Queen Mary University of London

Main Examination period 2019

MTH709U / MTH776P: Bayesian Statistics

Duration: 3 hours

Apart from this page, you are not permitted to read the contents of this question paper until instructed to do so by an invigilator.

You should attempt ALL questions. Marks available are shown next to the questions.

Only non-programmable calculators that have been approved from the college list of non-programmable calculators are permitted in this examination. Please state on your answer book the name and type of machine used.

Statistical functions provided by the calculator may be used provided that you state clearly where you have used them.
The New Cambridge Statistical Tables 2nd Eddition are provided.
A table of common distributions is provided as an appendix.

Complete all rough work in the answer book and cross through any work that is not to be assessed.

Possession of unauthorised material at any time when under examination conditions is an assessment offence and can lead to expulsion from QMUL. Check now to ensure you do not have any unauthorised notes, mobile phones, smartwatches or unauthorised electronic devices on your person. If you do, raise your hand and give them to an invigilator immediately.

It is also an offence to have any writing of any kind on your person, including on your body. If you are found to have hidden unauthorised material elsewhere, including toilets and cloakrooms, it will be treated as being found in your possession. Unauthorised material found on your mobile phone or other electronic device will be considered the same as being in possession of paper notes. A mobile phone that causes a disruption in the exam is also an assessment offence.

Exam papers must not be removed from the examination room.

Examiners: L I Pettit, J Griffin

Question 1. [21 marks]

(a) Show that a Poisson distribution with mean θ is an exponential family. [**4**]

(b) If a random sample of n Poisson distributed observations is collected write down a sufficient statistic for θ. [**1**]

(c) It is believed that the number of accidents in a new factory will follow a Poisson distribution with mean θ per month. The prior distribution of θ is given by a gamma distribution $Ga(\alpha, \beta)$.

(i) A safety inspector assesses that based on his experience at a similar factory $\alpha = 12, \beta = 4$. If there are 18 accidents in the first eight months, derive the posterior distribution of θ and find its mean and variance. [**5**]

(ii) Show the posterior mean can be written as a weighted average of the prior mean and the sample mean. [**3**]

(d) A manager believes that the new factory has new safety features and the experience from the other factory is not relevant. He suggests using a Jeffreys' prior.

(i) Find the Jeffreys' prior. [**4**]

(ii) Find the posterior mean using the Jeffreys' prior and comment on the difference between the two posterior means. [**4**]

Question 2. [20 marks] A horticulturalist is interested in the probability θ that a seed of a particular variety germinates successfully. Her prior distibution for the germination probability can be represented by the $Be(a, b)$ distribution. In an experiment she sows n seeds of which x germinate successfully.

(a) Find her posterior distribution for θ. [**3**]

(b) If she decides that $a = 3, b = 1$ and she observes $x = 7$ successes with $n = 10$ seeds, find her posterior distribution for θ. [**1**]

(c) If she wishes to estimate θ using a quadratic loss function

$$l(t, \theta) = (t - \theta)^2$$

derive the Bayes estimate of θ and the expected loss and calculate their values. [**6**]

(d) What is her predictive probability that all the seeds in a further batch of 10 would germinate? [**4**]

(e) She wishes to set up a further experiment using a batch of m seeds. Show that her probability for one or more seeds germinating is

$$1 - \frac{(m+3) \times (m+2) \times \cdots \times 4}{(m+13) \times (m+12) \times \cdots \times 14}.$$

Hence show that she would require $m \geq 5$ to ensure that the probability of one or more seeds germinating is at least 0.99. [**6**]

Question 3. [20 marks]

(a) A random sample of failure times $t_1, \ldots, t_n$ is observed for n machine components. Each t_i is assumed to have an exponential distribution with mean λ^{-1} and the prior distribution for λ is taken as $Ga(\alpha, \beta)$ with parameters α and β. Find the posterior distribution of λ. [**3**]

(b) Show that for the situation described in (a), the marginal likelihood (or prior predictive density) is given by

$$p(t_1, \ldots, t_n) = \frac{\beta^\alpha \Gamma(\alpha + n)}{\Gamma(\alpha)(\beta + S_n)^{\alpha+n}}$$

where $S_n = \sum_{i=1}^n t_i$. [**5**]

(c) Suppose $n = 10$, $\sum t_i = 50$, $\alpha = 5$ and $\beta = 20$. Find the Bayes factor to test the null hypothesis that $\lambda = \alpha/\beta$ against an alternative that λ has a $Ga(\alpha, \beta)$ prior. What is your conclusion? [**8**]

(d) What would be the effect on your answer to part (a) if the last of the components had not in fact failed and you had a censored value (t_n) for it. [**4**]

Question 4. [14 marks]

(a) A random sample $x_1, x_2, \ldots, x_n$ is taken from a distribution with density $p(x|\theta)$. Derive the form of the posterior distribution if the prior distribution is taken as a mixture of two priors $p_1(\theta)$ and $p_2(\theta)$ with weights w and $1-w$, that is

$$p(\theta) = wp_1(\theta) + (1-w)p_2(\theta).$$

[**4**]

(b) Define a $(1-\alpha) \times 100\%$ **credible interval**. [**2**]

(c) The plot below shows the density of the posterior distribution of parameter θ. Comment on what this shows. [**3**]

(d) Make a sketch of this posterior and include an example of a $100(1-\alpha)\%$ highest posterior density interval, this need not be to scale but should show the necessary properties such an interval has. [**3**]

(e) What is the advantage of a credible interval over a classical confidence interval? [**2**]

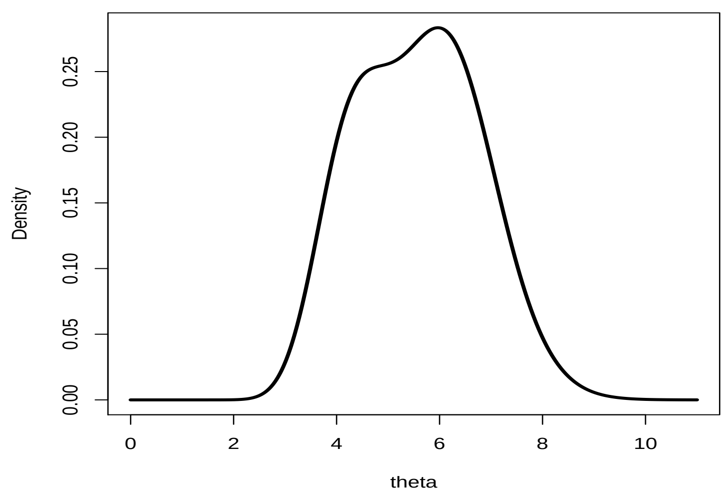

Figure 1: Plot of posterior density.

Question 5. [25 marks]

(a) Consider a two-stage linear model

$$\begin{aligned} \underline{y}|\underline{\theta}_1 &\sim N(A_1\underline{\theta}_1, C_1) \\ \underline{\theta}_1 &\sim N(\underline{\mu}, C_2) \end{aligned}$$

where A_1, C_1, C_2 and $\underline{\mu}$ are known.

Show that the posterior distribution of $\underline{\theta}_1$ is $N(B\underline{b}, B)$ where

$$\begin{aligned} B^{-1} &= A_1^T C_1^{-1} A_1 + C_2^{-1}, \\ \underline{b} &= A_1^T C_1^{-1} \underline{y} + C_2^{-1} \underline{\mu}. \end{aligned}$$

[**5**]

(b) Observations are taken from a model

$$y_{ij} \sim No(\alpha_i, \xi) \qquad \text{for } i = 1, 2 \text{ and } j = 1, \ldots, n_i.$$

Note that ξ is the precision. The priors for α_1 and α_2 are respectively $No(2, 2\xi)$ and $No(6, 4\xi)$ and they are assumed independent. The prior for ξ is $Ga(a/2, b/2)$.

(i) Find the posterior distributions of α_1 and α_2 given ξ, $p(\alpha_1|\underline{y}, \xi)$ and $p(\alpha_2|\underline{y}, \xi)$. [**7**]

(ii) Find the posterior distribution of ξ given α_1 and α_2, $p(\xi|\underline{y}, \alpha_1, \alpha_2)$. [**5**]

(iii) Explain how samples from the unconditional posterior distributions of α_1, α_2 and ξ can be simulated using Gibbs sampling. [**5**]

(iv) Given these samples how would you estimate the posterior median of $\alpha_1 - \alpha_2$? [**3**]

End of Paper – An appendix of 2 pages follows.

Bayesian Statistics – Common Distributions

Discrete Distributions

Distribution	Density	Range of Variates	Mean	Variance
Uniform	$\frac{1}{N}$	$N = 1, 2, \ldots$ $x = 1, 2, \ldots, N$	$\frac{N+1}{2}$	$\frac{N^2-1}{12}$
Bernoulli	$p^x(1-p)^{1-x}$	$0 \le p \le 1,\ x = 0, 1$	p	$p(1-p)$
Binomial	$\binom{n}{x}p^x(1-p)^{n-x}$	$0 \le p \le 1,\ n = 1, 2, \ldots$ $x = 0, 1, \ldots n$	np	$np(1-p)$
Poisson	$\frac{\exp(-\lambda)\lambda^x}{x!}$	$\lambda > 0,\ x = 0, 1, 2, \ldots$	λ	λ
Geometric	$p(1-p)^x$	$0 < p \le 1,\ x = 0, 1, 2, \ldots$	$\frac{(1-p)}{p}$	$\frac{(1-p)}{p^2}$
Negative Binomial	$\binom{r+x-1}{x}p^r(1-p)^x$	$0 < p \le 1,\ r > 0$ $x = 0, 1, 2, \ldots$	$\frac{r(1-p)}{p}$	$\frac{r(1-p)}{p^2}$

Continuous Distributions

Distribution	Density	Range of Variates	Mean	Variance
Uniform	$\frac{1}{b-a}$	$-\infty < a < b < \infty$ $a < x < b$	$\frac{a+b}{2}$	$\frac{(b-a)^2}{12}$
Normal $N(\mu, \sigma^2)$	$\frac{1}{\sqrt{2\pi\sigma^2}} \exp[\frac{-(x-\mu)^2}{2\sigma^2}]$	$-\infty < \mu < \infty$ $\sigma > 0,\ -\infty < x < \infty$	μ	σ^2
Normal $No(\mu, h)$	$\frac{\sqrt{h}}{\sqrt{2\pi}} \exp[\frac{-h(x-\mu)^2}{2}]$	$-\infty < \mu < \infty$ $h > 0,\ -\infty < x < \infty$	μ	h^{-1}
Exponential	$\lambda \exp(-\lambda x)$	$\lambda > 0,\ x \ge 0$	$\frac{1}{\lambda}$	$\frac{1}{\lambda^2}$
Gamma $Ga(\alpha, \beta)$	$\frac{\beta^\alpha x^{\alpha-1}\exp(-\beta x)}{\Gamma(\alpha)}$	$\alpha > 0,\ \beta > 0,\ x > 0$	$\frac{\alpha}{\beta}$	$\frac{\alpha}{\beta^2}$

Distribution	Density	Range of Variates	Mean	Variance
Beta $Be(a,b)$	$\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1}(1-x)^{b-1}$	$a>0,\ b>0,\ 0<x<1$	$\frac{a}{a+b}$	$\frac{ab}{(a+b+1)(a+b)^2}$
$t_\nu(m,g)$	$\frac{g^{1/2}\Gamma((\nu+1)/2)}{\sqrt{(\nu\pi)}\Gamma(\nu/2)}$	$-\infty<x<\infty$	location m	precision g,
	$\times\left[1+\frac{g}{\nu}(x-m)^2\right]^{-(\nu+1)/2}$	dof ν		
F_n^m	$\frac{\Gamma[(m+n)/2]}{\Gamma(m/2)\Gamma(n/2)}\left(\frac{m}{n}\right)^{\frac{m}{2}}$	$m,n=1,2,\ldots$	$\frac{n}{n-2}$	$\frac{2n^2(m+n-2)}{m(n-2)^2(n-4)}$
	$\times\frac{x^{(m-2)/2}}{[1+(m/n)x]^{(m+n)/2}}$	$x\geq 0$	for $n>2$	for $n>4$
χ_k^2	$\frac{1}{\Gamma(k/2)2^{k/2}} x^{k/2-1}\exp(-\frac{x}{2})$	$k=1,2,\ldots,\ x>0$	k	$2k$
Pareto	$\frac{\alpha\beta^\alpha}{x^{\alpha+1}}$	$\alpha>0,\beta>0,x>\beta$	$\frac{\beta\alpha}{(\alpha-1)}$	$\frac{\beta^2\alpha}{(\alpha-1)^2(\alpha-2)}$

End of Appendix.