Technische Universiteit Delft — Mekelweg 4
Faculteit Elektrotechniek, Wiskunde en Informatica — 2628 CD Delft

Exam Statistical Inference (WI4455)
April 11, 2017, 13.30–16.30

Using books or notes is not allowed at the exam.
Unless stated diferently, always add an explanation to your answer.

1. A random variable X has one of two possible densities:
$$f(x \mid \theta) = \theta e^{-\theta x}, \qquad x \in (0, \infty), \qquad \theta \in \{1, 2\}.$$
For $\mu \in [0, \infty]$, consider the family of decision rules
$$d_\mu(x) = \begin{cases} 1 & \text{if } x \geq \mu \\ 2 & \text{if } x < \mu \end{cases}.$$
   (a) Assume loss function $L(\theta, a) = |\theta - a|$. Show that
$$R(\theta, d_\mu) = |\theta - 1| e^{-\theta\mu} + |\theta - 2|(1 - e^{-\theta\mu}).$$
   (b) Sketch the parametrised curve
$$\mathcal{C} = \{(R(1, d_\mu), R(2, d_\mu)) \, : \, \mu \in [0, \infty]\}.$$
   (c) Derive the value of μ for which d_μ is minimax.

2. Suppose $X \sim Ber(\theta)$.
   (a) Find the Fisher information $I(\theta; X)$.
   (b) Show that if we use Jeffreys' prior, then the posterior distribution is the Beta-distribution with parameters $X + 1/2$ and $3/2 - X$.

   Reminder: the density of the Beta distribution with parameters α and β, evaluated at x, is proportional to $x^{\alpha-1}(1-x)^{\beta-1}$.

3. Suppose $X_1, \ldots, X_n$ are independent and identically distributed random variables with the $Ber(\theta)$ distribution.
   (a) Derive a sufficient statistic for θ.
   (b) We consider the following two estimators for θ:
$$\hat{\Theta}_1 = \bar{X}_n \quad \text{(the MLE)}$$
   and
$$\hat{\Theta}_2 = \frac{\sum_{i=1}^n X_i + \alpha}{\alpha + \beta + n} \quad \text{(the posterior mean under the } Be(\alpha, \beta) \text{ prior).}$$
   Calculate the risk of both estimators under squared error loss.

(c) Consider the special case $\alpha = \beta = \sqrt{n/4}$. Show that $\hat{\Theta}_2$ is minimax.

4. Suppose $X_1, \ldots, X_n \sim N(\theta, \sigma^2)$ and assume σ^2 is known. We consider testing $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$.

   (a) Assume a Bayesian setup in which $X_1, \ldots, X_n \mid \Theta = \theta \stackrel{\text{iid}}{\sim} N(\theta, \sigma^2)$ and Θ gets assigned a prior distribution. Take a flat prior on Θ: $f_\Theta(\theta) \propto 1$. Show that the posterior probability of H_0 equals
   $$\Phi\left(\sqrt{n}\frac{\theta_0 - \bar{X}_n}{\sigma}\right),$$
   where Φ denotes the cumulative distribution function of the standard normal distribution.

   (b) Derive an expression for the p-value when using a frequentist test with test statistic $T = \sqrt{n}(\bar{X}_n - \theta_0)/\sigma$.

   (c) Suppose $n = 10$ and that an experimenter has obtained the data $x_1, \ldots, x_{10}$ (which are considered realisations of $X_1, \ldots, X_{10}$. Suppose that the sampling design was such that observations were gathered sequentially until the average of the gathered observations exceeded some threshold. Comment on the validity of the p-value derived under (b).

5. Suppose $X \sim Unif\,(0, \theta)$ and we wish to estimate θ. Overestimation is considered twice as expensive as underestimation and for this reason the following loss function is used:
   $$L(\theta, a) = \begin{cases} \theta - a & \text{if } \theta > a \\ 2(a - \theta) & \text{if } \theta \leq a \end{cases}.$$
   Assume apriori $\Theta \sim Ga\,(2, 1)$, that is $f_\Theta(\theta) \propto \theta e^{-\theta} \mathbf{1}_{[0,\infty)}(\theta)$.

   (a) Derive the posterior density of Θ.

   (b) Derive the Bayes estimator for the given loss function.