Paper 4, Section II

27J Principles of Statistics

Suppose you have at hand a pseudo-random number generator that can simulate an i.i.d. sequence of uniform $U[0,1]$ distributed random variables $U_1^*, \ldots, U_N^*$ for any $N \in \mathbb{N}$. Construct an algorithm to simulate an i.i.d. sequence $X_1^*, \ldots, X_N^*$ of standard normal $N(0,1)$ random variables. [Should your algorithm depend on the inverse of any cumulative probability distribution function, you are required to provide an explicit expression for this inverse function.]

Suppose as a matter of urgency you need to approximately evaluate the integral

$$I = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} \frac{1}{(\pi + |x|)^{1/4}} e^{-x^2/2} dx.$$

Find an approximation I_N of this integral that requires N simulation steps from your pseudo-random number generator, and which has stochastic accuracy

$$\Pr(|I_N - I| > N^{-1/4}) \leqslant N^{-1/2},$$

where Pr denotes the joint law of the simulated random variables. Justify your answer.

Paper 3, Section II

27J Principles of Statistics

State and prove Wilks' theorem about testing the simple hypothesis $H_0 : \theta = \theta_0$, against the alternative $H_1 : \theta \in \Theta \setminus \{\theta_0\}$, in a one-dimensional regular parametric model $\{f(\cdot, \theta) : \theta \in \Theta\}, \Theta \subseteq \mathbb{R}$. [You may use without proof the results from lectures on the consistency and asymptotic distribution of maximum likelihood estimators, as well as on uniform laws of large numbers. Necessary regularity conditions can be assumed without statement.]

Find the maximum likelihood estimator $\hat{\theta}_n$ based on i.i.d. observations $X_1, \ldots, X_n$ in a $N(0, \theta)$-model, $\theta \in \Theta = (0, \infty)$. Deduce the limit distribution as $n \to \infty$ of the sequence of statistics

$$-n \left(\log(\overline{X^2}) - (\overline{X^2} - 1) \right),$$

where $\overline{X^2} = (1/n) \sum_{i=1}^n X_i^2$ and $X_1, \ldots, X_n$ are i.i.d. $N(0,1)$.

Paper 2, Section II

28J Principles of Statistics

In a general decision problem, define the concepts of a *Bayes rule* and of *admissibility.* Show that a unique Bayes rule is admissible.

Consider i.i.d. observations $X_1, \ldots, X_n$ from a Poisson(θ), $\theta \in \Theta = (0, \infty)$, model. Can the maximum likelihood estimator $\hat{\theta}_{MLE}$ of θ be a Bayes rule for estimating θ in quadratic risk for any prior distribution on θ that has a continuous probability density on $(0, \infty)$? Justify your answer.

Now model the X_i as i.i.d. copies of $X|\theta \sim$ Poisson(θ), where θ is drawn from a prior that is a Gamma distribution with parameters $\alpha > 0$ and $\lambda > 0$ (given below). Show that the posterior distribution of $\theta|X_1, \ldots, X_n$ is a Gamma distribution and find its parameters. Find the Bayes rule $\hat{\theta}_{BAYES}$ for estimating θ in quadratic risk for this prior. [The Gamma probability density function with parameters $\alpha > 0, \lambda > 0$ is given by

$$f(\theta) = \frac{\lambda^\alpha \theta^{\alpha-1} e^{-\lambda\theta}}{\Gamma(\alpha)}, \qquad \theta > 0,$$

where $\Gamma(\alpha)$ is the usual Gamma function.]

Finally assume that the X_i have actually been generated from a fixed Poisson(θ_0) distribution, where $\theta_0 > 0$. Show that $\sqrt{n}(\hat{\theta}_{BAYES} - \hat{\theta}_{MLE})$ converges to zero in probability and deduce the asymptotic distribution of $\sqrt{n}(\hat{\theta}_{BAYES} - \theta_0)$ under the joint law $P_{\theta_0}^{\mathbb{N}}$ of the random variables $X_1, X_2, \ldots$. [You may use standard results from lectures without proof provided they are clearly stated.]

Paper 1, Section II

28J Principles of Statistics

State without proof the inequality known as the Cramér–Rao lower bound in a parametric model $\{f(\cdot, \theta) : \theta \in \Theta\}, \Theta \subseteq \mathbb{R}$. Give an example of a maximum likelihood estimator that attains this lower bound, and justify your answer.

Give an example of a parametric model where the maximum likelihood estimator based on observations $X_1, \ldots, X_n$ is biased. State without proof an analogue of the Cramér–Rao inequality for biased estimators.

Define the concept of a *minimax decision rule*, and show that the maximum likelihood estimator $\hat{\theta}_{MLE}$ based on $X_1, \ldots, X_n$ in a $N(\theta, 1)$ model is minimax for estimating $\theta \in \Theta = \mathbb{R}$ in quadratic risk.