Markov chain Monte Carlo I

Bayesian Models for Ecologists

Tom Hobbs

June 5, 2024



What is this course about?



You can understand it.

Rules of probability

- Conditioning and independence
- Law of total probability
- Factoring joint probabilities
- Distribution theory

Markov chain Monte Carlo

Current concepts, theory, questions



New insight qualified by uncertainty Introduction Why MCMC? Intuition Writing an MCMC sampler

The Bayesian method



The MCMC algorithm

► Why MCMC?

- Some intuition about how it works for a single parameter model
- MCMC for multiple parameter models
 - Gibbs sampling (lab)
 - Full-conditional distributions (MCMC II)
 - Metropolis-Hastings algorithm (MCMC II)
 - MCMC software (JAGS tomorrow)

MCMC learning outcomes

- 1. Develop a big picture understanding of how MCMC allows us to approximate the marginal posterior distributions of parameters and latent quantities.
- 2. Understand and be able to code a simple MCMC algorithm.
- 3. Appreciate the different methods that can be used within MCMC algorithms to make draws from the posterior distribution.
 - 3.1 Metropolis
 - 3.2 Gibbs
 - 3.3 Metropolis-Hastings
- 4. Understand concepts of burn-in and convergence.
- 5. Understand and be able to write full-conditional distributions

Remember the marginal distribution of the data



We have simple solutions for the posterior for simple models:



Problems of high dimension do not have simple solutions:

$\begin{bmatrix} \theta_1, \theta_2, \theta_3, \theta_4, \mathbf{z} \mid \mathbf{y}, \mathbf{u} \end{bmatrix} = \\ \frac{\prod_{i=1}^n [y_i \mid \theta_1 z_i] [u_i \mid \theta_2, z_i] [z_i \mid \theta_3, \theta_4] [\theta_1] [\theta_2] [\theta_3] [\theta_4]}{\int \dots \int \prod_{i=1}^n [y_i \mid \theta_1 z_i] [u_i \mid \theta_2, z_i] [z_i \mid \theta_3, \theta_4] [\theta_1] [\theta_2] [\theta_3] [\theta_4] dz_i d\theta_1 d\theta_2 d\theta_3 d\theta_4}$

What we are doing in MCMC?

Recall that the posterior distribution is proportional to the joint: because the marginal distribution of the data $\int [y|\theta][\theta]d\theta$ is a constant after the data have been observed.



Factoring the joint distribution into a product of probability distributions using the chain rule of probability is where we start all Bayesian modeling. The factored joint distribution provides the basis for MCMC.

What we are doing in MCMC?



What we are doing in MCMC?



What are we doing in MCMC?

- The posterior distribution is unknown, but the likelihood is known as a likelihood profile and we know the priors.
- We want to accumulate many, many values that represent random samples proportionate to their density in the *marginal* posterior distribution of each unobserved quantity of interest.
- MCMC generates these samples using the likelihood and the priors to decide which samples to keep and which to throw away.
- We can then use these samples to calculate statistics describing the marginal distributions of the unobserved quantities: means, medians, variances, credible intervals etc.

What are we doing in MCMC?

The marginal posterior distribution of each unobserved quantity is approximated by samples accumulated in the chain.



What are we doing in MCMC?



We keep the more probable members of the posterior distribution by comparing a proposal with the current value in the chain.

$$\begin{array}{cccc} k & 1 & 2\\ \text{Proposal}\,\theta^{*k+1} & \theta^{*2}\\ \text{Test} & P(\theta^{*2}) > P(\theta^{1})\\ \text{Chain}(\theta^{k}) & \theta^{1} & \theta^{2} = \theta^{*2} \end{array}$$

We keep the more probable members of the posterior distribution by comparing a proposal with the current value in the chain.

$$k \qquad 1 \qquad 2 \qquad 3$$

Proposal θ^{*k+1} θ^{*2} θ^{*3}
Test $P(\theta^{*2}) > P(\theta^{1})$ $P(\theta^{2}) > P(\theta^{*3})$
Chain (θ^{k}) θ^{1} $\theta^{2} = \theta^{*2}$ $\theta^{3} = \theta^{2}$

We keep the more probable members of the posterior distribution by comparing a proposal with the current value in the chain.



The cunning idea behind Metropolis updates



When do we keep the proposal?

 $P_R = \min(1, R)$

Keep θ^{*k+1} as the next value in the chain with probability P_R and keep θ^k with probability $1 - P_R$.

When do we keep the proposal?

- 1. Calculate R based on likelihoods and priors.
- 2. Draw a random number, U from uniform distribution 0,1 If R > U, we keep the proposal θ^{*k+1} as the next value in the chain.
- 3. Otherwise, we retain θ^k as the next value.

A simple example for one parameter

- Joy is interested in estimating the proportion of frogs in a county that are infected with Chytrid fungus.
- She is not very ambitious, so he only checks 12 frogs, 3 of which are infected. She assumes there is no prior knowledge of this proportion.
- How could she calculate the parameters of the posterior distribution of the proportion of the frog population that is infected (prevalence) on the back of a cocktail napkin?

The model

$[\phi|y] \propto \mathsf{binomial}(y|n, \phi)\mathsf{beta}(\phi|1, 1)$



Simulated and Calculated Distribution, iterations = 100

φ



Simulated and Calculated Distribution, iterations = 1000

φ



Simulated and Calculated Distribution, iterations = 10000



Simulated and Calculated Distribution, iterations = 100000





The chain has *converged* when adding more samples does not change the shape of the posterior distribution. We throw away samples that are accumulated before convergence (burn-in).

Intuition for MCMC for multi-parameter models



$$\begin{split} g(\beta_0,\beta_1,x_i) &= \beta_0 + \beta_1 x_i \\ [\beta_0,\beta_1,\sigma^2 \mid y_i] \propto [\beta_0,\beta_1,\sigma^2,y_i] \\ \text{factoring rhs using DAG:} \\ [\beta_0,\beta_1,\sigma^2 \mid y_i] \propto [y_i \mid g(\beta_0,\beta_1,x_i),\sigma^2][\beta_0],[\beta_1][\sigma^2] \\ \text{joint for all data :} \\ \end{split}$$

$$egin{aligned} &[eta_0,eta_1,\sigma^2\midoldsymbol{y}]\propto\prod_{i=1}^n[y_i\mid g(eta_0,eta_1,x_i),\sigma^2][eta_0][eta_1][\sigma^2] \end{aligned}$$

choose specific distributions:

$$\begin{split} [\beta_0, \beta_1, \sigma^2 \mid \pmb{y}] \propto \prod_{i=1}^n \operatorname{normal}(y_i \mid g(\beta_0, \beta_1, x_i), \sigma^2) \\ \times \operatorname{normal}(\beta_0 \mid 0, 10000) \operatorname{normal}(\beta_1 \mid 0, 10000) \\ \times \operatorname{uniform}(\sigma^2 \mid 0, 500) \end{split}$$

Intuition for MCMC for multi-parameter models

$$\begin{aligned} [\beta_0, \beta_1, \sigma^2 \mid \mathbf{y}] &\propto \prod_{i=1}^n \operatorname{normal}(y_i | g(\beta_0, \beta_1, x_i), \sigma^2) \\ &\times \operatorname{normal}(\beta_0 \mid 0, 10000) \operatorname{normal}(\beta_1 \mid 0, 10000) \\ &\times \operatorname{uniform}(\sigma^2 \mid 0, 100) \end{aligned}$$

- 1. Set initial values for $\beta_0, \beta_1, \sigma^2$
- 2. Assume β_1, σ^2 are known and constant. Make a draw for β_0 . Store the draw.
- 3. Assume β_0, σ^2 are known and constant. Make a draw for β_1 . Store the draw.
- 4. Assume β_0, β_1 are known and constant. Make a draw for σ^2 . Store the draw.
- 5. Do this many times. The stored values for each parameter approximate its marginal posterior distribution after convergence.

Do Lab

Implementing MCMC for multiple parameters

- Write an expression for the posterior and joint distribution using a DAG as a guide. Always.
- If you are using MCMC software (e.g. JAGS) use expression for the posterior and joint distribution as template for writing code. You are done.
- If you are writing your own MCMC sampler or you simply want to understand what JAGS is doing for you:
 - Decompose the expression of the multivariate joint distribution into a series of univariate distributions called *full-conditional distributions*.
 - Choose a sampling method for each full-conditional distribution.
 - Cycle through each unobserved quantity, sampling from its full-conditional distribution, treating the others as if they were known and constant.
 - The accumulated samples approximate the marginal posterior distribution of each unobserved quantity.
 - Note that this takes a complex, multivariate problem and turns it into a series of simple, univariate problems that we solve, as in the example above, one at a time.

Definition of full-conditional distribution

Let $\boldsymbol{\theta}$ be a vector of length k containing all of the unobserved quantities we seek to understand. Let $\boldsymbol{\theta}_{_j}$ be a vector of length k-1 that contains all of the unobserved quantities *except* θ_j . The full-conditional distribution of θ_j is

$$[\theta_j|y,\boldsymbol{\theta}_{j}],$$

which we notate as

 $[\theta_j|\cdot].$

It is the posterior distribution of θ_j conditional on all of the other parameters and the data, which we assume are *known*.

Writing full-conditional distributions

- You will have one full-conditional for each unobserved quantity in the posterior.
- For each unobserved quantity, write the distributions where it appears.
- Ignore the other distributions.
- Simple.

Example

- Clark 2003 considered the problem of modeling fecundity of spotted owls and the implication of individual variation in fecundity for population growth rate.
- Data were number of offspring produced by per pair of owls with sample size n = 197.

Clark, J. S. 2003. Uncertainty and variability in demography and population growth: A hierarchical approach. Ecology 84:1370-1381.

Example



$$[\boldsymbol{\lambda}, \boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{y}] \propto \prod_{i=1}^{n} \text{Poisson}(y_i | \lambda_i) \text{gamma}(\lambda_i | \boldsymbol{\alpha}, \boldsymbol{\beta})$$

 \times gamma (α |.001,.001) gamma (β |.001,.001)

Full-conditionals

 $[\boldsymbol{\lambda}, \alpha, \beta | \mathbf{y}] \propto \prod_{i=1}^{n} \operatorname{Poisson}\left(y_i | \lambda_i\right) \operatorname{gamma}\left(\lambda_i | \alpha, \beta\right) \operatorname{gamma}\left(\beta | .001, .001\right) \operatorname{gamma}\left(\alpha | .001, .001\right)$



We use the multivariate joint distribution to find univariate fullconditional distributions for all unobserved quantities.

How many full conditionals are there?

Writing full-conditional distributions

- You will have one full-conditional for each unobserved quantity in the posterior.
- For each unobserved quantity, write the distributions (including products) where it appears.
- Ignore the other distributions.
- Simple.

Full-conditional for each λ_i

 $[\boldsymbol{\lambda}, \alpha, \beta | \mathbf{y}] \propto \prod_{i=1}^{n} \text{Poisson}\left(y_i | \lambda_i\right) \text{gamma}\left(\lambda_i | \alpha, \beta\right) \text{gamma}\left(\beta | .001, .001\right) \text{ gamma}\left(\alpha | .001, .001\right)$ Writing the full-conditional distribution for λ_i : $[\lambda_i \mid .] \propto \text{Poisson}(y_i \mid \lambda_i) \text{gamma}(\lambda_i \mid \alpha, \beta)$ y_i α

Full-conditional for β





Full-conditional for lpha



Writing the full-conditional distribution for α :

$$[\alpha \mid \cdot] \propto \prod_{i=1}^{n} \operatorname{gamma}(\lambda_{i} \mid \alpha, \beta) \operatorname{gamma}(\alpha \mid .001, .001)$$



Full-conditionals for the model

Posterior and joint:

$$\begin{aligned} [\boldsymbol{\lambda}, \boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{y}] &\propto \prod_{i=1}^{n} \mathsf{Poisson}\left(y_{i} | \lambda_{i}\right) \mathsf{gamma}\left(\lambda_{i} | \boldsymbol{\alpha}, \boldsymbol{\beta}\right) \\ &\times \mathsf{gamma}\left(\boldsymbol{\alpha} | .001, .001\right) \mathsf{gamma}\left(\boldsymbol{\beta} | .001, .001\right) \end{aligned}$$

Full conditionals:

$$[\lambda_i|.] \propto \text{Poisson}(y_i|\lambda_i) \text{gamma}(\lambda_i|\alpha,\beta)$$

 $[\beta|.] \propto \prod_{i=1}^n \text{gamma}(\lambda_i|\alpha,\beta) \text{gamma}(\beta|.001,.001)$
 $[\alpha|.] \propto \prod_{i=1}^n \text{gamma}(\lambda_i|\alpha,\beta) \text{gamma}(\alpha|.001,.001)$

Implementing MCMC for multiple parameters and latent quantities

- Write an expression for the posterior and joint distribution using a DAG as a guide. Always.
- If you are using MCMC software (e.g. JAGS) use expression for posterior and joint as template for writing code.
- If you are writing your own MCMC sampler:
 - Decompose the expression of the multivariate joint distribution into a series of univariate distributions called *full-conditional distributions*.
 - Choose a sampling method for each full-conditional distribution.
 - Cycle through each unobserved quantity, sampling from the its full-conditional distribution, treating the others as if they were known and constant.
 - Note that this takes a complex, multivariate problem and turns it into a series of simple, univariate problems that we solve, as in the example above, one at a time.

Choosing a sampling method

1. Acept-reject:

- 1.1 Metropolis: requires a symmetric proposal distribution (e.g., normal, uniform). This is what we used above in the example for one parameter.
- Metropolis-Hastings: allows asymmetric proposal distributions (e.g., beta, gamma, lognormal). A minor modification of Metropolis. See optional notes.
- 2. Gibbs: accepts all proposals because they are especially well chosen. Requires conjugates. In lab today.

Why do you need to understand conjugate priors?

- A easy way to find parameters of posterior distributions for simple problems.
- Critical to understanding Gibbs updates in Markov chain Monte Carlo as you are about to learn.

What are conjugate priors?

Assume we have a likelihood and a prior:

 $\overbrace{[\theta|v]}^{\text{posterior}}$

$$\overline{[y]} = \frac{[y|\theta]}{[y]}$$

likelihood prior

If the form of the distribution of the posterior

 $[\boldsymbol{\theta}|y]$

is the same as the form of the distribution of the prior,

 $[\theta]$

then the likelihood and the prior are said to be conjugates $[y|\theta][\theta]$

congugates

and the prior is called a conjugate prior for θ .

Gibbs updates

When priors and likelihoods are conjugate, we *know* all but one of the parameters of the full-conditional because they are *assumed to be known* at each iteration. We make a draw of the single unknown *directly* from its posterior distribution as if the other parameters were fixed.

Wickedly clever.

Gamma-Poisson conjugate relationship for λ

The conjugate prior distribution for a Poisson likelihood is $gamma(\lambda | \alpha, \beta)$. Given *n* observations y_i of new data, the posterior distribution of λ is



Sampling from the Poisson-gamma conjugate:

Full conditional:

$$[\lambda_i \mid .] \propto \mathsf{Poisson}(y_i \mid \lambda_i) \mathsf{gamma}(\lambda_i \mid \alpha, \beta) \tag{4}$$

Gibbs sample:

$$[\lambda_i^k | y_i] = \operatorname{gamma} \left(\begin{array}{cc} \operatorname{The \ current} \alpha & \operatorname{The \ current} \beta \\ \hline \alpha^{k-1} & +y_i, & \overrightarrow{\beta^{k-1}} & +1 \end{array} \right).$$
(5)
In R, this would be:
shape[k] = alpha[k-1] + y_i
rate[k] = beta[k-1] + 1
lambda[k,i] = rgamma(1, shape[k], rate[k])

Gamma-gamma conjugate relationship

The conjugate prior distribution for the β parameter (rate) in a gamma likelihood gamma $(y_i | \alpha, \beta)$ is a gamma distribution gamma $(\beta | \alpha_0, \beta_0)$. Given *n* observations y_i of new data, the posterior distribution of β (assuming that α (shape) is known) is given by:



We can substitute any "known" quantity for $\mathbf{y},$ e.g., $\pmb{\lambda}$ in MCMC.

Sampling from the gamma-gamma conjugate:

The full conditional:

$$[\beta|.] \propto \prod_{i=1}^{n} \operatorname{gamma}(\lambda_{i}|\alpha,\beta) \operatorname{gamma}(\beta|.001,.001)$$

Gibbs sample:
$$eta^k \sim \mathsf{gamma}\left(.001 + n \alpha^{k-1}, .001 + \sum_{i=1}^n \lambda_i^k
ight)$$

In R this would be: shape[k] = .001 + length(y) * alpha[k-1] rate[k] = .001 + sum(lambda[,k])beta[k] = rgamma(1, shape[k], rate[k])

MCMC algorithm

- 1. Iterate over i = 1...197
- 2. At each *i*, make a draw from

$$\lambda_i^k \sim (7)$$

Gibbs update using gamma - Poisson conjugate for $each\lambda_i$

$$\beta^k \sim \operatorname{gamma}\left(.001 + \alpha^{k-1}n, .001 + \sum_{i=1}^n \lambda_i^k\right)$$
 (8)

Gibbs update using gamma - gamma conjugate for β

$$\alpha^{k} \propto \prod_{i=1}^{n} \operatorname{gamma}\left(\lambda_{i}^{k} | \alpha^{k-1}, \beta^{k}\right) \operatorname{gamma}\left(\alpha^{k-1} | .001, .001\right)$$
(9)

No conguate for α . Use Metropolis - Hastings update

3. Repeat for k = 1...K iterations, storing λ_i^k, α^k and β^k . Store the value of each parameter at each iteration in a vector.

Inference from MCMC

Make inference on each unobserved quantity using the elements of their vectors stored after convergence. These post-convergence vectors, (i.e., the "rug" described above) approximate the marginal posterior distributions of unobserved quantities.