STAT 530: In advance of Ch. 10

Mar. 17, 2010

Need a Monte Carlo technique that works when Gibbs Sampling is not applicable

Recall, we argued that when it is applicable, Gibbs sampling works because:

$\theta^{(s)} \longrightarrow \theta^{(s+1)}$ is Markov, stationary distribution equal to the posterior distribution of (θ).

Or deconstruct further:

For $(\theta_1^{(s+1)}, \ldots, \theta_{j-1}^{(s+1)}, \theta_{j+1}^{(s)}, \ldots, \theta_p^{(s)})$ fixed:

$\theta_j^{(s)} \longrightarrow \theta_j^{(s+1)}$ is Markov, stationary distribution equal to the posterior full conditional for $(\theta_j | \theta_{-j})$

So, seek other Markov chain 'recipes' (transitions) having the posterior (whole, or full cond.) as the stationary dist.

To explain, pretend everything is discrete/finite/univariate.

Markov transition (on state Z) represented as $T_{ij} = Pr(Z^{(t+1)} = j | Z^{(t)} = i)$

Distribution π is stationary for T if $\pi T = \pi$.

So we want to set π to be the posterior distribution and find a T for which this holds.

But first, two asides

Intuition for why Markov chain converges to its stationary distribution regardless of its initial distribution

Write initial distribution as linear combination of eigenvectors...

For a given π and T, how to verify that stationarity holds?

Show that **detailed balance** holds, i.e., for any a, b:

$$\pi_a T_{ab} = \pi_b T_{ba}$$