Petter Mostad
Applied Mathematics and Statistics
Chalmers and GU

MSA101 / MVE187 Computational methods for Bayesian statistics

Re-exam 8 January 2019, 8:30 - 12:30
Allowed aids: None.
The appendix of this exam contains informaion about some probability distributions.
Total number of points: 30. To pass, at least 12 points are needed

1. (2 points) Define precisely a credibility interval and a confidence interval, and explain in detail the difference between the two.

2. (6 points)

   (a) Define the density $\pi(x) = \frac{\pi}{2}\cos\left(\frac{\pi}{2}x\right)$ for $x \in [0,1]$. Derive and describe exactly the most efficient way to simulate a random sample from this density.

   (b) Define the density $\pi(x) \propto \cos\left(\frac{\pi}{2}x\right)e^{-x}$ for $x \in [0,1]$. Explain exactly how to use rejection sampling to simulate a random sample from this density.

   (c) Define as above $\pi(x) \propto \cos\left(\frac{\pi}{2}x\right)e^{-x}$ for $x \in [0,1]$. Now explain exactly how to use slice sampling to obtain an (approximate) random sample from the density.

3. (3 points)

   (a) Let X be a continuous random variable and let $\pi(x)$ denote its density. Write down the definition of the entropy of X .

   (b) Assume X has an Exponential distribution with parameter λ. Compute its entropy.

4. (8 points) For real x, consider the model

$$x \mid \tau \sim \text{Normal}(0, \tau^{-1})$$

   with parameter $\tau > 0$.

   (a) Show that the family $\tau \sim \text{Gamma}(\alpha, \beta)$ is a conjugate family, and find the density for $\tau \mid x$.

   (b) Given[1] $x \mid \tau \sim \text{Normal}(0, \tau^{-1})$ and $\tau \sim \text{Gamma}(\alpha, \beta)$, find the marginal density $\pi(x)$ up to a constant not depending on x.

   (c) Find the name and parameters of the density you found in (b).

[1] In the original exam, this was written erroneously as $x \mid \tau \sim \text{Normal}(0, \tau)$

(d) Assume

$$\pi(\tau) = \frac{1}{2} \cdot \text{Gamma}(\tau; 3, 7) + \frac{1}{2} \cdot \text{Gamma}(\tau; 1, 1)$$

Find the formula for the posterior density $\pi(\tau \mid x)$. (You may express it using the density found in (b) and (c)).

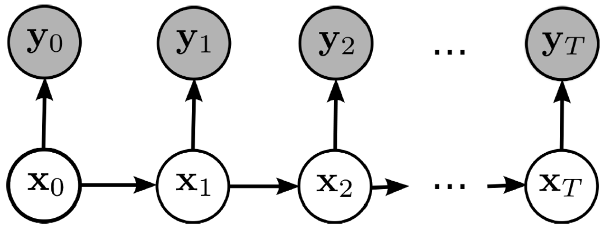

Figure 1: The hidden Markov model used in question 5.

5. (2 points) Consider the graphical model pictured in Figure 1, where $x_0, \ldots, x_T$ and $y_0, \ldots, y_T$ are random variables.

   (a) Assume you want to find the marginal distribution of x_2 given observations of $y_0, \ldots, y_T$. Name an algorithm that allows you to compute this.

   (b) Assume the variables $x_0, \ldots, x_T$ are discrete with a finite number of possible values. Name an algorithm that allows you to compute the value of x_2 in the vector $(x_0, \ldots, x_T)$ maximizing $\pi(x_0, \ldots, x_T \mid y_0, \ldots, y_T)$.

6. (8 points) Consider the following model, where $\theta \in [0, 1]$, and $z, y_1, y_2, \ldots, y_{14}$ are non-negative integers:

$$\begin{aligned} \theta &\sim \text{Uniform}(0, 1) \\ z \mid \theta &\sim \text{Binomial}(9, \theta) \\ y_1 \ldots, y_{14} \mid z, \theta &\sim \text{Poisson}(z) \end{aligned}$$

   where the last equation means that the y_i are *iid* Poisson(z) given z (and θ).

   (a) Compute and simplify the function $\ell(z, \theta) = \log\left[\pi(y_1, \ldots, y_{14}, z, \theta)\right]$.

   (b) Name the conditional distribution $\pi(\theta \mid z, y_1, \ldots, y_{14})$ and find its parameters.

   (c) Describe a possible way to make computations[2] to simulate from $\pi(z \mid \theta, y_1, \ldots, y_{14})$.

   (d) Describe how to use Gibbs sampling to generate an approximate sample from $\pi(z, \theta \mid y_1, \ldots, y_{14})$.

   (e) Describe the computational steps in an EM-algorithm which computes the θ maximizing the marginal posterior $\pi(\theta \mid y_1, \ldots, y_{14})$.

7. (1 point) What is Approximate Bayesian Computation (ABC)?

[2]In other words, how would you make computations for example in R? However, you do not need to make R code, just describe how the computations should be done.

Appendix: Some probability distributions

The Bernoulli distribution

If $x \in \{0, 1\}$ has a Bernoulli distribution with parameter $0 \leq p \leq 1$, then the probability mass function is

$$\pi(x) = p^x(1-p)^{1-x}.$$

We write $x \mid p \sim \text{Bernoulli}(p)$ and $\pi(x \mid p) = \text{Bernoulli}(x; p)$.

The Beta distribution

If $x \geq 0$ has a Beta distribution with parameters $\alpha > 0$ and $\beta > 0$ then the density is

$$\pi(x \mid \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}.$$

We write $x \mid \alpha, \beta \sim \text{Beta}(\alpha, \beta)$ and $\pi(x \mid \alpha, \beta) = \text{Beta}(x; \alpha, \beta)$.

The Beta-Binomial distribution

If $x \in \{0, 1, 2, \ldots, n\}$ has a Beta-Binomial distribution, with n a positive integer and parameters $\alpha > 0$ and $\beta > 0$, then the probability mass function is

$$\pi(x \mid n, \alpha, \beta) = \binom{n}{x} \frac{\Gamma(x+\alpha)\Gamma(n-x+\beta)\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(n+\alpha+\beta)}.$$

We write $x \mid n, \alpha, \beta \sim \text{Beta-Binomial}(n, \alpha, \beta)$ and $\pi(x \mid n, \alpha, \beta) = \text{Beta-Binomial}(x; n, \alpha, \beta)$.

The Binomial distribution

If $x \in \{0, 1, 2, \ldots, n\}$ has a Binomial distribution, with n a positive integer and parameter $0 \leq p \leq 1$, then the probability mass function is

$$\pi(x \mid n, p) = \binom{n}{x} p^x(1-p)^{n-x}.$$

We write $x \mid n, p \sim \text{Binomial}(n, p)$ and $\pi(x \mid n, p) = \text{Binomial}(x; n, p)$.

The Exponential distribution

If $x \geq 0$ has an Exponential distribution with $\lambda > 0$ as parameter, then the density is

$$\pi(x \mid \lambda) = \lambda \exp(-\lambda x).$$

We write $x \mid \lambda \sim \text{Exponential}(\lambda)$ and $\pi(x \mid \lambda) = \text{Exponential}(x; \lambda)$.

The Gamma distribution

If $x > 0$ has a Gamma distribution, with parameters $\alpha > 0$ and $\beta > 0$, then the density is

$$\pi(x \mid \alpha\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp(-\beta x).$$

We write $x \mid \alpha, \beta \sim \text{Gamma}(\alpha, \beta)$ and $\pi(x \mid \alpha, \beta) = \text{Gamma}(x; \alpha, \beta)$.

The Geometric distribution

If the non-negative integer x has a Geometric distribution with parameter $p \in [0, 1]$, its probability mass function is given by

$$\pi(x \mid p) = (1 - p)^x p.$$

We write $x \mid p \sim \text{Geometric}(p)$ and $\pi(x \mid p) = \text{Geometric}(x; p)$.

The Logistic distribution

If x has a Logistic distribution, with parameters μ and $s > 0$, then the density is

$$\pi(x \mid \mu, s) = \frac{\exp\left(-\frac{x-\mu}{s}\right)}{s\left(1 + \exp\left(-\frac{x-\mu}{s}\right)\right)^2}.$$

We write $x \mid \mu, s \sim \text{Logistic}(\mu, s)$ and $\pi(x \mid \mu, s) = \text{Logistic}(x; \mu, s)$.

The Normal distribution

If the real x has a Normal distribution with parameters μ and σ^2, its density is given by[3]

$$\pi(x \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x - \mu)^2\right).$$

We write $x \mid \mu, \sigma^2 \sim \text{Normal}(\mu, \sigma^2)$ and $\pi(x \mid \mu, \sigma^2) = \text{Normal}(x; \mu, \sigma^2)$.

The Poisson distribution[4]

If the non-negative integer x has a Poisson distribution wiht parameter $\lambda > 0$, its probability mass function is

$$\pi(x \mid \lambda) = e^{-\lambda} \frac{\lambda^x}{x!}.$$

We write $x \mid \lambda \sim \text{Poisson}(\lambda)$ and $\pi(x \mid \lambda) = \text{Poisson}(x; \lambda)$.

[3]In the original exam, the density was given erroneously as $\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{\sigma^2}(x - \mu)^2\right)$.

[4]In the original exam, this distribution was missing.

The t distribution

If the real x has a t distribution with parameter ("degrees of freedom") $\nu > 0$, its density is given by

$$\pi(x \mid \nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}.$$

We write $x \mid \nu \sim \mathrm{t}(\nu)$ and $\pi(x \mid \nu) = \mathrm{t}(x; \nu)$.

The non-standardized t distribution

If the real x has a non-standardised t distribution with parameters $\nu > 0$, μ and σ^2, then its density is given by

$$\pi(x \mid \nu, \mu, \sigma^2) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi\sigma^2)}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{1}{\nu} \cdot \frac{(x-\mu)^2}{\sigma^2}\right)^{-\frac{\nu+1}{2}}$$

We write $x \mid \nu, \mu, \sigma^2 \sim \mathrm{t}(\nu, \mu, \sigma^2)$ and $\pi(x \mid \nu, \mu, \sigma^2) = \mathrm{t}(x; \nu, \mu, \sigma^2)$.

The Uniform distribution

If $x \in [a, b]$ has a Uniform(a, b) distribution with parameters $b > a$, then the density is given by

$$\pi(x \mid a, b) = \frac{1}{b-a}.$$

We write $x \mid a, b \sim \mathrm{Uniform}(a, b)$ and $\pi(x \mid a, b) = \mathrm{Uniform}(x; a, b)$.