UNIVERSITY OF CAMBRIDGE

MATHEMATICAL TRIPOS Part III

Friday, 8 June, 2012 9:00 am to 11:00 am

PAPER 38

TIME SERIES AND MONTE CARLO INFERENCE

Attempt no more than ***THREE*** *questions.*

There are ***FOUR*** *questions in total.*

The questions carry equal weight.

STATIONERY REQUIREMENTS
Cover sheet
Treasury Tag
Script paper

SPECIAL REQUIREMENTS
None

You may not start to read the questions printed on the subsequent pages until instructed to do so by the Invigilator.

1

(i). Let (X_t) be a stationary time series with autocovariance function $(\gamma_k : k = 0, 1, \ldots)$. Under what condition on (γ_k) does a spectral density function $f_X : [0, \pi] \to [0, \infty)$ exist? Assuming this condition is satisfied, write down the expressions for $f_X(\omega)$ in terms of the (γ_k), and for γ_k in terms of f_X.

(ii). The process (Y_t) is defined by:

$$Y_t = \sum_{s=-\infty}^{\infty} a_s X_{t-s}$$

where (a_t) is a sequence of real numbers such that $\sum_{s=-\infty}^{\infty} |a_s| < \infty$. Show that the spectral density function f_Y of Y is given by

$$f_Y(\omega) = |\alpha(\omega)|^2 f_X(\omega)$$

where $\alpha(\omega) = \sum_{s=-\infty}^{\infty} a_s e^{is\omega}$.

(iii). What is meant by saying that (X_t) is a *ARMA*(p, q) *process* with autoregressive coefficients $(\phi_1, \ldots, \phi_p)$ and moving average coefficients $(\theta_1, \ldots, \theta_q)$? Derive a formula for the spectral density function f_X of (X_t).

(iv). Show that the spectral density of a AR(1) process with autoregressive coefficient $\phi \in (-1, 1)$ is

$$f(\omega) = \frac{\sigma^2}{\pi(1 - \phi^2)} \sum_{k=-\infty}^{\infty} \phi^{|k|} z^k$$

where $z = e^{i\omega}$ and σ^2 is the variance of the driving white noise process.

(v). Hence, or otherwise, compute the autocovariance function of a ARMA$(1, 1)$ process having autoregressive coefficient ϕ and moving average coefficient θ.

2

The MA(1) process (X_t) is generated as

$$X_t = \theta\epsilon_{t-1} + \epsilon_t$$

where $\theta \in \mathbb{R}$ and (ϵ_t) is a white noise process with variance σ^2. What is the autocovariance function $(\gamma_k : k \in \mathbb{N})$ of (X_t)?

The sequence is observed from time 1. Uncorrelated variables (U_t), and scalars λ_t, ϕ_t $(t = 1, 2, \ldots)$, are determined recursively by:

$$\begin{aligned} U_1 &= X_1 && (1) \\ U_t &= X_t - \lambda_t U_{t-1} \quad (t = 2, 3, \ldots) && (2) \\ \phi_t &= \mathrm{var}(U_t). && (3) \end{aligned}$$

Let $\widehat{X}_t$ be the best linear predictor of X_t based on $(X_1, \ldots, X_{t-1})$. Justify the recursive formulae:

$$\begin{aligned} \lambda_t &= \gamma_1/\phi_{t-1} && (4) \\ \phi_t &= \gamma_0 - \gamma_1\lambda_t && (5) \\ \widehat{X}_t &= \lambda_t U_{t-1} && (6) \\ U_t &= X_t - \widehat{X}_t && (7) \end{aligned}$$

with the starting conditions $\lambda_1 = 0$, $U_0 \equiv 0$.

You may assume that $\lambda = \lim_{t\to\infty} \lambda_t$ exists and lies in $[-1, 1]$. Show that $\lambda = \theta$ if $|\theta| \leqslant 1$. What is λ when $|\theta| > 1$?

We observe $X_1 = -1.3$, $X_2 = 0.8$. Assuming $\sigma^2 = 1, \theta = -1$, find the best linear predictors of X_3 and X_4 based on these observations, and the variance of the prediction error in each case.

3

Describe, without proof, how the *polar rejection* method could be applied to generate a random variable Y from the chi-squared distibution with 1 degree of freedom, $Y \sim \chi^2_1$.

[*Hint: If* Z *has the standard normal distribution, and* $Y = Z^2$*, then* $Y \sim \chi^2_1$.]

The variable X is said to have the *inverse Gaussian distribution* with mean $\mu > 0$ and shape parameter $\lambda > 0$, and we write $X \sim IG(\mu, \lambda)$, if its probability density function is

$$f_X(x) = \left(\frac{\lambda}{2\pi x^3}\right)^{1/2} \exp\left\{\frac{-\lambda(x-\mu)^2}{2\mu^2 x}\right\} \qquad (x > 0). \tag{1}$$

Show that the following algorithm will generate a variable $X \sim IG(\mu, \lambda)$:

(i). Generate, independently, $Y \sim \chi^2_1$, and U, uniformly distributed on $[0, 1]$.
Note: Y *has density*

$$f_Y(y) = \frac{1}{\sqrt{2\pi}} y^{-\frac{1}{2}} \exp\left(-\tfrac{1}{2}y\right) \qquad (y > 0).$$

(ii). Define X_-, X_+ to be, respectively, the smaller and larger roots of the quadratic equation

$$\lambda(x-\mu)^2 - Y\mu^2 x = 0. \tag{2}$$

(iii). If $U \leqslant \mu/(\mu + X_-)$ set $X = X_-$.

(iv). Otherwise set $X = X_+$.

[*Hint: Compute* $\Pr\{X \in (x, x + dx)\}$*, treating the cases* $x \leqslant \mu$ *and* $x > \mu$ *separately. Note that* $X_- \times X_+ = \mu^2$*, and so* $X \leqslant \mu$ *if and only if* $X = X_-$*, while* $\mu/(\mu + X_-) = X_+/(\mu + X_+)$.]

Using equation (2), or otherwise, show that the distribution of $\lambda(X-\mu)^2/\mu^2 X$ is χ^2_1.

4

Describe the *Gibbs sampler* for generating a random configuration from the joint distribution P of variables $(X_1, \ldots, X_k)$. Show that P is a stationary distribution of the associated Markov chain.

A statistical model for a two-dimensional image is as follows. With each pixel (i, j) $(1 \leqslant i, j \leqslant N)$ in a $N \times N$ array is associated a random variable, X_{ij}, having possible values 0 and 1. Two pixels (i, j) and (i', j') are said to be *neighbours* if *either* $i = i'$ and $|j - j'| = 1$, *or* $j = j'$ and $|i - i'| = 1$. The probability of any configuration $\boldsymbol{x} = (x_{ij} : 1 \leqslant i, j \leqslant N)$ for the collection $\boldsymbol{X} = (X_{ij} : 1 \leqslant i, j \leqslant N)$ of all these variables has the form:

$$p(\boldsymbol{x}) = \exp(\lambda + \alpha n_1 + \beta n_2) \tag{1}$$

where λ, α, β are constants, n_1 is the total number of 1's in the configuration $\boldsymbol{x}$, and n_2 is the total number of pairs of neighbouring pixels (i, j), (i', j') such that $x_{ij} = x_{i'j'} = 1$.

Compute the probability p_{ij} that $X_{ij} = 1$, given the values of the variables at all other pixels, as a function of the total number z_{ij} of 1's at the pixels neighbouring (i, j). Hence devise a Gibbs sampler scheme for generating a random configuration from the distribution (1).

Suppose the pixels are coloured alternately black and white, as on a chessboard, so that neighbouring pixels are of opposite colours. Let $\boldsymbol{B}$ denote the collection of X's at black pixels, and $\boldsymbol{W}$ the collection of X's at white pixels. Describe the Gibbs sampler for generating a draw from the distribution of $(\boldsymbol{B}, \boldsymbol{W})$.

The image we actually observe is a noisy version $\boldsymbol{Y} = (Y_{ij})$ of $\boldsymbol{X}$: given $\boldsymbol{X} = \boldsymbol{x}$, the Y's are independent, Y_{ij} having the normal distribution with mean x_{ij} and variance 1. How would you modify the above Gibbs schemes to sample from the posterior distribution of $\boldsymbol{X}$, given $\boldsymbol{Y} = \boldsymbol{y}$?

END OF PAPER