De los ejercicios de abajo (sacados del libro de *Georgii, Stochastics*) se proponen los siguientes:

7.2, 7.3, 7.4, 7.5 b) , 7.7, 7.9, 7.10, 7.12, 7.13, 7.14, 7.18, 7.25, 7.26, 7.29.

Aclaración 7.2: Dicho en otras palabras, es lo mismo que tener una muestra $X_1, \cdots, X_n$ i.i.d con $X_1 \sim \mathcal{U}[\theta - \frac{1}{2}, \theta + \frac{1}{2}]$.

(7.39) Example. *Bayes estimate of the expectation of a normal distribution when the variance is known.* Let $(\mathbb{R}^n, \mathscr{B}^n, \mathcal{N}_{\vartheta,v}{}^{\otimes n} : \vartheta \in \mathbb{R})$ be the n-fold Gaussian product model with fixed variance $v > 0$, which has the likelihood function

$$\varrho(x, \vartheta) = (2\pi v)^{-n/2} \exp\Big[-\frac{1}{2v} \sum_{i=1}^{n} (x_i - \vartheta)^2 \Big].$$

We choose a prior distribution which is also normal, namely $\boldsymbol{\alpha} = \mathcal{N}_{m,u}$ for $m \in \mathbb{R}$ and $u > 0$. Using the maximum likelihood estimator $M(x) = \frac{1}{n} \sum_{i=1}^{n} x_i$ and appropriate constants $c_x,\ c'_x > 0$, we can then write

$$\begin{aligned}
\pi_x(\vartheta) &= c_x \ \exp\Big[-\frac{1}{2u}(\vartheta - m)^2 - \frac{1}{2v} \sum_{i=1}^{n} (x_i - \vartheta)^2 \Big] \\
&= c'_x \ \exp\Big[-\frac{\vartheta^2}{2}\Big(\frac{1}{u} + \frac{n}{v}\Big) + \vartheta\Big(\frac{m}{u} + \frac{n}{v} M(x)\Big)\Big] \\
&= \phi_{T(x),u^*}(\vartheta)\,;
\end{aligned}$$

here $u^* = 1/\big(\frac{1}{u} + \frac{n}{v}\big)$ and

$$T(x) = \frac{\frac{1}{u}\, m + \frac{n}{v}\, M(x)}{\frac{1}{u} + \frac{n}{v}}\,.$$

(Since π_x and $\phi_{T(x),u^*}$ are both probability densities, the factor c''_x that appears in the last step is necessarily equal to 1.) Hence, we obtain $\boldsymbol{\pi}_x = \mathcal{N}_{T(x),u^*}$ and, in particular, $T(x) = \mathbb{E}(\boldsymbol{\pi}_x)$. Theorem (7.38) thus tells us that T is the Bayes estimator corresponding to the prior distribution $\boldsymbol{\alpha} = \mathcal{N}_{m,u}$. Note that T is a convex combination of m and M, which gives M more and more weight as either the number n of observations or the prior uncertainty u increases. In the limit as $n \to \infty$, we obtain an analogue of the consistency statement (7.37); see also Problem 7.28.

Problems

7.1. Forest mushrooms are examined in order to determine their radiation burden. For this purpose, n independent samples are taken, and, for each, the number of decays in a time unit is registered by means of a Geiger counter. Set up a suitable statistical model and find an unbiased estimator of the radiation burden.

7.2. *Shifted uniform distributions.* Consider the product model $(\mathbb{R}^n, \mathscr{B}^n, \mathcal{U}_\vartheta{}^{\otimes n} : \vartheta \in \mathbb{R})$, where $\mathcal{U}_\vartheta$ is the uniform distribution on the interval $\big[\vartheta - \frac{1}{2}, \vartheta + \frac{1}{2}\big]$. Show that

$$M = \frac{1}{n} \sum_{i=1}^{n} X_i \quad \text{and} \quad T = \frac{1}{2}\Big(\max_{1\le i\le n} X_i + \min_{1\le i\le n} X_i \Big)$$

are unbiased estimators of ϑ. *Hint:* Use the symmetry of $\mathcal{U}_\vartheta$ for $\vartheta = 0$.

7.3. *Discrete uniform distribution model.* A lottery drum contains N lots labelled with the numbers $1, 2, \dots, N$. Little Bill, who is curious about the total number N of lots, uses an unwatched moment to take a lot at random, read off its number, and return it to the drum. He repeats this n times.

(a) Find a maximum likelihood estimator T of N that is based on the observed numbers $X_1, \dots, X_n$. Is it unbiased? *Hint:* Use Problem 4.5.

(b) Find an approximation to the relative expectation $\mathbb{E}_N(T)/N$ for large N. *Hint:* Treat a suitable expression as a Riemann sum.

7.4. Consider again the setting of Problem 7.3. This time little Bill draws n lots *without* replacing them. Find the maximum likelihood estimator T of N, calculate $\mathbb{E}_N(T)$, and give an unbiased estimator of N.

7.5. Determine a maximum likelihood estimator

(a) in the situation of Problem 7.1,

(b) in the real product model $(\mathbb{R}^n, \mathscr{B}^n, Q_\vartheta^{\otimes n} : \vartheta > 0)$, where $Q_\vartheta = \boldsymbol{\beta}_{\vartheta,1}$ is the probability measure on $(\mathbb{R}, \mathscr{B})$ with density $\rho_\vartheta(x) = \vartheta x^{\vartheta-1}\, 1_{]0,1[}(x)$,

and check whether it is unique.

7.6. *Phylogeny.* When did the most recent common ancestor V of two organisms A and B live? In the 'infinite-sites mutation model', it is assumed that the mutations occur along the lines of descent from V to A and V to B at the times of independent Poisson processes with known intensity (i.e., mutation rate) $\mu > 0$. It is also assumed that each mutation changes a different nucleotide in the gene sequence. Let x be the observed number of nucleotides that differ in the sequences of A and B. What is your maximum likelihood estimate of the age of V? First specify the statistical model!

7.7. A certain butterfly species is split into three types 1, 2 and 3, which occur in the genotypical proportions $p_1(\vartheta) = \vartheta^2$, $p_2(\vartheta) = 2\vartheta(1-\vartheta)$ and $p_3(\vartheta) = (1-\vartheta)^2$, $0 \le \vartheta \le 1$. Among n butterflies of this species you have caught, you find n_i specimens of type i. Determine a maximum likelihood estimator T of ϑ. (Do not forget to consider the extreme cases $n_1 = n$ and $n_3 = n$.)

7.8. At the summer party of the rabbit breeders' association, there is a prize draw for K rabbits. The organisers print $N \ge K$ lots, of which K are winning lots, the remaining ones are blanks. Much to his mum's dismay, little Bill brings x rabbits home, $1 \le x \le K$. How many lots did he probably buy? Give an estimate using the maximum likelihood method.

7.9. Consider the geometric model $(\mathbb{Z}_+, \mathscr{P}(\mathbb{Z}_+), \mathcal{G}_\vartheta : \vartheta \in]0,1])$. Determine a maximum likelihood estimator of the unknown parameter ϑ. Is it unbiased?

7.10. Consider the statistical product model $(\mathbb{R}^n, \mathscr{B}^n, Q_\vartheta^{\otimes n} : \vartheta \in \mathbb{R})$. Suppose that Q_ϑ is the so-called *two-sided exponential distribution* or *Laplace distribution* centred at ϑ, i.e., the probability measure on $(\mathbb{R}, \mathscr{B})$ with density

$$\varrho_\vartheta(x) = \tfrac{1}{2}\, e^{-|x-\vartheta|}, \quad x \in \mathbb{R}.$$

Find a maximum likelihood estimator of ϑ and show that it is unique for even n only. *Hint:* Use Problem 4.15.

7.11. *Estimate of a transition matrix.* Let $X_0, \dots, X_n$ be a Markov chain with finite state space E, known initial distribution α and unknown transition matrix Π. For $a, b \in E$, let $L^{(2)}(a,b) = |\{1 \le i \le n : X_{i-1} = a,\ X_i = b\}|/n$ be the relative frequency of the letter pair (a,b) in the 'random word' $(X_0, \dots, X_n)$. The random matrix $L^{(2)} = \big(L^{(2)}(a,b)\big)_{a,b\in E}$ is called the *empirical pair distribution*. Define the empirical transition matrix T on E by

$$T(a,b) = L^{(2)}(a,b)/L(a) \quad \text{if } L(a) := \sum_{c\in E} L^{(2)}(a,c) > 0\,,$$

and arbitrarily otherwise. Specify the statistical model and show that T is a maximum likelihood estimator of Π. *Hint:* You can argue as in Example (7.7).

7.12. Consider the binomial model of Example (7.14). For any given n, find an estimator of ϑ for which the mean squared error does not depend on ϑ.

7.13. *Unbiased estimators can be bad.* Consider the model $(\mathbb{N}, \mathscr{P}(\mathbb{N}), P_\vartheta : \vartheta > 0)$ of the conditional Poisson distributions

$$P_\vartheta(\{n\}) = \mathcal{P}_\vartheta(\{n\}|\mathbb{N}) = \frac{\vartheta^n}{n!\,(e^\vartheta - 1)}\,, \quad n \in \mathbb{N}.$$

Show that the only unbiased estimator of $\tau(\vartheta) = 1 - e^{-\vartheta}$ is the (useless) estimator $T(n) = 1 + (-1)^n$, $n \in \mathbb{N}$.

7.14. *Uniqueness of best estimators.* In a statistical model $(\mathcal{X}, \mathscr{F}, P_\vartheta : \vartheta \in \Theta)$, let S, T be two best unbiased estimators of a real characteristic $\tau(\vartheta)$. Show that $P_\vartheta(S = T) = 1$ for all ϑ. *Hint:* Consider the estimators $S + c\,(T - S)$ for $c \approx 0$.

7.15. Consider the negative binomial model $(\mathbb{Z}_+, \mathscr{P}(\mathbb{Z}_+), \overline{\mathcal{B}}_{r,\vartheta} : 0 < \vartheta < 1)$ for given $r > 0$. Determine a best estimator of $\tau(\vartheta) = 1/\vartheta$ and determine its variance explicitly for each ϑ.

7.16. *Randomised response.* In a survey on a delicate topic ('Do you take hard drugs?') it is difficult to protect the privacy of the people questioned and at the same time to get reliable answers. That is why the following 'unrelated question method' was suggested. A deck of cards is prepared such that half of the cards contain the delicate question A and the other half a harmless question B, which is unrelated to question A ('Did you go to the cinema last week?'). The interviewer asks the candidate to shuffle the cards, then to choose a card without showing it to anyone, and to answer the question found on this card. The group of people questioned contains a known proportion p_B of people affirming question B (cinema-goers). Let $\vartheta = p_A$ be the unknown probability that the sensitive question A is answered positively. Suppose n people are questioned independently. Specify the statistical model, find a best estimator of ϑ, and determine its variance.

7.17. Consider the n-fold Gaussian product model $(\mathbb{R}^n, \mathscr{B}^n, \mathcal{N}_{m,\vartheta}{}^{\otimes n} : \vartheta > 0)$ with known expectation $m \in \mathbb{R}$ and unknown variance. Show that the statistic

$$T = \sqrt{\frac{\pi}{2}}\,\frac{1}{n}\sum_{i=1}^{n} |X_i - m|$$

on $\mathbb{R}^n$ is an unbiased estimator of $\tau(\vartheta) = \sqrt{\vartheta}$, but that there is no ϑ at which $\mathbb{V}_\vartheta(T)$ reaches the Cramér–Rao bound $\tau'(\vartheta)^2/I(\vartheta)$.

7.18. *Shifted uniform distributions.* Consider the situation of Problem 7.2. Compute the variances $\mathbb{V}_\vartheta(M)$ and $\mathbb{V}_\vartheta(T)$ of the estimators M and V, and decide which of them you would recommend for practical use. *Hint:* For $n \geq 3$ and $\vartheta = 1/2$, determine first the joint distribution density of $\min_{1\leq i\leq n} X_i$ and $\max_{1\leq i\leq n} X_i$, and then the distribution density of T. Use also (2.23).

7.19. *Sufficiency and completeness.* Let $(\mathcal{X}, \mathscr{F}, P_\vartheta : \vartheta \in \Theta)$ be a statistical model and $T : \mathcal{X} \to \Sigma$ a statistic with (for simplicity) countable range Σ. T is called *sufficient* if there exists a family $\{Q_s : s \in \Sigma\}$ of probability measures on $(\mathcal{X}, \mathscr{F})$ that do not depend on ϑ and satisfy $P_\vartheta(\,\cdot\,|T = s) = Q_s$ whenever $P_\vartheta(T = s) > 0$. T is called *complete* if $g \equiv 0$ is the only function $g : \Sigma \to \mathbb{R}$ such that $\mathbb{E}_\vartheta(g \circ T) = 0$ for all $\vartheta \in \Theta$. Let τ be a real characteristic. Show the following.

(a) *Rao–Blackwell.* If T is sufficient, then every unbiased estimator S of τ can be improved as follows: Let $g_S(s) := \mathbb{E}_{Q_s}(S)$ for $s \in \Sigma$; then the estimator $g_S \circ T$ is unbiased and satisfies $\mathbb{V}_\vartheta(g_S \circ T) \leq \mathbb{V}_\vartheta(S)$ for all $\vartheta \in \Theta$.

(b) *Lehmann–Scheffé.* If T is sufficient and complete and S is an arbitrary unbiased estimator of τ, then $g_S \circ T$ is in fact a best estimator of τ. *Hint:* Argue by contradiction.

7.20. Let $(\mathcal{X}, \mathscr{F}, P_\vartheta : \vartheta \in \Theta)$ be an exponential model relative to a statistic T, and suppose for simplicity that T takes values in $\Sigma := \mathbb{Z}_+$. Show that T is sufficient and complete.

7.21. Recall the situation of Problem 7.3 and show that the maximum likelihood estimator T to be determined there is sufficient and complete.

7.22. *Relative entropy and Fisher information.* Let $(\mathcal{X}, \mathscr{F}, P_\vartheta : \vartheta \in \Theta)$ be a regular statistical model with finite sample space $\mathcal{X}$. Show that

$$\lim_{\varepsilon\to 0} \varepsilon^{-2} H(P_{\vartheta+\varepsilon}; P_\vartheta) = I(\vartheta)/2 \quad \text{for all } \vartheta \in \Theta.$$

7.23. *Estimation of the mutation rate in the infinite alleles model.* For given $n \geq 1$, consider Ewens' sampling distribution $\varrho_{n,\vartheta}$ with unknown mutation rate $\vartheta > 0$, as defined in Problem 6.5(a). Show the following.

(a) $\{\varrho_{n,\vartheta} : \vartheta > 0\}$ is an exponential family and $K_n(x) := \sum_{i=1}^n x_i$ (the number of different clans in the sample) is a best unbiased estimator of $\tau_n(\vartheta) := \sum_{i=0}^{n-1} \frac{\vartheta}{\vartheta+i}$.

(b) The maximum likelihood estimator of ϑ is $T_n := \tau_n^{-1} \circ K_n$. (Note that τ_n is strictly increasing.)

(c) The sequence $(K_n/\log n)_{n\geq 1}$ of estimators of ϑ is asymptotically unbiased and consistent. However, the squared error of $K_n/\log n$ is of order $1/\log n$, so it converges to 0 very slowly.

7.24. *Estimation by the method of moments.* Let $(\mathbb{R}, \mathscr{B}, Q_\vartheta : \vartheta \in \Theta)$ be a real-valued statistical model, and let $r \in \mathbb{N}$ be given. Suppose that for each $\vartheta \in \Theta$ and every $k \in \{1, \dots, r\}$, the kth moment $m_k(\vartheta) := \mathbb{E}_\vartheta(\mathrm{Id}_{\mathbb{R}}^k)$ of Q_ϑ exists. Furthermore, let $g : \mathbb{R}^r \to \mathbb{R}$ be continuous, and consider the real characteristic $\tau(\vartheta) := g(m_1(\vartheta), \dots, m_r(\vartheta))$. In the associated infinite product model $(\mathbb{R}^{\mathbb{N}}, \mathscr{B}^{\otimes\mathbb{N}}, Q_\vartheta^{\otimes\mathbb{N}} : \vartheta \in \Theta)$, one can then define the estimator

$$T_n := g\Big(\frac{1}{n}\sum_{i=1}^n X_i,\ \frac{1}{n}\sum_{i=1}^n X_i^2, \dots,\ \frac{1}{n}\sum_{i=1}^n X_i^r\Big)$$

of τ, which is based on the first n observations. Show that the sequence (T_n) is consistent.

7.25. Consider the two-sided exponential model of Problem 7.10. For each $n \geq 1$, let T_n be a maximum likelihood estimator based on n independent observations. Show that the sequence (T_n) is consistent.

7.26. Verify the consistency statement (7.37) for the posterior distributions in the binomial model of Example (7.36).

7.27. *Dirichlet and multinomial distributions.* As a generalisation of Example (7.36), consider an urn model in which each ball has one of a finite number s of colours (instead of only two). Let Θ be the set of all probability densities on $\{1, \dots, s\}$. Suppose that the prior distribution $\boldsymbol{\alpha}$ on Θ is the *Dirichlet distribution* $\mathcal{D}_\varrho$ for some parameter $\varrho \in \,]0, \infty[^s$, which is defined by the equation

$$\mathcal{D}_\varrho(A) = \frac{\Gamma\big(\sum_{i=1}^s \varrho(i)\big)}{\prod_{i=1}^s \Gamma\big(\varrho(i)\big)} \int 1_A(\vartheta) \prod_{i=1}^s \vartheta_i^{\varrho(i)-1} d\vartheta_1 \dots d\vartheta_{s-1}\,, \quad A \in \mathscr{B}_\Theta.$$

(The integral runs over all $(\vartheta_1, \dots, \vartheta_{s-1})$ for which $\vartheta := (\vartheta_1, \dots, \vartheta_{s-1}, 1 - \sum_{i=1}^{s-1} \vartheta_i)$ belongs to Θ. The fact that $\mathcal{D}_\varrho$ is indeed a probability measure will follow for instance from Problem 9.8. In the case $\varrho \equiv 1$, $\mathcal{D}_\varrho$ is the uniform distribution on Θ.) We take a sample of size n with replacement. For each colour composition $\vartheta \in \Theta$, the colour histogram then has the multinomial distribution $\mathcal{M}_{n,\vartheta}$. Determine the associated posterior distribution.

7.28. *Asymptotics of the residual uncertainty as the information grows.* Consider Example (7.39) in the limit as $n \to \infty$. Let $x = (x_1, x_2, \dots)$ be a sequence of observed values in $\mathbb{R}$ such that the sequence of averages $M_n(x) = \frac{1}{n} \sum_{i=1}^n x_i$ remains bounded. Let $\pi_x^{(n)}$ be the posterior density corresponding to the outcomes $(x_1, \dots, x_n)$ and the prior distribution $\mathcal{N}_{m,u}$. Let $\theta_{n,x}$ be a random variable with distribution $\pi_x^{(n)}$. Show that the rescaled random variables $\sqrt{n/v}\,(\theta_{n,x} - M_n(x))$ converge in distribution to $\mathcal{N}_{0,1}$.

7.29. *Gamma and Poisson distribution.* Let $(\mathbb{Z}_+^n, \mathscr{P}(\mathbb{Z}_+^n), \mathcal{P}_\vartheta^{\otimes n} : \vartheta > 0)$ be the n-fold Poisson product model. Suppose the prior distribution is given by $\boldsymbol{\alpha} = \Gamma_{a,r}$, the gamma distribution with parameters $a, r > 0$. Find the posterior density π_x for each $x \in \mathbb{Z}_+^n$, and determine the Bayes estimator of ϑ.