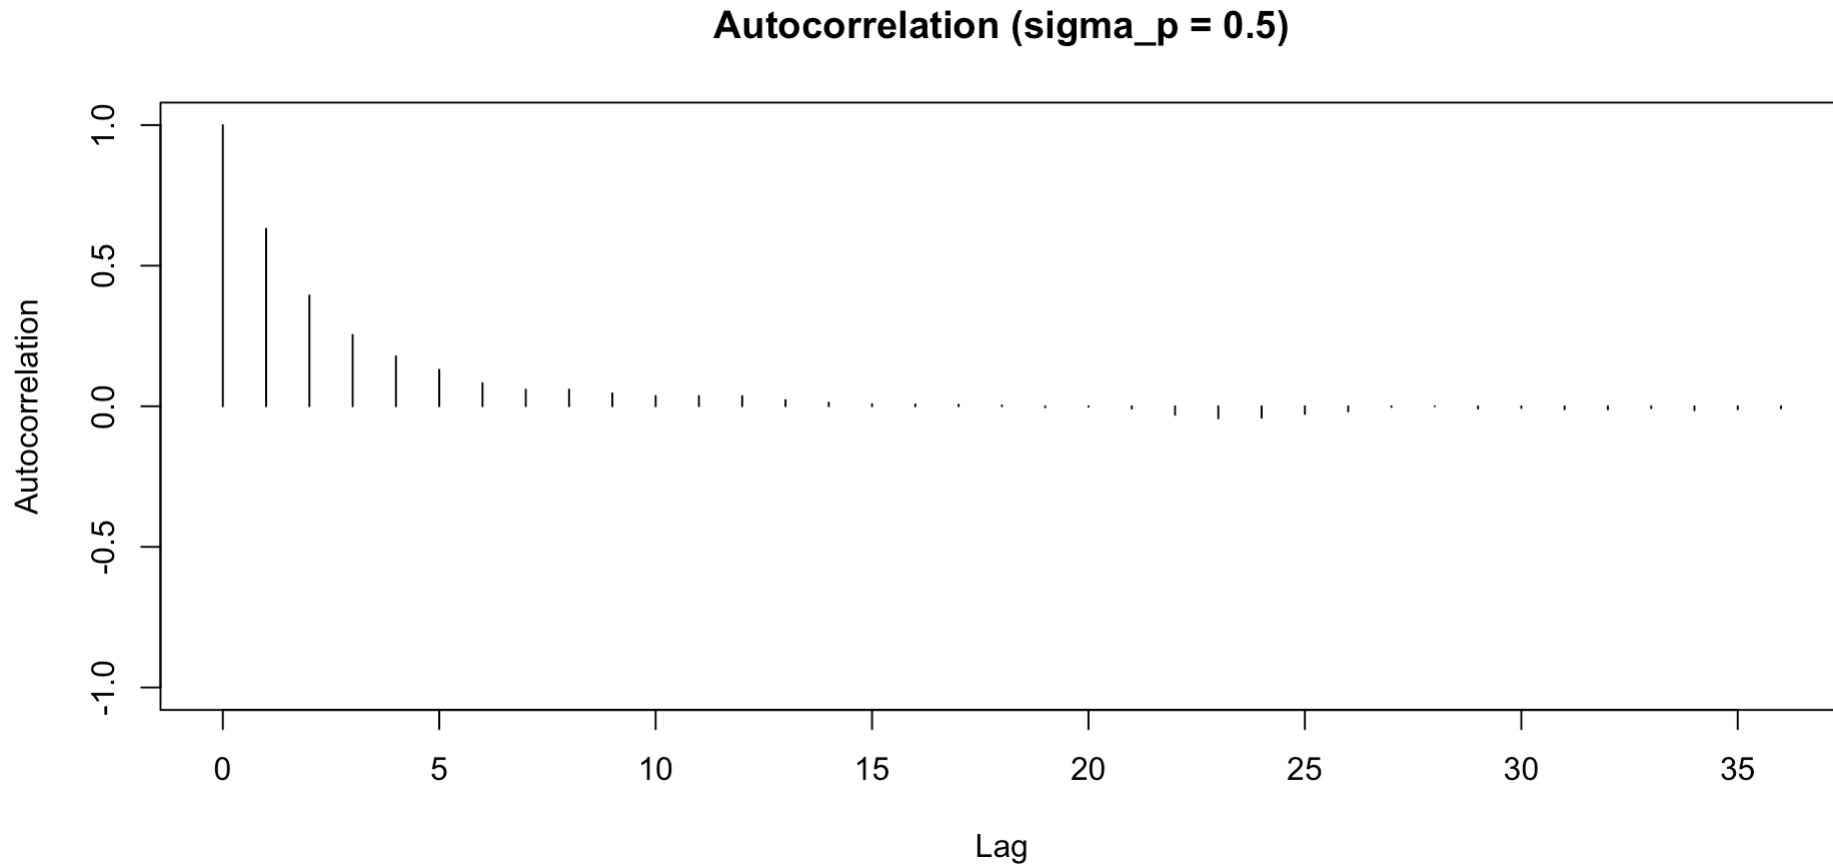
MCMC - Metropolis-Hastings Algorithm(HW problem 1, contd)

```
1 densplot(chain_small, main = "Posterior Density (sigma_p = 0.5)")
```



MCMC - Metropolis-Hastings Algorithm(HW problem 1, contd)

```
1 autocorr.plot(chain_small, main = "Autocorrelation (sigma_p = 0.5)")
```



MCMC - Metropolis-Hastings Algorithm(HW problem 1, contd)

(b) Provide a trace plot and a plot of the histogram and/or kernel density estimate for each chain.

To answer this, you can run the following code and then compare the plots:

```
1 set.seed(123)
2 chain1 <- mh_sampler(sigma_p = 0.5)
3 set.seed(456)
4 chain2 <- mh_sampler(sigma_p = 0.5)
5 set.seed(789)
6 chain3 <- mh_sampler(sigma_p = 0.5)
```

- `summary()`, `autocorr.plot()`, `densplot()`
- I have demonstrated the case when $\sigma_p = 0.5$. Don't forget to try $\sigma_p = 2$!

MCMC - Metropolis-Hastings Algorithm(HW problem 1, cont'd)

(c) After removing an appropriate burn-in, evaluate the acceptance rate and the autocorrelation for varying lags. Which value of σ_p provides a better choice for the proposal function?

```
1 # Acceptance rate function
2 acceptance_rate <- function(chain) {
3   draws <- as.numeric(chain)
4   mean(diff(draws) != 0)
5 }
6 acceptance_rate(chain_small)
```

```
[1] 0.4211053
```

(This is just an arbitrary number I came up with — you should check your own chain.)

The pattern is likely to be:

- $\sigma_p = 0.5$ accepts $\sim 82\%$ proposals \rightarrow too conservative.
- $\sigma_p = 2$ accepts $\sim 27\%$ proposals \rightarrow better exploration.

MCMC - Gibbs sampler(HW problem 2)

Gibbs sampler. Consider the file [school1.dat](#) (from the Book) which contains data on the amount of times students at a high school spent on studying or homework during an exam period. Suppose we fit the following model:

$$\begin{aligned} Y_i &\sim \mathcal{N}(\mu, \sigma^2), \quad i = 1, \dots, n \\ \mu &\sim \mathcal{N}(\mu_0, \tau_0^2) \\ \sigma^2 &\sim \text{Inv-Gamma}(\alpha, \beta) \end{aligned}$$

with $\mu_0 = 5, \tau_0^2 = 0.5, \alpha = 1, \beta = 4$.

(a) Write down the full conditionals, $p(\mu | \sigma^2, \mathbf{y})$ and $p(\sigma^2 | \mu, \mathbf{y})$.

$$\begin{aligned} p(\mu, \sigma^2 | \mathbf{y}) &\propto p(\mathbf{y} | \mu, \sigma^2) p(\mu) p(\sigma^2) \\ &\propto (\sigma^2)^{-n/2} \exp\left[-\frac{1}{2\sigma^2} \sum (y_i - \mu)^2\right] \times \exp\left[-\frac{1}{2\tau_0^2} (\mu - \mu_0)^2\right] \times (\sigma^2)^{-(\alpha+1)} \exp\left(-\frac{\beta}{\sigma^2}\right) \\ &\propto (\sigma^2)^{-(\alpha+n/2+1)} \exp\left[-\frac{1}{2\sigma^2} \sum (y_i - \mu)^2 - \frac{(\mu - \mu_0)^2}{2\tau_0^2} - \frac{\beta}{\sigma^2}\right]. \end{aligned}$$

MCMC - Gibbs sampler(HW problem 2, cont'd)

Say the full conditionals are:

- $\mu \mid \sigma^2, \mathbf{y} \sim N(\mu_n, \tau_n^2)$, where $\mu_n = \frac{\frac{n\bar{y}}{\sigma^2} + \frac{\mu_0}{\tau_0^2}}{\frac{n}{\sigma^2} + \frac{1}{\tau_0^2}}$, $\tau_n^2 = \left(\frac{n}{\sigma^2} + \frac{1}{\tau_0^2} \right)^{-1}$

- $\sigma^2 \mid \mu, \mathbf{y} \sim \text{Inv-Gamma} \left(\alpha + \frac{n}{2}, \beta + \frac{1}{2} \sum (y_i - \mu)^2 \right)$

Let's again setup our parameters first: $\mu_0 = 5, \tau_0^2 = 0.5, \alpha = 1, \beta = 4$

```

1 # Parameters
2 mu0 <- 5
3 tau0_sq <- 0.5
4 alpha <- 1
5 beta <- 4
6
7 # Data
8 y <- scan(url("https://www2.stat.duke.edu/~pdh10/FCBS/Exercises/school1.dat"))
9 n <- length(y)
10 ybar <- mean(y)

```

MCMC - Gibbs sampler(HW problem 2, cont'd)

```
1 # Storage
2 n_iter <- 10000
3 mu <- numeric(n_iter)
4 sigma2 <- numeric(n_iter)
5
6 # Initial values
7 mu[1] <- mean(y)
8 sigma2[1] <- var(y)
9
10 # Gibbs sampling
11 for (t in 2:n_iter) {
12   # 1. Sample mu | sigma2, y
13   tau_n_sq <- 1 / (n / sigma2[t-1] + 1 / tau0_sq)
14   mu_n <- tau_n_sq * (n * ybar / sigma2[t-1] + mu0 / tau0_sq)
15   mu[t] <- rnorm(1, mu_n, sqrt(tau_n_sq))
16
17   # 2. Sample sigma2 | mu, y
18   alpha_n <- alpha + n/2
19   beta_n <- beta + 0.5 * sum((y - mu[t])^2)
20   sigma2[t] <- 1 / rgamma(1, shape = alpha_n, rate = beta_n)
21 }
22
23 # Burn-in and analysis
24 burn_in <- 2000
25 mu_samp <- mu[(burn_in + 1):n_iter]
26 sigma2_samp <- sigma2[(burn_in + 1):n_iter]
27
```

MCMC - Gibbs sampler(HW problem 2, cont'd)

```
1 effectiveSize(mcmc(mu_samp))
```


```
var1  
4980.3
```

```
1 effectiveSize(mcmc(sigma2_samp))
```


```
var1  
5249.669
```

```
1 par(mfrow = c(1, 2))  
2 acf(mu_samp, main = "ACF of mu")  
3 acf(sigma2_samp, main = "ACF of sigma2")
```

ACF of mu



ACF of sigma2




```
1 par(mfrow = c(1, 1)) # reset layout
```


MCMC - Gibbs sampler(HW problem 2, cont'd)

```
1 plot(mu_mcmc, main = "Trace & Density of mu")
```

Trace & Density of mu




Trace & Density of mu




MCMC - Gibbs sampler(HW problem 2, cont'd)

```
1 plot(sigma2_mcmc, main = "Trace & Density of sigma2")
```

Trace & Density of sigma2

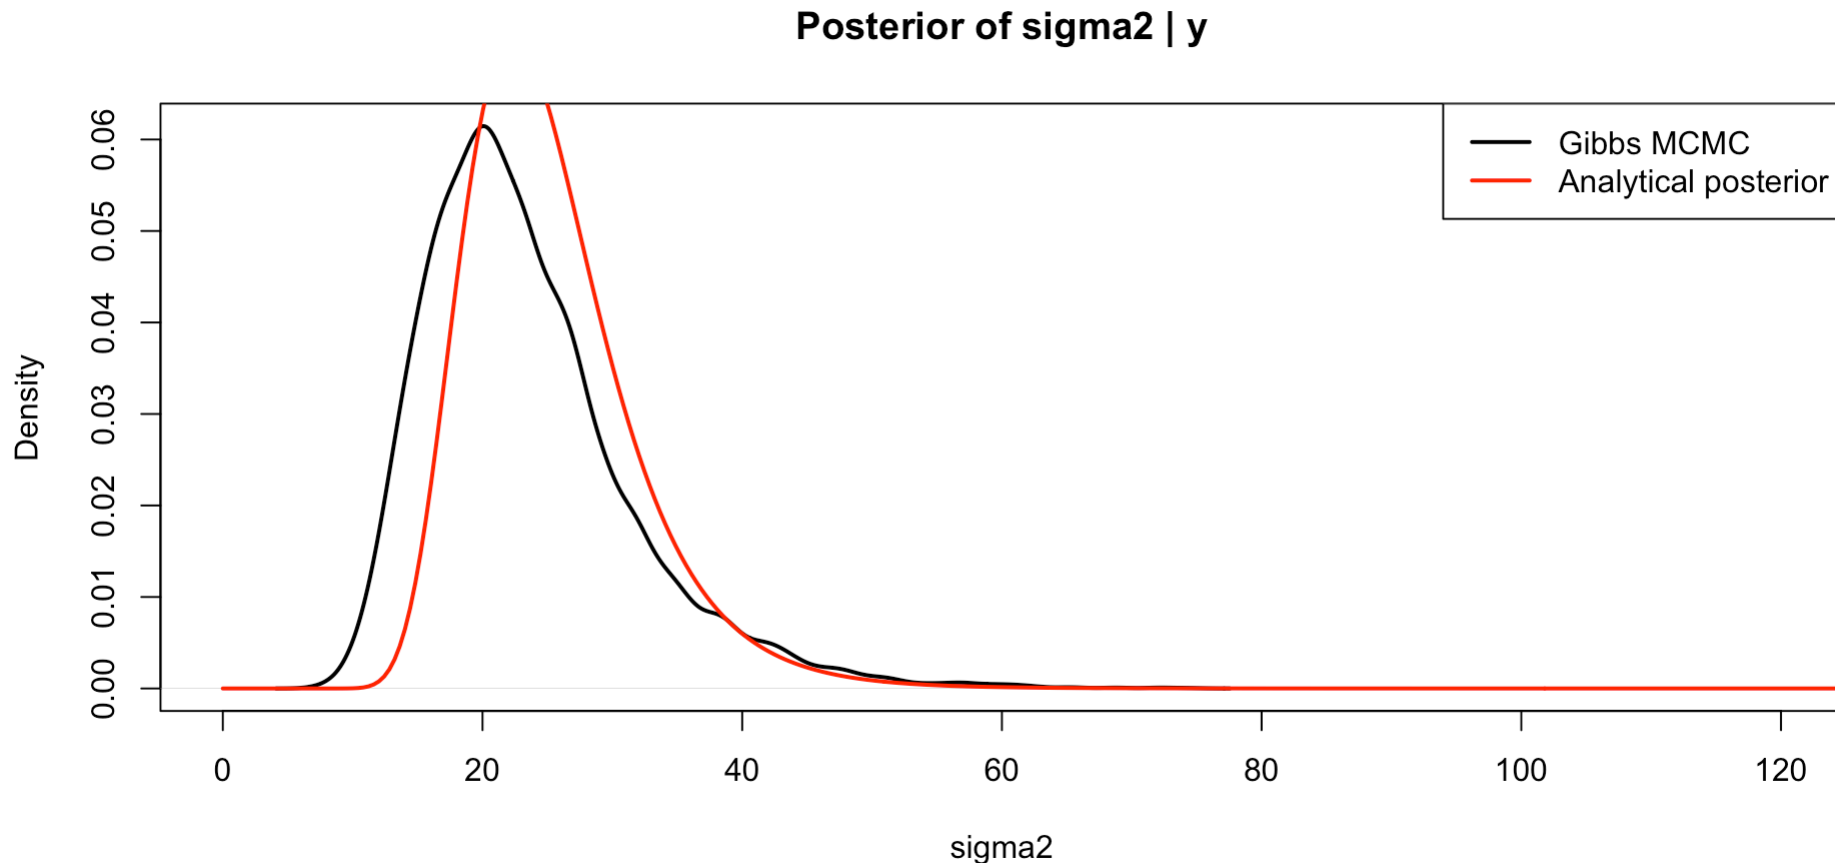


Trace & Density of sigma2



MCMC - Gibbs sampler(HW problem 2, cont'd)

(d) Evaluate the marginal posterior density $p(\sigma^2 | \mathbf{y})$ analytically up to a normalizing constant and overlay this density on your plot from (c) [Hint: you can perform the comparison by evaluating the unnormalized density on a grid of values].



MCMC - Gibbs sampler & MH(package(coda))

- You can technically wrap any numeric vector or matrix(`object` here) in `mcmc_obj <- mcmc(object)`, even if it isn't a "real" MCMC sequence. However - and this is key - doing so doesn't make it an MCMC chain in a statistical sense. It just labels it as one so that functions like `plot()`, `acf()`, and `effectiveSize()` can run without error.

Function	Purpose
<code>effectiveSize(mcmc_obj)</code>	Computes Effective Sample Size (ESS) — measures chain independence.
<code>gelman.diag(mcmc.list_obj)</code>	Gelman–Rubin diagnostic for checking convergence across multiple chains.
<code>geweke.diag(mcmc_obj)</code>	Geweke diagnostic for within-chain convergence (compares early vs. late samples).
<code>heidel.diag(mcmc_obj)</code>	Heidelberger–Welch test for stationarity and convergence.
<code>raftery.diag(mcmc_obj)</code>	Raftery–Lewis diagnostic estimating chain length needed for precision.

JAGS - Just Another Gibbs Sampler

The Posterior That Looks Easy but Isn't — JAGS Is Your Best Friend!

Model:

$$y_i \sim \text{Cauchy}(\theta, \sigma), \quad \theta \sim \mathcal{N}(0, 10^2), \quad \sigma \sim \text{Uniform}(0, 10)$$

Posterior:

$$p(\theta, \sigma | y) \propto \left[\prod_{i=1}^n \frac{1}{\pi\sigma \left(1 + \frac{(y_i - \theta)^2}{\sigma^2}\right)} \right] \exp\left(-\frac{\theta^2}{2 \times 10^2}\right)$$

It only has **two parameters**, yet the posterior is:

- **✗** No full conditional — not a known distribution
- **✗** Non-conjugate — prior and likelihood don't simplify
- **✗** No closed-form normalization — integral can't be solved

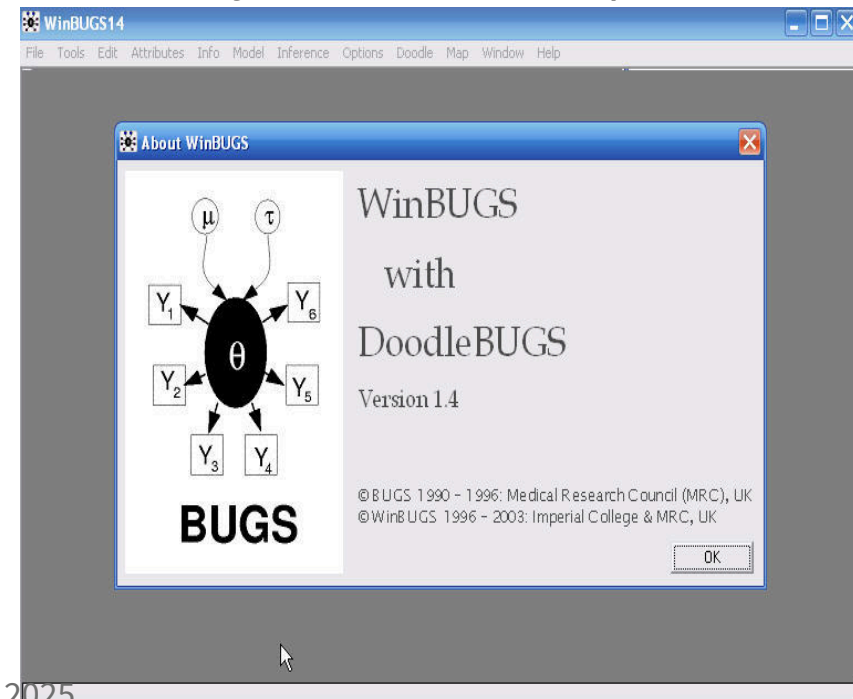
Result: $p(\theta, \sigma | y)$ has **no analytic form**, no direct sampler, and no closed expression.

🤔 Manual derivation stops here.

What Are BUGS and JAGS?

- **BUGS (Bayesian inference Using Gibbs Sampling):**
Early Bayesian software (1990s, Cambridge) using Gibbs sampling. Versions include *Classic BUGS(1990s)*, *WinBUGS(Windows-only)*, and *OpenBUGS(open-source successor)*.
- **JAGS (Just Another Gibbs Sampler):**
Developed by Martyn Plummer in 2003 as a **cross-platform, open-source, and R-integrated** version of BUGS.

💡 JAGS keeps the same model syntax as BUGS but adds flexibility, automatic Metropolis-within-Gibbs sampling, and modern extensibility.






Why JAGS

1. Doing MCMC by hand means:

- Writing out **priors** and **likelihood**
- Deriving **full conditional posteriors** for Gibbs sampling
- Specifying a **proposal distribution** for Metropolis–Hastings
- Tracking the **update order** for every parameter

2. As models grow complex, this becomes

-  *tedious*,
-  *algebraically messy*, and
-  *often intractable*.

3. You want to spend time **analyzing**, not deriving.

JAGS can help you solve this once you specify the prior and the likelihood!

Enter JAGS

- **JAGS** automates MCMC sampling.
- You define your model using **BUGS-like syntax**.
- It automatically:
 - Constructs the **posterior**
 - Performs **Gibbs / Metropolis-Hastings sampling**
 - Returns results to **R** through **rjags** or **R2jags**.

How JAGS Samples Behind the Scenes

Even within Gibbs sampling, not all parameters have closed-form conditionals.

JAGS handles this **automatically**:

Parameter Type	Sampling Strategy	Description
Conjugate blocks	Pure Gibbs	Exact draws from known conditionals
Non-conjugate blocks	Metropolis-within-Gibbs	Uses adaptive random-walk MH inside Gibbs
Difficult continuous nodes	Slice sampling	Efficient for highly nonlinear posteriors

So Gibbs and Metropolis **run in parallel**:

- Each iteration updates some parameters by **Gibbs**, others by **Metropolis**
- JAGS adapts proposal scales automatically during burn-in
- You don't need to code or tune anything

 *You define the model; JAGS decides the best sampler for each part.*

How to install JAGS(Mac)

- Option 1: [Download JAGS manually](#)
- Option 2: Homebrew (Mac)

```
1 brew install jags #run this command in terminal
2 brew list --versions jags #check version
3 brew uninstall jags #uninstall jags
4 brew upgrade #upgrade
```

- If you download manually, how to remove?

```
1 #run this command in terminal
2 which jags
3
4 #if you get this, then JAGS was installed manually
5 /usr/local/bin/jags
6
7 # To remove it, do
8 sudo rm -rf /usr/local/bin/jags
9 sudo rm -rf /usr/local/lib/JAGS
10 sudo rm -rf /usr/local/share/jags
11 sudo rm -rf /usr/local/include/jags
12
13 # Verify JAGS is removed
14 which jags
```

How to install JAGS(Windows)

1. Download JAGS manually

2. Using Command Prompt (CMD)

```
1 # Step 1 – Check if JAGS is installed
2 jags
3
4 # If JAGS is installed, you will see: Welcome to JAGS 4.3.2 (official binary)
5
6 # Type "exit" to close JAGS
7 exit
8
9 # Step 2 – Check the installation path (optional)
10 where jags
11
12 # Step 3 – To uninstall JAGS
13 # Go to Control Panel → Programs → Programs and Features
14 # Find "JAGS 4.x.x" → Click Uninstall
15
16 # Or delete manually (if Control Panel fails)
17 rmdir /S /Q "C:\Program Files\JAGS"
18
19 # Step 4 – Verify removal. It should show: 'jags' is not recognized as an internal or external command
20 jags
```

Loading JAGS(library(rjags)) in R

```
1 # install.packages(c("rjags", "coda"))
2
3 # Ensure dependencies are installed
4 # (coda must be loaded before rjags, since rjags depends on it)
5 install.packages(
6   "rjags",
7   type = "source",
8   repos = "https://cloud.r-project.org",
9   configure.args = "--with-jags-prefix=/opt/homebrew/opt/jags"
10 )
11
12 # Load required packages
13 library(coda) # For MCMC diagnostics and output handling
14 library(rjags) # Interface to the JAGS (Just Another Gibbs Sampler) engine
```

BUGS syntax - Example 1

$$y_i \sim N(\mu, \sigma^2), \quad \mu \sim N(0, 100), \quad \sigma \sim \text{Uniform}(0, 10)$$

```

1 # BUGS syntax (general form)
2 "
3 model {
4   for (i in 1:N) {
5     # Likelihood
6     y[i] ~ distribution(parameter1, parameter2) # stochastic: likelihood for data
7   }
8
9   # Priors for parameters
10  parameter1 ~ prior_distribution1(hyperparameter_set1) # prior for parameter1
11  parameter2_raw ~ prior_distribution2(hyperparameter_set2) # prior for parameter2 (before transformation)
12
13  # Deterministic relationship (if needed)
14  parameter2 <- g(parameter2_raw) # e.g., transformation such as 1 / variance
15 }
16 "
```

Then, append the `model{}` block into the R code

```

1 # Model specification
2 model_code <-
3 "
4 model {
5   for (i in 1:N) {
6     y[i] ~ dnorm(mu, tau) # likelihood for data
7   }
8   # Priors for parameters
9   mu ~ dnorm(0, 0.01) # Prior: mu ~ N(0, 100), in JAGS, always use precision = 1/100 = 0.01
10  sigma ~ dunif(0, 10) # Prior: sigma ~ Uniform(0, 10)
11  tau <- pow(sigma, -2) # Precision tau = 1 / sigma^2
12 }
13 "
14
15 # Write model to file
16 writeLines(model_code, "model.txt")
```

BUGS syntax - Example 2

$$y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \beta_0 \sim N(0, 1000), \beta_1 \sim N(0, 1000), \sigma^2 \sim \text{Inv-Gamma}(0.001, 0.001)$$

```

1 # BUGS syntax (general linear regression form)
2 "
3 model {
4   for (i in 1:N) {
5     # Likelihood
6     y[i] ~ dnorm(mu[i], tau)           # stochastic: likelihood for data
7     mu[i] <- beta0 + beta1 * x[i]     # deterministic: mean structure
8   }
9
10  # Priors for regression coefficients
11  beta0 ~ dnorm(0, 0.001)              # Prior:  $\beta_0 \sim N(0, 1000)$ 
12  beta1 ~ dnorm(0, 0.001)              # Prior:  $\beta_1 \sim N(0, 1000)$ 
13
14  # Prior for error variance
15  tau <- 1 / pow(sigma, 2)             # Precision =  $1 / \sigma^2$ 
16  sigma ~ dunif(0, 100)                # Equivalent to Inv-Gamma(0.001, 0.001)
17 }
18 "
```

Then, append the `model{}` block into the R code

```

1 # Model specification
2 model_code_2 <-
3 "
4 model {
5   for (i in 1:N) {
6     y[i] ~ dnorm(mu[i], tau)           # likelihood for data
7     mu[i] <- beta0 + beta1 * x[i]     # mean structure
8   }
9   beta0 ~ dnorm(0, 0.001)              # prior for intercept
10  beta1 ~ dnorm(0, 0.001)              # prior for slope
11  sigma ~ dunif(0, 100)                # prior for standard deviation
12  tau <- pow(sigma, -2)                 # precision =  $1 / \sigma^2$ 
13 }
14 "
15
16 # Write model to file
17 writeLines(model_code_2, "model_2.txt")
```

Hands-On with rjags: Let's Build a Bayesian Model

Example 1 (Cont'd)

Suppose we would like to run a model to draw inference about the mean for the following model:

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

where y_i are the data include 30 observations.

Priors are:


$$\mu \sim N(0, 100)$$

$$\sigma \sim \text{Uniform}(0, 10)$$

```
1 # Ground Truth
2 rm(list=ls())
3 n = 30
4 mu = 1.12
5 sigma = 0.38
6
7 # Generate values (stochastic part of the model)
8 # In reality, you only observe 30 data points, and the true population mean
9 # and variance are unknown. Our goal is to estimate these parameters.
10 set.seed(425)
11 y = rnorm(n, mean = mu, sd = sigma)
```

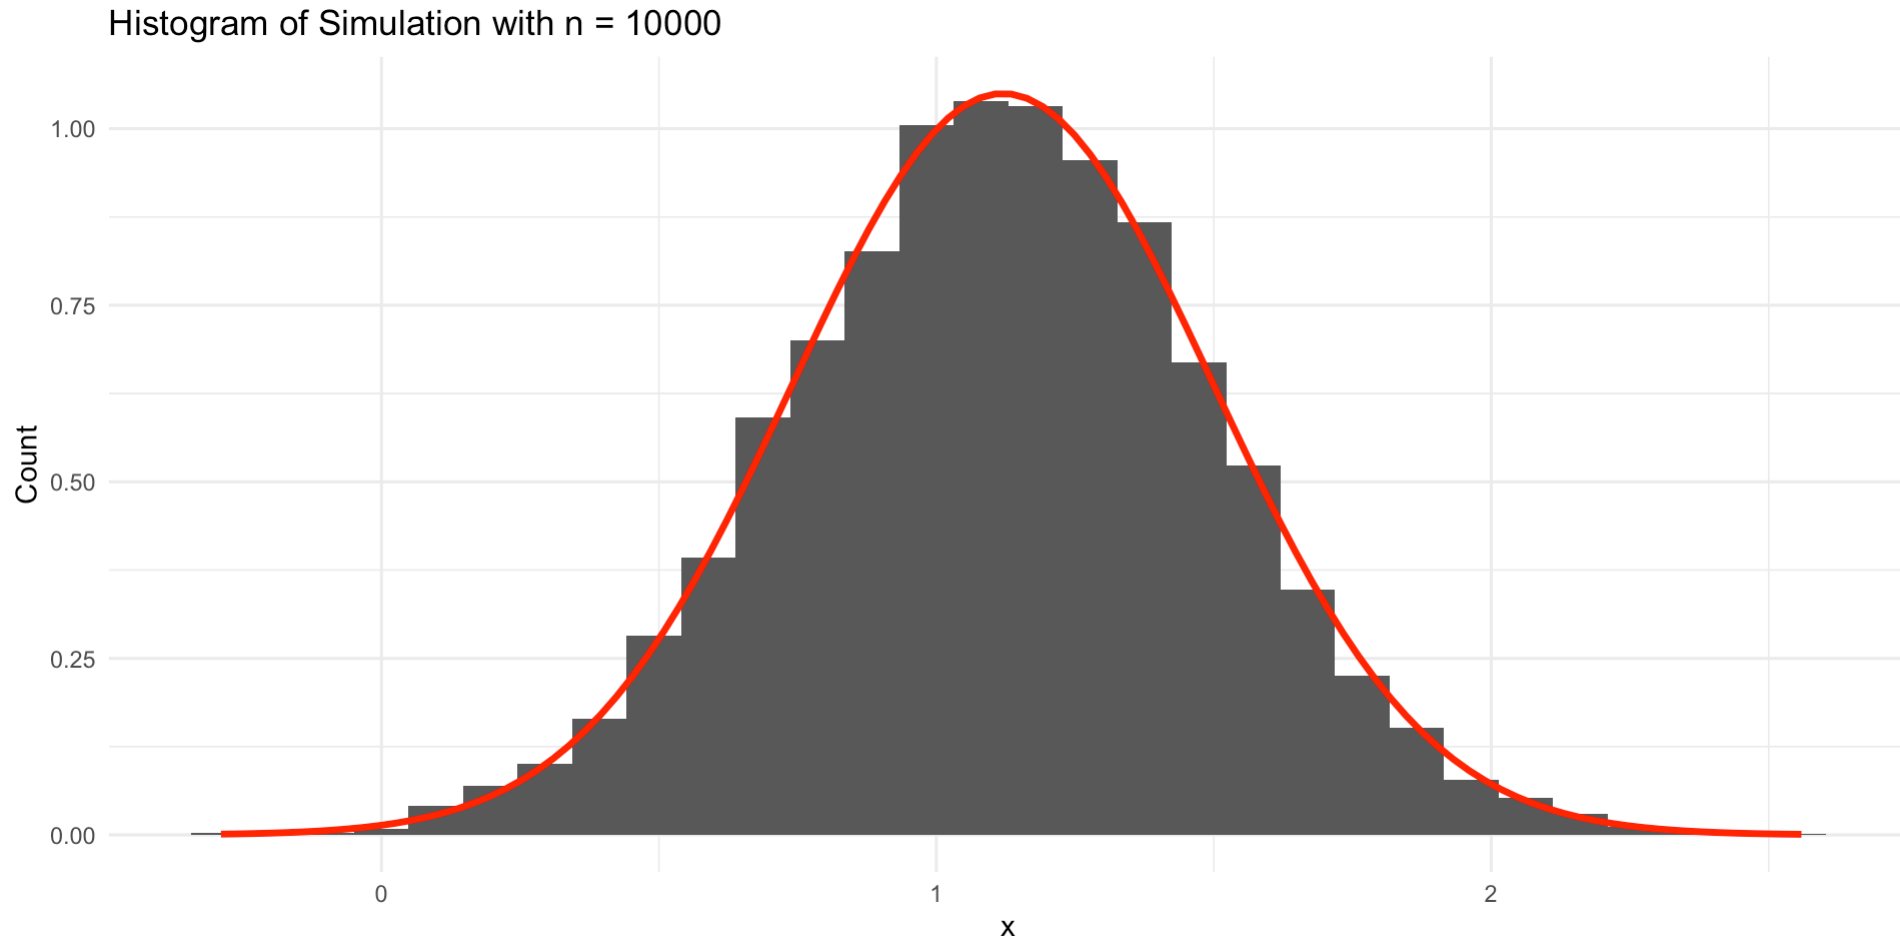
Example 1 (Cont'd)

```
1 df = data.frame(y)
2 library(ggplot2)
3 ggplot(df, aes(x = y)) +
4   geom_histogram(bins = 30) +
5   labs(title = "Histogram of Simulated y", x = "y", y = "Count") +
6   theme_minimal()
```



Example 1 (Cont'd)

If we put $n = 10000$, it is indeed the bell shape!



Example 1 (Cont'd)

- Recall our BUGS syntax: `writeln(model_code, "model.txt")`
- You don't really need to specify `inits`

```

1 # Prepare data for JAGS
2 jagsData <- list(
3   y = y,
4   N = length(y)
5 )
6
7 # Initial values for the model
8 jags_inits <- function() {
9   list(mu = rnorm(1, 0, 10), sigma = runif(1, 0, 10)) # priors
10 }
11
12 # Parameters to monitor
13 jags_params <- c("mu", "sigma", "tau")
14
15 # Load rjags and run the model
16 library(rjags)
17 model_1 <- jags.model(file = "model.txt", data = jagsData, inits = jags_inits, n.chains = 3)

```

Compiling model graph

Resolving undeclared variables

Allocating nodes

Graph information:

Observed stochastic nodes: 30

Unobserved stochastic nodes: 2

Total graph size: 39

Initializing model

Example 1 (Cont'd)

```
1 # Sample from posterior
2 samples <- coda.samples(model_1, variable.names = jags_params, n.iter = 2000, n.burnin = 1000)
3
4 # Convert to matrix for easier manipulation
5 output <- as.matrix(samples)
6
7 # Extract results
8 mu_samples <- output[, "mu"]
9 sigma_samples <- output[, "sigma"]
10 tau_samples <- output[, "tau"]
11
12 # Summary statistics
13 # posterior prediction
14 mean(mu_samples)
```

```
[1] 1.082436
```

```
1 mean(sigma_samples)
```

```
[1] 0.3665139
```







```
1 mean(tau_samples)
```

```
[1] 7.876675
```

- Recall our ground truth: $\mu = 1.12$, $\sigma = 0.38$. What will you conclude about the result from JAGS.

Example 1 (Cont'd) - Trace Plot

- 1 # Plot MCMC chains
- 2 `plot(samples)`



Example 1 (Cont'd) - Autocorrelation Plot & Effective Sample Size

```
1 autocorr.diag(samples)
```

	mu	sigma	tau
Lag 0	1.000000e+00	1.00000000	1.00000000
Lag 1	1.772588e-02	0.33924252	0.262195963
Lag 5	9.288274e-03	0.01052557	0.002459708
Lag 10	-3.721319e-05	-0.01503355	-0.022877238
Lag 50	-7.350899e-04	-0.03318907	-0.027529983

```
1 effectiveSize(samples)
```


mu	sigma	tau
5835.195	2712.066	3251.849

```
1 # autocorr.plot(samples)
```


Example 1 (Cont'd) - Gelman-Rubin

- 1 #Examine convergence of the Markov chains using the Gelman-Brooks-Rubin diagnostic
- 2 `gelman.plot(samples)`


mu



sigma



tau



Example 2 (Cont'd)

Simple Linear Regression: We have the data of 48 pigs with body weight measured at 9 successive weeks. Suppose we would like to fit the model: for $i = 1, \dots, 432$

$$y_i = \beta_0 + \beta_1 x_i + \epsilon_i \quad \epsilon_i \sim N(0, \sigma^2)$$

Priors are:

$$\beta_0 \sim N(0, 1000)$$

$$\beta_1 \sim N(0, 1000)$$

$$\sigma^2 \sim \text{Inv-Gamma}(0.001, 0.001)$$

```

1 rm(list=ls())
2 # Ground Truth
3 set.seed(425)
4 N <- 48
5 beta0_true <- 2
6 beta1_true <- 0.5
7 sigma_true <- 1
8 # Generate synthetic data
9 x <- rnorm(N, 5, 2)
10 y <- rnorm(N, beta0_true + beta1_true * x, sigma_true)

```

Example 2 (Cont'd)

Recall our BUGS syntax: `writeLines(model_code_2, "model_2.txt")`

```

1 # Prepare data for JAGS
2 jags_data <- list(
3   N = N,
4   x = x,
5   y = y
6 )
7
8 # Initial values for chains
9 inits <- function() {
10  list(beta0 = rnorm(1, 0, 1),
11        beta1 = rnorm(1, 0, 1),
12        sigma2 = runif(1, 0, 5))
13 }
14
15 # Parameters to monitor
16 params <- c("beta0", "beta1", "sigma2", "tau2")
17
18 # Load rjags and run the model
19 library(rjags)
20 model_2 <- jags.model("model_2.txt",
21                       data = jags_data,
22                       inits = inits,
23                       n.chains = 3)

```

Compiling model graph

Resolving undeclared variables

Allocating nodes

Graph information:

Observed stochastic nodes: 48

Unobserved stochastic nodes: 3

Total graph size: 202

Initializing model

Example 2 (Cont'd)

```

1 samples_2 <- coda.samples(model_2, variable.names = params, n.iter = 3000, n.burnin = 1000)
2
3 # Convert to matrix for easier manipulation
4 output_2 <- as.matrix(samples_2)
5
6 # Extract results
7 beta0_samples <- output_2[, "beta0"]
8 beta1_samples <- output_2[, "beta1"]
9 sigma_samples_2 <- output_2[, "sigma2"]
10 tau_samples_2 <- output_2[, "tau2"]
11
12 # Summary statistics
13 # posterior prediction
14 mean(beta0_samples)

```

```
[1] 1.978377
```

```
1 mean(beta1_samples)
```

```
[1] 0.51849
```

```
1 mean(sigma_samples_2)
```

```
[1] 0.9782063
```







```
1 mean(tau_samples_2)
```

```
[1] 1.079328
```

- Recall our ground truth: $\beta_0_{\text{true}} = 2$, $\beta_1_{\text{true}} = 0.5$, $\sigma_{\text{true}} = 1$. What will you conclude about the result from JAGS.

Example 2 (Cont'd) - Trace Plot

```
1 # Plot MCMC chains
2 plot(samples_2)
```



Example 2 (Cont'd) - Autocorrelation Plot & Effective Sample Size

```
1 autocorr.diag(samples_2)
```

	beta0	beta1	sigma2	tau2
Lag 0	1.000000000	1.000000000	1.00000000000	1.0000000000
Lag 1	0.87054836	0.87023150	0.2882158506	0.254044303
Lag 5	0.49285648	0.49501826	0.0233869895	0.022817377
Lag 10	0.22422971	0.22620577	0.0002910298	-0.002729308
Lag 50	0.02637208	0.02833683	-0.0314088728	-0.030339463

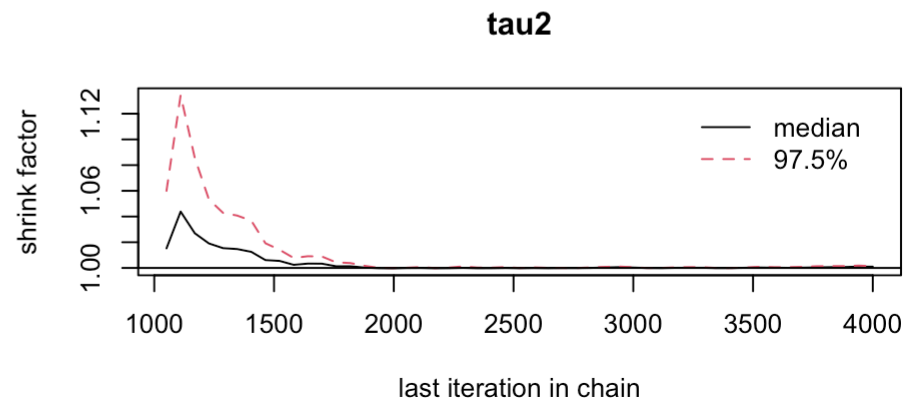
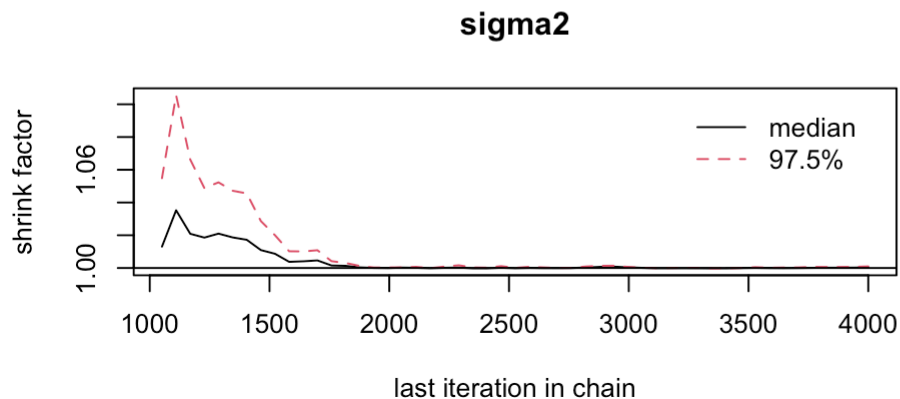
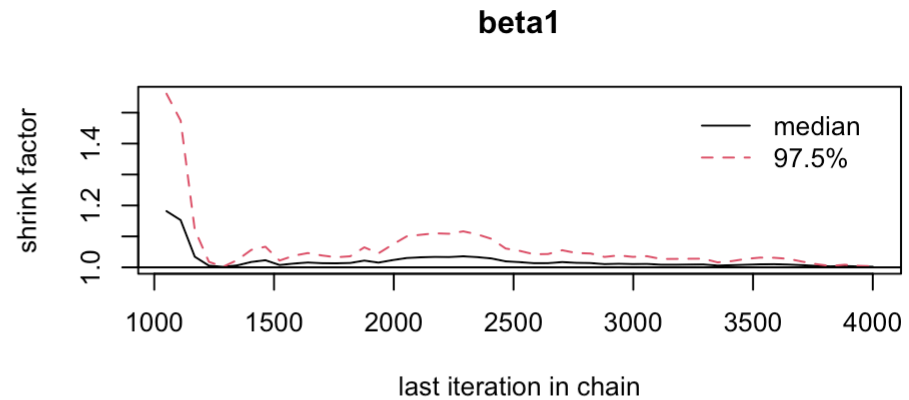

```
1 effectiveSize(samples_2)
```

	beta0	beta1	sigma2	tau2
	629.0657	624.3691	4972.6959	5351.8985


```
1 # autocorr.plot(samples)
```

Example 2 (Cont'd) - Gelman-Rubin

- 1 #Examine convergence of the Markov chains using the Gelman-Brooks-Rubin diagnostic
- 2 `gelman.plot(samples_2)`



Finally



Resources

- [JAGS \(official site\)](#)
- *JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling* by Martyn Plummer — University of Warwick
- Rebecca Steorts, *Introduction to Just Another Gibbs Sampler (JAGS)* — Duke University
- *An Introduction to Bayesian Reasoning and Methods* by Kevin Davis
- *A First Course in Bayesian Statistical Methods* by Dr. Peter Hoff
- Ritvik Kharkar
- Original slide materials were prepared by [Dr. Scott Liang](#) and [Enyu Li](#).
- Updated and expanded under the supervision of [Dr. Marina Vannucci](#).

Example 3: Random-Intercept Model in Action (Optional practice)

Same pig dataset, assume intercept is different for each pig Data likelihood: for $i = 1, \dots, 48, j = 1, \dots, 9$

$$y_{ij} = \beta_0 + \theta_i + \beta_1 x_{ij} + \varepsilon_{ij}, \quad \theta_i \sim N(0, \tau^2), \quad \varepsilon_{ij} \sim N(0, \sigma^2)$$

$$\beta_0 \sim N(0, 1000), \quad \beta_1 \sim N(0, 1000)$$

$$\tau^2 \sim \text{Inv-Gamma}(0.001, 0.001), \quad \sigma^2 \sim \text{Inv-Gamma}(0.001, 0.001)$$

```

1 # Model specification
2 model_code <-
3 "
4   model {
5     for (i in 1:N_pig) {
6       for (j in 1:N_obs) {
7         y[i, j] ~ dnorm(mu[i, j], tau_eps)
8         mu[i, j] <- beta0 + theta[i] + beta1 * x[i, j]
9       }
10      theta[i] ~ dnorm(0, tau_theta)
11    }
12    beta0 ~ dnorm(0, 0.001)
13    beta1 ~ dnorm(0, 0.001)
14    sigma ~ dunif(0, 100)
15    tau_eps <- pow(sigma, -2)
16    sigma_theta ~ dunif(0, 100)
17    tau_theta <- pow(sigma_theta, -2)
18  }
19 "
20
21 # Write model to file
22 writeLines(model_code, "model.txt")

```

Note

Feel free to email me at jk201@rice.edu if you want to check your answer or go over the solution together.

Thank you!

Have a great rest of your day!