UC Berkeley
Department of Statistics

STAT 210A: Introduction to Mathematical Statistics

Problem Set 5
Fall 2007

Issued: Thursday, September 27 **Due:** Thursday, October 4

Reading: For this problem set: Chapter 9 of Keener (Bayesian methods); §3.1, 3.2 of Bickel and Doksum.

Problem 5.1
Find conjugate distributions for the following family of distributions:

(a) the gamma family of distribution functions $\Gamma(\lambda, p)$

(b) the beta family of distribution functions $\beta(\lambda, p)$

(c) the family defined on the parameters β_1 and β_2 by the linear regression $y_i = \beta_1 + \beta_2 x_i + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ and σ^2 known.

Problem 5.2
This problem addresses the issue of implementing Bayes estimators for exponential family models. Suppose that we have a (conditional) exponential family model

$$p(x \mid \theta) \quad = \quad h(x) \; \exp\left\{ \sum_{i=1}^{d} \theta_i T_i(x) - A(\theta) \right\},$$

where $x = (x_1, \ldots, x_n)$ and the random vector Θ has density $\lambda(\cdot)$.

(a) Define the marginal density $m(x) = \int p(x;\theta)\lambda(\theta) d\theta$ induced by this Bayesian model. Show that for $j = 1, \ldots, n$, we have

$$\mathbb{E}\left[\sum_{i=1}^{d} \Theta_i \frac{\partial T_i(x)}{\partial x_j} \;\mid\; x \right] \quad = \quad \frac{\partial}{\partial x_j} \log m(x) - \frac{\partial}{\partial x_j} \log h(x).$$

(Assume here that all relevant quantities are suitably differentiable.)

(b) Suppose that $X = (X_1, \ldots, X_n)$ has density $p(x; \theta) = h(x) \; \exp\left(\sum_{i=1}^{d} \theta_i x_i - A(\theta) \right)$. Use part (a) to conclude that the Bayes estimator δ_λ of θ_j under quadratic loss is given by

$$\delta_\lambda(x) \quad = \quad \frac{\partial}{\partial x_j} \log m(x) - \frac{\partial}{\partial x_j} \log h(x).$$

(c) Apply your result from (b) to rederive the Bayes estimator under quadratic loss for the normal-normal model with $X_i \mid \theta \sim N(\theta, \sigma^2), i = 1, \ldots n$ i.i.d., and $\Theta \sim N(\mu, b^2)$.

Problem 5.3
Consider the i.i.d. linear observation model

$$Y_i = x_i^T \beta^* + W_i, \qquad i = 1, \ldots n \tag{1}$$

where $\beta^* \in \mathbb{R}^p$, the design vectors $x_i \in \mathbb{R}^p$ are fixed and known, and $W_i \sim N(0, \sigma^2)$ is observation noise. Assume that $p < n$.

(a) Assume that $\sigma^2 > 0$ is known, and that β^* is modeled as a fixed but unknown vector. Find the least-squares estimate of β^* based on Y.

(b) Now assume that β is distributed according to the prior $N(\beta_0, \sigma^2 Q)$, where $Q \in \mathbb{R}^{p \times p}$ is a known, positive definite symmetric matrix. Find the Bayes estimator of $z^T\beta$ under least-squares loss where $z \in \mathbb{R}^p$ is a known fixed vector.

(c) Suppose that the minimum eigenvalue of $\sum_{i=1}^n x_i x_i^T$ tends to ∞. Under the sampling model (1), show that the estimator from part (b) converges to $z^T\beta^*$ as $n \to +\infty$.

Problem 5.4
Let $\Theta = (0, +\infty)$ and $\mathcal{A} = [0, \infty)$, and suppose that $X \sim \text{Poi}(\theta)$. Consider the loss function $L(\theta, a) = (\theta - a)^2/\theta$.

(a) Find Bayes estimators with respect to the family of Gamma(a, b) priors.

(b) Show that the estimator $\delta(X) = X$ can be obtained by a suitable limit from (a).

Problem 5.5
Let $(X_1, \ldots, X_n)$ be an i.i.d. sample from the uniform distribution on $(0, \theta)$, where $\theta > 0$ is unknown. Suppose that the prior distribution of θ is log-normal with parameters (μ_0, σ_0^2) where $\mu_0 \in \mathbb{R}$ and $\sigma_0^2 > 0$ are known constants.

(a) Find the posterior density of $\log \theta$.

(b) Suppose that we are interested in estimating θ under the loss function

$$L(\delta, \theta) = \begin{cases} 0 & \text{if } \delta = \theta \\ 1 & \text{otherwise.} \end{cases}$$

Find the Bayes estimator of θ under this loss function. (*Hint*: Part (a) is related.)

Problem 5.6
Consider the Bayesian model in which Θ has distribution Λ, and conditioned on $\Theta = \theta$, the random variable X has distribution $\mathbb{P}_\theta$. Suppose that we are interested in estimating $g(\theta)$ under quadratic loss. Prove that no unbiased estimator $\delta(X)$ of $g(\theta)$ can be a Bayesian estimator unless the Bayesian risk $r(\Lambda, \delta) = 0$. This shows that Bayes estimators and unbiased estimators agree only in pathological cases.