

Gibbs sampling algorithm for mixture models

Given a set of n observations, $\mathbf{y} = (y_1, \dots, y_n)$, a mixture model with K components is usually defined as follows: $y_i \propto \sum_{j=1}^K \pi_j p(y_i | \theta_j)$, $i = 1, \dots, n$.

We use auxiliary variables $\mathbf{z} = (z_1, \dots, z_n)$ to indicate to which component observation i belongs. The prior distribution on \mathbf{z} is $p(\mathbf{z} | \pi)$ is defined as : $p(\mathbf{z} | \pi) = \prod_{j=1}^K \pi_j^{n_j}$, where n_j is the number of observations assigned to component j . Note that observations y_i are independent given z_i and θ_{z_i} : $p(\mathbf{y} | \mathbf{z}, \theta) = \prod_{i=1}^n p(y_i | \theta_{z_i})$.

1. Draw a graphical model that represents this mixture model. You should include variables $\pi, \theta, \mathbf{z} = (z_1, \dots, z_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$ in your graph.
2. Suppose we want to model traffic on a given road. At each hour i , we measure the number of cars y_i . We think the number of cars is well modeled by a Poisson distribution. However, we know that some times are rush hour, so we decide to use a mixture of Poisson distributions with two components (normal and rush hour traffic respectively.) For simplicity, imagine that traffic at each hour is independent of one another.

That is, we imagine each y_i is drawn from a Poisson with mean λ_j

$$p(y_i | z_i = j, \lambda_j) = \frac{\lambda_j^{y_j} e^{-\lambda_j}}{y_j!}.$$

We use a Gamma prior for λ :

$$p(\lambda_j) \sim \text{Gamma}(a, b) = \frac{b^a \lambda_j^{a-1} e^{-b\lambda_j}}{(a-1)!}.$$

We would like to do inference in this model using Gibbs sampling. Derive the full conditional distribution that you could use to sample λ_j . You do not need to calculate the normalization factor. Note: It may help to use $n_j = \sum_i I(z_i = j)$ to denote the number of observations assigned to component j and $s_j = \sum_i y_i I(z_i = j)$ the sum of their values.