STK 4021 - Applied Bayesian Analysis and Numerical Methods
Thordis L. Thorarinsdottir
Fall 2014

Problem Set 11

Problem 1 (Bayes rule). In the decision theoretic framework, suppose that the parameter space Θ is an open subset of $\mathbb{R}^k$ and that the sampling distribution P_θ has density function $p(x|\theta)$ with respect to some σ-finite measure ν. Let $p(\theta)$ denote a prior density for $\theta \in \Theta$, on which you may impose regularity conditions if needed. Let $p(\theta \,|\, x)$ denote the posterior density.

(a) Show that a nonrandomized decision rule δ_p is Bayes for the prior density $p(\theta)$ if

$$\delta_p(x) = \arg\min_a \int_\Theta L(\theta, a)\, p(x|\theta) p(\theta)\, \mathrm{d}\theta$$

for all possible observations x.

(b) Find the Bayes rule δ_p for a univariate estimation problem under the quadratic loss function $L(\theta, a) = (\theta - a)^2$.

(c) Find the Bayes rule $\delta_{p,\alpha}$ for a univariate estimation problem under the asymmetric piecewise linear loss function

$$L_\alpha(\theta, a) = \begin{cases} (1-\alpha)\,|\theta - a| & \text{if } \theta < a, \\ \alpha\,|\theta - a| & \text{if } \theta \geq a, \end{cases}$$

where $\alpha \in (0, 1)$.

(d) Suppose that the posterior density $p(\theta \,|\, x)$ is continuous, unimodal about its mode μ, and strictly decreasing as θ moves away from μ. Find the Bayes rule $\delta_{p,c}$ for a univariate estimation problem under the zero-one loss function

$$L_c(\theta, a) = \begin{cases} 0 & \text{if } |\theta - a| \leq c, \\ 1 & \text{if } |\theta - a| > c, \end{cases}$$

where $c > 0$.

Problem 2 (estimation in a binomial family). This problem supplements one of the examples in lecture. We consider strategies for estimating the parameter θ from a binomial sample of size n under a quadratic loss function. Nature picks some θ from a prior density $p(\theta)$ on the parameter space $\Theta = (0, 1)$. The statistician observes $X = k$ with sampling density

$$p_\theta(k) = P_\theta(X = k) = \binom{n}{k} \theta^k (1-\theta)^{n-k}, \qquad k = 0, 1, \ldots, n,$$

and takes a decision $\delta(k) \in (0, 1)$, depending on the realization k of the random variable X.

(a) If the prior density is $p = \beta(a, b)$, show that the posterior density is

$$p(\theta \mid k) = \beta(k + a, n - k + b).$$

Show that the posterior density compromises between the prior density and the data, in the sense that the posterior mean of θ, $\delta_p(k) = (k + a)/(n + a + b)$, lies between the prior mean, $a/(a + b)$, and the observed relative frequency, k/n.

(b) Give an example of a prior density and an observation k, in which the posterior variance of θ is higher than the prior variance. Is this a counterintuitive situation?

(c) Show that if $a = b = \sqrt{n}/2$, the Bayes rule

$$\delta_p(k) = \frac{k + \sqrt{n}/2}{n + \sqrt{n}}, \qquad k = 0, 1, \ldots, n,$$

has constant risk, namely

$$R(\theta, \delta_p) = \frac{n}{4(n + \sqrt{n})^2}.$$

Show that this decision rule is admissible and minimax.

Solutions will be discussed in class on November 14.