Stat 535C - CS 535D Assignement #2, due in class on April, 6.

Hidden Markov Models

Consider the following model finite state-space hidden Markov model; i.e. the process $\{X_n\}$ is an unobserved/hidden finite state space Markov chain with K states such that

$$\Pr(X_1 = i) = \pi_i$$

and for $n > 1$

$$\Pr(X_n = j | X_{n-1} = i) = \pi_{i,j}.$$

We only have access to another process $\{Y_n\}$ such that, conditional upon $\{X_n\}$, the observations $\{Y_n\}$ are statistically independent and distributed. We consider here a simple case where the observations are counting data modelled as

$$Y_n | X_n = i \sim \text{Poisson}(\lambda_i).$$

The parameters $\theta = (\{\pi_i\}, \{\pi_{i,j}\}, \{\lambda_i\})$ are unknown. We (obviously) follow a Bayesian approach and set some conjugate priors on them; i.e.

$$\begin{aligned}(\pi_1, ..., \pi_K) &\sim \text{Dirichlet}(1, ..., 1),\\ (\pi_{i,1}, ..., \pi_{i,K}) &\sim \text{Dirichlet}(1, ..., 1),\\ \lambda_i &\sim \text{Gamma}(a_i, b_i)\end{aligned}$$

where (a_i, b_i) are fixed hyperparameters. It is strongly recommended to read [1] before starting the assignement. This paper can be downloaded on the course page.

1. Given T observations $y_{1:T}$, establish the expressions for the full conditional distributions $\pi(\theta | y_{1:T}, x_{1:T})$, $\pi(x_k | y_{1:T}, \theta, x_{1:k-1}, x_{k+1:T})$ and $\pi(x_{1:T} | y_{1:T}, \theta)$.

2. Implement the resulting Gibbs sampling algorithms, i.e. the "one-at-a time" Gibbs sampler which updates the hidden states using $\pi(x_k | y_{1:T}, \theta, x_{1:k-1}, x_{k+1:T})$ and the block Gibbs sampler which updates the hidden states simultaneously using $\pi(x_{1:T} | y_{1:T}, \theta)$. Apply these algorithms to the Fetal Lamb dataset downloadable on the course page. For sake of simplicity, consider only the cases where $K = 2$ and $K = 3$ using a prior ensuring that $E[\lambda_1] < ... < E[\lambda_K]$ (as in [1, p. 87]) and an exchangeable prior (i.e. $a_1 = ... = a_K = 1, b_1 = ... = b_K = 1$).

In the case of an exchangeable prior, discuss the problems faced by the Gibbs samplers and propose some strategies to improve the performance of the algorithm.

In the rest of the assignement, we will limit ourselves to the case where $K = 2$ and the prior is exchangeable.

3. Establish the expression of the marginal distribution $\pi\left(x_{1:T} \mid y_{1:T}\right)$ up to a normalizing constant where

$$\pi\left(x_{1:T} \mid y_{1:T}\right) = \int \pi\left(x_{1:T}, \theta \mid y_{1:T}\right) d\theta.$$

   Based on this expression, propose and implement an alternative MCMC algorithm to sample from $\pi\left(x_{1:T} \mid y_{1:T}\right)$. Make sure that your algorithm handles the label switching problem.

4. Propose and implement an MCMC algorithm to sample from $\pi\left(\theta \mid y_{1:T}\right)$ which does not required sampling the missing data $x_{1:T}$.

5. Propose and implement an SMC algorithm to sample from the sequence of distributions $\left\{\pi\left(x_{1:k} \mid y_{1:k}\right)\right\}_{k=1,\ldots,T}$.

Hand in any command scripts and programs you used, along with some informative output and plots, and your discussion.

References

[1] S. Chib, Calculating posterior distributions and modal estimates in Markov mixture models, *J. Econometrics*, vol. 75, pp. 79-97, 1996.