

Gibbs Sampling

Sargur N. Srihari
srihari@cedar.buffalo.edu

Topics in Monte Carlo Methods

1. Sampling and Monte Carlo Methods
2. Importance Sampling
3. Markov Chain Monte Carlo Methods
4. Gibbs Sampling
5. Mixing between separated modes

4. Gibbs Sampling

- The conceptually simplest approach for drawing samples from an undirected graph
- Suppose we have a graphical model over an n -dimensional vector of random variables \mathbf{x}
- We iteratively visit each variable x_i and draw a sample conditioned on all the other variables, i.e., from $p(x_i | \mathbf{x}_{-i})$
- Due to the separation properties of the graphical model, we can equivalently condition on only the neighbors of x_i

Gibbs Sampling

- Gibbs Sampling is an MCMC that samples each random variable of a PGM, one at a time
 - GS is a special case of the MH algorithm
- GS advantages
 - Are fairly easy to derive for many graphical models
 - e.g. mixture models, Latent Dirichlet allocation
 - Have reasonable computation and memory requirements, since they sample one r.v. at a time
 - Can be Rao-Blackwellized (integrate out some r.v.s)
 - to decrease sampling variance – called collapsed Gibbs sampling

The Gibbs Sampling Algorithm

- 1. Suppose the PGM contains variables x_1, \dots, x_n
- 2. Initialize starting values for x_1, \dots, x_n
- 3. Do until convergence:
 - 1. Order the n variables (can be fixed or random)
 - 2. For each variable x_i in order:
 - 1. Sample $x \sim p(x_i \mid x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$,
 - from conditional of x_i given current values of other variables
 - 2. Update $x_i \leftarrow x$
- When we update x_i , we immediately use its new value for sampling other variables x_j

Gibbs Sampling with M variables


- Initialize first sample: $\{z_i, i = 1, \dots, M\}$
- For $t=1, \dots, T$, $T = \text{no of samples}$
 - Sample $z_1^{(\tau+1)} \sim p(z_1 | z_2^{(\tau)}, z_3^{(\tau)}, \dots, z_M^{(\tau)})$
 - Sample $z_2^{(\tau+1)} \sim p(z_2 | z_1^{(\tau+1)}, z_3^{(\tau)}, \dots, z_M^{(\tau)})$
 -
 - Sample $z_j^{(\tau+1)} \sim p(z_j | z_1^{(\tau+1)}, \dots, z_{j-1}^{(\tau+1)}, z_{j+1}^{(\tau)}, \dots, z_M^{(\tau)})$
 -
 - Sample $z_M^{(\tau+1)} \sim p(z_M | z_1^{(\tau+1)}, z_2^{(\tau+1)}, \dots, z_{M-1}^{(\tau+1)})$
- $p(z_j | z_{-j})$ is called a *full conditional* for variable j

Markov Blanket to the rescue

- The conditional $p(x_i | x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ looks intimidating, but recall Markov Blankets
- Let $MB(x_i)$ be the Markov Blanket of x_i , then

$$p(x_i | x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) = p(x_i | MB(x_i))$$

- For a BN, the Markov Blanket of x_i is the set containing its parents, children, and co-parents




- For an MRF, the Markov Blanket of x_i is its immediate neighbors

Gibbs as a case of Metropolis-Hastings

- Gibbs sampling is a special case of MH with proposal $q(x'_i, x_{-i} | x_i, x_{-i}) = p(x'_i | x_{-i})$
 - It is easy check that the acceptance probability simplifies to one
- Assuming the right transition operator, both Gibbs and MH produce samples from their stationary distribution, which is p

Block Gibbs Sampling

- All visible units of an RBM may be sampled simultaneously because they are conditionally independent given all of the hidden units



- Likewise, all hidden units may be sampled simultaneously because they are conditionally independent given all of the visible units
- Gibbs sampling approaches that update many variables simultaneously in this way are called block Gibbs sampling

Prevalence of Gibbs Sampling

- Alternate MCMC approaches to sample from p_{model} are possible
 - E.g., M-H is widely used in other disciplines
- In deep learning for undirected modeling, it is rare to use any approach other than Gibbs
- Improved sampling techniques is a research frontier