MSc/Part B

SB1b/M2 Computational Statistics
Hidden Markov models
Specimen questions

Date of this version: 12-04-2017

Do not turn this page until you are told that you may do so

1. Let $(X_0, X_1, Y_1, \ldots, X_T, Y_T)$ be a discrete-state hidden Markov model, with $X_i \in \{1, \ldots, K\}$ and $Y_i \in \{1, \ldots, M\}$ with joint probability mass function (pmf)

$$\begin{aligned} p(x_0, x_1, y_1, \ldots, x_T, y_T) &:= \Pr(X_0 = x_0, X_1 = x_1, Y_1 = y_1, \ldots, X_T = x_T, Y_T = y_T) \\ &= p(x_0) \prod_{t=1}^{T} p(x_t|x_{t-1}) p(y_t|x_t) \end{aligned}$$

   (i) Show carefully that, for any $t = 0, \ldots, T-1$

$$p(x_t|x_{t+1}, \ldots, x_T) = p(x_t|x_{t+1})$$

   (ii) Deduce that the joint pmf admits the following alternative decomposition

$$p(x_0, x_1, y_1, \ldots, x_T, y_T) = p(x_T) \prod_{t=T}^{1} p(x_{t-1}|x_t) p(y_t|x_t)$$

   (iii) Let $\mu_i = \Pr(X_T = i)$, $B_{i,k} := \Pr(Y_t = k|X_t = i)$ and $C_{i,j} := \Pr(X_{t-1} = j|X_t = i)$. Derive a backward recursion to calculate $p(y_{t:T}, x_t)$ for $t = T, T-1, \ldots, 1$. Justify carefully any conditional independence assumption.

   (iv) What is the computational complexity of this recursion?

2. Suppose you have a state-space model of the form

$$\begin{aligned} X_t &= \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} X_{t-1} + \begin{pmatrix} V_{t,1} \\ 2V_{t,2} \end{pmatrix} \\ Y_t &= \begin{pmatrix} 1 & 0 \end{pmatrix} X_t + \frac{1}{2} W_t \end{aligned}$$

where $V_{t,1}$, $V_{t,2}$ and W_t are independent standard normal random variables for each t. Assume that you have observed $Y_1 = y_1, \ldots, Y_{10} = y_{10}$ and you have used the Kalman filter to calculate $\mu_{10|10} := \mathbb{E}[X_{10}|y_{1:10}]$ and $\Sigma_{10|10} := cov[X_{10}|y_{1:10}]$, and obtain the values

$$\begin{aligned} \mu_{10|10} &= \begin{pmatrix} 1 \\ -1 \end{pmatrix} \\ \Sigma_{10|10} &= \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \end{aligned}$$

   (i) Calculate the state prediction $\mu_{11|10}$ and its covariance $\Sigma_{11|10}$

   (ii) Calculate the predicted observation $\widehat{y}_{11|10} := \mathbb{E}[Y_{11}|y_{1:10}]$ and its variance $S_{11} := Var(Y_{11}|y_{1:10})$

3. Let $(X_0, X_1, Y_1, \ldots, X_T, Y_T)$ be a discrete-state hidden Markov model, with $X_i \in \{1, \ldots, K\}$ and $Y_i \in \{1, \ldots, M\}$ with joint probability mass function (pmf)

$$\begin{aligned} p(x_0, x_1, y_1, \ldots, x_T, y_T) &:= \Pr(X_0 = x_0, X_1 = x_1, Y_1 = y_1, \ldots, X_T = x_T, Y_T = y_T) \\ &= p(x_0) \prod_{t=1}^{T} p(x_t|x_{t-1}) p(y_t|x_t) \end{aligned}$$

(i) Show that

$$p(x_{t-1}, x_t|y_{1:T}) = \frac{p(x_{t-1}, y_{1:t-1}) p(y_t|x_t) p(x_t|x_{t-1}) p(y_{t+1:T}|x_t)}{p(y_{1:T})}$$

(ii) Assume that you have used a forward-backward recursion to calculate $\alpha_t(x_t) := p(x_t, y_{1:t})$ and $\beta_t(x_t) := p(y_{t+1:T}|x_t)$ for each $t = 1, \ldots, T$. Describe an algorithm to simulate N realizations $x_{0:T}^{(i)}$, $i = 1, \ldots, N$, from the posterior distribution $p(x_{0:T}|y_{1:T})$.

(iii) What is the overall complexity of simulating realizations from $p(x_{0:T}|y_{1:T})$?

End of Last Page