TD4 - (Likelihood Ratio) Testing

Exercise 1 We collect one sample X from a Poisson distribution with parameter λ. We recall that its probability mass function is given by

$$\forall k \in \mathbb{N}, \quad f_\lambda(x) = \frac{\lambda^k}{k!} e^{-\lambda} .$$

We want to test $\mathcal{H}_0 : (\lambda = 5)$ against $\mathcal{H}_1 : (\lambda = 10)$ at level $\alpha = 0.05$, based on X.

1. Prove that a randomized Neyman-Pearson test can be formulated as

$$\begin{aligned} \widetilde{D}(X) &= 1, \quad \text{if } X > t \\ \widetilde{D}(X) &= \gamma, \quad \text{if } X = t \\ \widetilde{D}(X) &= 0, \quad \text{if } X < t. \end{aligned}$$

2. Using that $P_{Z\sim\mathcal{P}(5)}(Z > 9) = 0.032$ and $P_{Z\sim\mathcal{P}(5)}(Z > 8) = 0.068$, deduce that $t = 9$ and $\gamma = 1/2$.

3. What is the power of this test?

Exercise 2 We collect iid data $X_1, \ldots, X_n$ from an exponential distribution with parameter θ. We recall that its density is given by

$$\forall x \in \mathbb{R}, \quad f_\theta(x) = \theta \exp(-\theta x) \mathbb{1}_{[0,+\infty[}(x).$$

1. Propose a Uniformly More Powerful test of level α for the test

$$\mathcal{H}_0 : (\theta \le \theta_0) \quad \text{against} \quad \mathcal{H}_1 : (\theta > \theta_0) .$$

2. Can we propose a UMP(α) test for

$$\mathcal{H}_0 : (\theta = \theta_0) \quad \text{against} \quad \mathcal{H}_1 : (\theta \neq \theta_0) \ ?$$

Exercise 3 (exam 2024) We observe a n sample $X_1, \ldots, X_n$ from a geometric distribution $\mathcal{G}(\theta)$, whose probability mass function is given by

$$f_\theta(k) = \theta(1-\theta)^{k-1} \quad \text{for all } k \in \mathbb{N}\backslash\{0\}$$

for some $\theta \in (0,1)$. Given $0 < \theta_0 < \theta_1 < 1$ we are first interested in testing

$$\mathcal{H}_0 : (\theta = \theta_0) \quad \text{against} \quad \mathcal{H}_1 : (\theta = \theta_1)$$

1. Prove that a Likelihood Ratio Test can be expressed as $D(X) = \mathbb{1}(\sum_{i=1}^n X_i < t)$ for some threshold $t \in \mathbb{N}$ with $t > n + 1$.

2. For what values of $\alpha \in (0,1)$ can we find a LRT that is Uniformly More Powerful at level α?

Now we turn our attention to the composite hypothesis testing problem

$$\mathcal{H}_0 : (\theta = \theta_0) \quad \text{against} \quad \mathcal{H}_1 : (\theta \neq \theta_0)$$

where $\theta_0 \in (0,1)$.

3. Compute $\log \widetilde{\Lambda}(X)$, the (generalized) log-likelihood ratio statistic associated to the above test.
4. Propose a test that is asymptotically of level α.

Exercise 4 (exam 2025) We consider a regression model in which we observe iid pairs $(X_i, Y_i)_{1\leq i\leq n}$ with $X_i = (X_{i,1}, X_{i,2}) \in \mathbb{R}^2$ and $Y_i \in \mathbb{R}$. The distribution of (X_i, Y_i) is the following:

$$\begin{aligned} X_i &\sim P_X \\ Y_i &= \beta_0 + \beta_1 X_{i,1} + \beta_2 X_{i,2} + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2) \end{aligned}$$

where P_X is some distribution on $\mathbb{R}^2$. We assume that the variance σ^2 is known and we are interested in the parameter $\beta = (\beta_0, \beta_1, \beta_2) \in \mathbb{R}^3$.

1. We recall the the density of a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ is given by

   $$f_{\mu,\sigma^2}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) .$$

   Give the expression of $\widehat{\beta}_n$, the maximum likelihood estimator of β.

   A closed-form solution is not needed, as long as you can write it as the solution to some optimization problem, made as simple as possible.

2. We are interested in testing the significativity of the variable $X_{i,2}$ by testing

   $$\mathcal{H}_0 : (\beta_2 = 0) \quad \text{against} \quad \mathcal{H}_1 : (\beta_2 \neq 0) .$$

   Propose a test that is asymptotically of level α.

Exercise 5 We consider a so-called one-parameter canonical exponential family, in which the density wrt to some reference measure is

$$f_\theta(x) = h(x) \exp(\theta x - b(\theta))$$

where b is some twice diffenrentiable function that is furhtermore stricly convex ($b'' > 0$). We admit that these assumption are sufficient to be in a regular model. We denote by $\mu(\theta) = \mathbb{E}_\theta[X]$ the expectation of the distribution parameterized by θ.

1. Prove that $b'(\theta) = \mathbb{E}_\theta[X]$ and $b''(\theta) = \mathrm{Var}_\theta[X]$.
2. Deduce that the mapping $\theta \mapsto \mu(\theta)$ is one-to-one. We denote by μ^{-1} its inverse.
3. Compute $\widehat{\theta}_n$, the maximum likelihood estimator of θ.

4. We introduce $\mathrm{K}(\theta,\theta')$, the Kullback-Leibler divergence between P_θ and $P_{\theta'}$, defined as

$$\mathrm{K}(\theta,\theta') = \mathbb{E}_\theta\left[\log\frac{f_\theta(X)}{f_{\theta'}(X)}\right]$$

   Prove that

$$K(\theta,\theta') = (\theta-\theta')\mu(\theta) - b(\theta) + b(\theta')$$

5. Deduce the following inequality: for all $\theta \in \Theta$,

$$\log\frac{L(X_1,\dots,X_n;\widehat{\theta}_n)}{L(X_1,\dots,X_n;\theta)} = n\mathrm{K}(\widehat{\theta}_n,\theta)$$

6. Prove that the (generalized) log-likelihood ratio associated to the test

$$\mathcal{H}_0 : (\theta \le \theta_0) \quad \text{against} \quad \mathcal{H}_1 : (\theta > \theta_0)$$

   satisfies

$$\log\widetilde{\Lambda}(X) = n\mathrm{K}(\widehat{\theta}_n,\theta_0)\mathbb{1}(\widehat{\theta}_n \ge \theta_0).$$

7. We admit the following concentration inequality (called a Chernoff inequality):

$$\forall\theta\in\Theta, \forall x>0, \quad \mathbb{P}_\theta\left(\widehat{\theta}_n > \theta, n\mathrm{K}(\widehat{\theta}_n,\theta) > x\right) \le e^{-x}\ .$$

   Propose a LRT of level α. Is this test UMP(α)?

8. Propose a test of level α for testing

$$\mathcal{H}_0 : (\mu = \mu_0) \quad \text{against} \quad \mathcal{H}_1 : (\mu \neq \mu_0)\ .$$

   Compare it to the asymptotic test of level α obtained using Wilk's theorem, for $\alpha = 0.05$. Which one will have the largest power?

Exercise 6 The scientist Mendel (considered as the father of genetics) did the following experiment. He bred two different kind of peas: one with round yellow seeds and one with wrinkled green seeds. There are four types of progeny: round yellow (1), wrinkled yellow (2), round green (3) and wrinkled green (4). For each individual in the progeny, we denote by p_i the probability that it is of type i.

Assuming that the individual are independent when there are n individuals, the number of individual of each kind $N_n = (N_{n,1}, N_{n,2}, N_{n,3}, N_{n,4})$ follows a so-called multinomial distribution with parameter n and $p = (p_1,p_2,p_3,p_4)$, for which

$$\mathbb{P}(N_n = (n_1,n_2,n_3,n_4)) = \frac{n!}{n_1!n_2!n_3!n_4!}\prod_{i=1}^4 p_i^{n_i}\ .$$

Mendel's ineheritance theory predicts that p is equal to

$$p_0 = \left(\frac{9}{16},\frac{3}{16},\frac{3}{16},\frac{1}{16}\right)$$

He did $n = 556$ experiments and observed $N_n = (315, 101, 108, 32)$. We want to test

$$\mathcal{H}_0 : (p = p_0) \quad \text{against} \quad \mathcal{H}_1 : (p \neq p_0)$$

1. Perform a Likelihood Ratio Test of this hypothesis. Does it reject $\mathcal{H}_0$?
2. For the above testing problem with multinomial data it is common to use another test, called Pearson's χ^2 test. When there are k possible types, this test is based on the test statistic

$$T_n = \sum_{i=1}^{k} \frac{(N_{n,i} - np_{0,i})^2}{np_{0,i}}$$

which is proved to satisfy $T_n \rightsquigarrow \chi^2_{k-1}$. Perform a χ^2 test. Does it reject the hypothesis?
3. Do you think that using statistical testing is appropriate to valide a theory?

Exercise 7 (exam 2025) We collect an iid sample $X_1, \dots, X_n$ distributed under a Binomial distribution with parameters $m \in \mathbb{N}\backslash\{0\}$ and $p \in [0,1]$ for which

$$\mathbb{P}(X_i = k) = \binom{m}{k} p^k (1-p)^{m-k}, \quad k \in \{0, \dots, m\} .$$

We assume that the parameter $m > 0$ is known.

1. Prove that the family of distributions $\{B(m,p), p \in [0,1]\}$ forms an exponential family.
2. Let p_0, p_1 be two values $(0,1)$ such that $p_1 < p_0$. Give the form of the Likelihood Ratio Test of

$$\mathcal{H}_0 : (p = p_0) \text{ against } \mathcal{H}_1 : (p = p_1)$$

3. Explain to which extend this test can be considered optimal.
4. Illustrate graphically the distribution of the test statistic $T(X)$ used in 2. under $\mathcal{H}_0$ and under $\mathcal{H}_1$. Represent graphically the type I error and the power for a given threshold.

   Plausible distributions are enough.