


Gibbs sampling, continued



Sample $\theta_1^{(1)}$ from $p(\theta_1|\theta_2^{(0)}, \mathbf{y})$, then sample $\theta_2^{(1)}$ from $p(\theta_2|\theta_1^{(1)}, \mathbf{y})$, then repeat.

forms a Markov chain with stationary distribution .

For large n , $\theta^{(n)}$ has distribution close to $p(\theta|\mathbf{y})$, and more to the point, the sample $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(n)}$ 'looks like' a sample from $p(\theta|\mathbf{y})$.

77

Example: Normal Random Sample

Suppose data $\mathbf{y} = (y_1, \dots, y_n)$ are a random sample from

Assume independent priors on μ and ϕ : $(\theta = (\mu, \phi))$

$$\sim \text{Normal}(\gamma, \kappa^{-1})$$

$$\sim \text{Gamma}(\alpha, \beta)$$

The posterior distribution $p(\mu, \phi|\mathbf{y}) =$ is (up to a constant of proportionality)

$$\phi^{n/2} \exp\left(-(\phi/2) \sum_{i=1}^n (y_i - \mu)^2\right) \phi^{\alpha-1} e^{-\beta\phi}$$

78

Gathering all terms which contain μ gives

$$\propto \exp \left\{ -\frac{\phi}{2} \left(n\mu^2 - \sum y_i \mu \right) - \frac{\kappa}{2} (\mu^2 - \gamma) \right\}$$

so that (completing the square)

$$\mu | \phi, \mathbf{y} \sim \text{Normal} \left(\frac{\phi \sum y_i + \gamma \kappa}{n\phi + \kappa}, \frac{1}{n\phi + \kappa} \right)$$

Doing the same for ϕ gives the full conditional distribution

$$\sim \text{Gamma} \left(\alpha + n/2, \beta + \sum (y_i - \mu)^2 / 2 \right)$$

we implement the Gibbs sampler by alternately drawing μ and ϕ from these distributions.

79

Systematic or random scan?

You do not have to cycle round the components
e.g. you can choose which component to update next

Other variations?

You can update a subset of several components at once, for example, draw (θ_2, θ_3) from $\pi(\theta_2, \theta_3 | \theta_1, \theta_4, \dots, \theta_k)$.

How does the Gibbs sampler work?

- Each step has a $\mathcal{O}(k)$ cost: we will spell this out later when we consider a generalisation of the Gibbs sampler.
- It is usually (i.e. except in pathological cases) $\mathcal{O}(k)$.

The BUGS program

Bayesian inference Using Gibbs Sampling

- Language for specifying complex Bayesian models
- Constructs object-oriented internal representation of the model by identifying parents and children
- Builds up an complex model through specification of local structure
- Simulation from full conditionals using Gibbs sampling
- Current version (WinBUGS 1.4.1) runs in Windows, and incorporates the graphical model editor and a language for running in batch mode
- 'Classic' BUGS available for UNIX but this is an old version

81

WinBUGS is freely available from

<http://www.mrc-bsu.cam.ac.uk/bugs>

- An open source version of BUGS (called) is under development, and includes versions of BUGS that run under LINUX (LinBUGS) and that can be run directly from R (BRugs). See <http://www.rni.helsinki.fi/openbugs>

82

The Metropolis method

Sometimes it is to use the Gibbs sampler, e.g. when it is difficult to simulate a random value from a full conditional (e.g. if it is not a distribution of standard form – there’s an example on slide).

In the Metropolis method, when the current state is θ , you draw a *proposed* new value from an distribution $q(\theta'; \theta)$, that satisfies the symmetry condition $q(\theta'; \theta) =$.

Then you the new value with probability

$$= \min \left[1, \frac{\pi(\theta')}{\pi(\theta)} \right]$$

and otherwise stay in the state θ .

83

Note that you *always* accept a proposed move to a *more probable* state.

Very commonly, the proposal is chosen so that, as in the Gibbs sampler, only component of θ' is different from that of θ .

The Metropolis-Hastings method

You can generalise by the requirement that $q(\theta'; \theta) = q(\theta; \theta')$. Then the correct probability becomes

$$\alpha(\theta'; \theta) = \min \left[1, \frac{\pi(\theta')}{\pi(\theta)} \right]$$

84