MATH 559: Bayesian Theory and Methods

December 13th, 2023
9.00am – 12.00pm.

MATH 559: Specific Instructions.

Please write your answers on the exam in the boxes provided; only answers written inside the boxes will be marked. Please write clearly. You may request a regular blank answer booklet (not a Crowdmark booklet) for rough working, but answers written on that booklet will not be marked.

This paper contains five questions. Each question carries 20 marks. Credit will be given for all questions attempted. The total mark available is 100 but rescaling of the final mark may occur.

A formula sheet is provided. No crib sheet is allowed.

Only non-programmable calculators may be used. Dictionaries and Translation dictionaries are permitted.

The usual abbreviations are used: pmf indicates the probability mass function; pdf indicates the probability density function; cdf indicates the cumulative distribution function; i.i.d. indicates independent and identically distributed. The function $\mathbb{1}_A(x)$ is the indicator function for set A that takes the value 1 if $x \in A$ and zero otherwise.

This page intentionally left blank.

1. Suppose $\{Y_n\}$ is an infinitely exchangeable sequence. For $n \geq 1$, suppose $Y_1, \ldots, Y_n$ is a finite number of elements from this sequence. Then if $f_{Y_1,\ldots,Y_n}(y_1, \ldots, y_n)$ is the *joint pdf* for $Y_1, \ldots, Y_n$, it can be deduced that

$$f_{Y_1,\ldots,Y_n}(y_1, \ldots, y_n) = \int \prod_{i=1}^{n} f(y_i; \theta)\ \pi_0(\theta)\, d\theta. \tag{1}$$

where $\pi_0(\theta)$ is a *prior density* for the unknown finite-dimensional parameter θ, and $f(y; \theta)$ is a conditional pdf in y.

(a) Using (1), derive an expression for the *posterior predictive distribution* for Y_{n+1} (another element of the infinitely exchangeable sequence) conditional on $Y_1 = y_1, \ldots, Y_n = y_n$, and explain how this defines the *posterior density*, $\pi_n(\theta)$, as an updated version of $\pi_0(\theta)$.

6 MARKS

Q1(a):

(b) Suppose that in (1), it is specified that

$$f(y;\theta) = \mathbb{1}_{(0,\infty)}(y)\ \theta \exp\{-\theta y\} \quad y \in \mathbb{R}$$

where $\theta \in \Theta \equiv \mathbb{R}^+$. Suppose that data $y_1, \ldots, y_n$ are observed.

(i) Find the form of the posterior distribution if a *conjugate prior* is specified. 5 MARKS

Q1(b)(i):

(ii) Find the *Jeffreys prior* for this model. 4 MARKS

Q1(b)(ii):

(iii) Find the *posterior predictive density* for Y_{n+1} under a conjugate prior. 5 MARKS

Q1(b)(iii):

2. Suppose that $Y_1, \ldots, Y_n$ are presumed conditionally independent and distributed as $Uniform(0, \theta)$ for some $\theta \in \Theta \equiv \mathbb{R}^+$, that is

$$f_Y(y; \theta) = \frac{\mathbb{1}_{(0,\theta)}(y)}{\theta} \qquad y \in \mathbb{R}$$

Bayesian inference for θ is to be carried out. A prior $\pi_0(\theta)$ with support Θ will be specified.

(a) Find the form of likelihood

$$\mathcal{L}_n(\theta) = \prod_{i=1}^{n} f_Y(y_i; \theta)$$

for observed data $y_1, \ldots, y_n$. 6 MARKS

Q2(a):

(b) Show that the statistic

$$T \equiv T(Y_1, \ldots, Y_n) = \max\{Y_1, \ldots, Y_n\}$$

is *sufficient in the Bayesian sense*, that is, that the posterior distribution depends on the data only through the observed value of T. 4 MARKS

Q2(b):

(c) Let $\psi = 1/\theta$, and suppose the prior density for ψ takes the form

$$\pi_0(\psi) = c\psi \qquad 0 < \psi < 1$$

and zero otherwise, for some constant $c > 0$. Find the posterior for ψ. 6 MARKS

Q2(c):

(d) For the posterior in (c), find the *posterior mode*. 4 MARKS

Q2(d):

3. Suppose that, for $n \geq 1$.

$$Y_{11}, \ldots, Y_{1n} \sim Normal(\mu_1, 1)$$
$$Y_{21}, \ldots, Y_{2n} \sim Normal(\mu_2, 1)$$

are conditionally independent random variables. Observed data $y_{11}, \ldots, y_{1n}$ and $y_{21}, \ldots, y_{2n}$ are to be used for Bayesian inference.

(a) Suppose that $\pi_0(\mu_1) \equiv Normal(\eta, 1)$. Show that the posterior for μ_1 is also a Normal distribution. 6 MARKS

Q3(a):

(b) Suppose that

$$\pi_0(\mu_1, \mu_2) = \pi_0(\mu_1)\pi_0(\mu_2)$$

where $\pi_0(\mu_1) \equiv \pi_0(\mu_2) \equiv Normal(\eta, 1)$. Find the posterior for

$$\phi = \mu_2 - \mu_1$$

based on the observed data. 8 MARKS

Q3(b):

(c) Suppose that only the data $z_1, \ldots, z_n$

$$z_i = y_{2i} - y_{1i} \qquad i = 1, \ldots, n$$

are observed. Find the posterior for ϕ as defined in (b) based on $z_1, \ldots, z_n$ under the implied prior for ϕ specified in (b). 6 MARKS

Q3(c):

4. Suppose that in a Bayesian model, we have that

$$f_Y(y;\theta) = \mathbb{1}_{(0,\infty)}(y)\frac{y}{\theta}\exp\left\{-\frac{y^2}{2\theta}\right\} \quad y \in \mathbb{R}$$

for $\theta \in \Theta \equiv \mathbb{R}^+$.

(a) Find the posterior, $\pi_n(\theta)$ based on a sample $y_1, \ldots, y_n$, drawn conditionally independently from $f_Y(y;\theta)$, if

$$\pi_0(\theta) \equiv InvGamma(a_0, b_0/2).$$

for hyperparameters $a_0, b_0 > 0$. 8 MARKS

Q4(a):

(b) Suppose $\lambda = 1/\theta$ is to be estimated using Bayesian estimation under *quadratic loss*, that is

$$L(t, \lambda) = (t - \lambda)^2.$$

Find the Bayesian estimate of λ. 6 MARKS

Q4(b):

(c) Describe one procedure to derive a 95% *credible interval* for θ. 6 MARKS

Q4(c):

5. (a) Describe how to carry out Monte Carlo sampling from the pdf

$$f(x) = c\mathbb{1}_{(0,\infty)}(x)\sqrt{x}\exp\left\{-\frac{1}{2}x^2\right\} \qquad x \in R$$

using the following approaches.

(i) *Rejection sampling:* give a specific recommendation for the proposal distribution $f_0(x)$, and give details of the sampling procedure. 5 MARKS

Q5(a)(i):

(ii) *The Metropolis-Hastings algorithm:* give a specific recommendation for the proposal density $q(x, \cdot)$, and give details of the sampling procedure. 5 MARKS

Q5(a)(ii):

(b) The Bayesian analysis of the Generalized Linear Model with $Y_i \sim Poisson(\mu_i)$ and

$$\log \mu_i = \beta_0 + \beta_1 x_i$$

for outcomes $Y_1, \ldots, Y_n$ which are conditionally independent given predictor values $x_1, \ldots, x_n$ and parameters (β_0, β_1) is to be considered.

(i) Write down the joint posterior distribution $\pi_n(\beta_0, \beta_1)$ up to proportionality for a suitably chosen prior distribution. 5 MARKS

Q5(b)(i):

(ii) Describe one method for performing Monte Carlo sampling from $\pi_n(\beta_0, \beta_1)$, giving details of each step of the approach. 5 MARKS

Q5(c)(ii):

Additional space for answers or rough working

Additional space for answers or rough working

Additional space for answers or rough working

Additional space for answers or rough working