Exercises for the Bayesian Learning course

Mattias Villani

August 2025

Contents

1 Exercise session 1 (preparatory)

Problem 1.1. Let $Y_1, \ldots, Y_n|\theta$ be a sample from an Bernoulli distribution with parameter θ.

(a) Derive the maximum likelihood estimator (MLE) for θ.

(b) Derive a formula for the observed information $\mathcal{J}_n(\hat{\theta})$.

(c) Derive a formula for the Fisher information $\mathcal{I}_n(\theta)$.

(d) Assume that you have observed $n = 20$ trials with a total of $s = 8$ successes. Plot the likelihood function. Compute the MLE. Compute an asymptotic normal approximation of the MLE based on the observed information and plot it.

Problem 1.2. Let $X_1, X_2, \ldots, X_n$ be independent random variables with density function

$$p(x) = \alpha x^{\alpha-1} \quad \text{for } 0 < x < 1,$$

and $p(x) = 0$ otherwise. The parameter α is positive.

(a) Derive the maximum likelihood estimator of α.

(b) Determine a normal large sample approximation of the sampling distribution of the maximum likelihood estimator $\hat{\alpha}$ based on the Fisher information.

(c) Assume that you have observed the data

$$x_1 = 0.89, x_2 = 0.81, x_3 = 0.27, x_4 = 0.96, x_5 = 0.74.$$

Plot the likelihood function. Compute the MLE.

Problem 1.3.

(a) Use the inverse CDF method to simulate 1000 draws from the distribution in Problem 1.2 with $\alpha = 5$. Plot a histogram and overlay a kernel density estimate using the `density` function in R.

(b) Use Monte Carlo integration to compute $\mathbb{E}(\exp(X))$, where X is a random variable following the distribution in Problem 1.2 with $\alpha = 5$. Convince yourself that the Monte Carlo estimate converges by a suitable plot.

(c) Compute the numerical standard error of the Monte Carlo estimate.

(d) Use the central limit theorem to obtain an approximate distribution for the Monte Carlo estimator.

Problem 1.4. Implement rejection sampling to sample from the distribution in Problem 1.2. Use the uniform distribution on $[0, 1]$ as the proposal density $q(x)$.

(a) Determine a suitable bounding constant M, assuming that $\alpha > 1$.

(b) Implement rejection sampling for this distribution for $\alpha = 5$. Sample 10000 draws and plot a histogram of draws. Overlay a plot of the density function.

(c) Is rejection sampling with this proposal distribution efficient when $\alpha = 5$? Relate your answer to the value for M.

(d) Use importance sampling to compute $\mathbb{E}(\exp(X))$ with the uniform distribution on $[0, 1]$ as the proposal/importance density $q(x)$. Plot the 10 largest importance weights.

2 Exercise session 2

3 Exercise session 3

4 Exercise session 4

5 Exercise session 5

6 Exercise session 6

7 Exercise session 7

8 Exercise session 8