Lab 6

Questions with a (*) next to them form part of homework 2. Due in 2/11/2015. [Q 2,3,4]

1. A p-value is a statistic $P : X \to [0,1]$ such that under H_0,

$$\mathbb{P}(P \leq t) \leq t \text{ for } t \in [0,1].$$

   The p-value should be designed such that small values indicate a divergence between the data and the null hypothesis; the critical region for a test of size α is then $\{x : P(x) < \alpha\}$.

   (a) A coin is tossed 100 times. Find a p-value to test if the coin is biased.

   (b) Let X be a continuous random variable with strictly positive p.d.f. $f : \mathbb{R} \to (0, \infty)$. Show that the c.d.f.

$$F(x) = \mathbb{P}(X \leq x), \qquad F : \mathbb{R} \to (0,1)$$

   is an invertible function and that the distribution of $F(X)$ is Uniform$(0,1)$.

2. (*) Simple tests: When testing $H_0 : \theta = \theta_0$ against $H_1 : \theta = \theta_1$, the size of the test is $\alpha = \mathbb{P}(\text{Reject } H_0 \mid H_0)$. The power of the test is $\beta = \mathbb{P}(\text{Reject } H_0 \mid H_1)$. Suppose $X_1, \ldots X_n \sim N(\mu, 1)$ and you want to test $H_0 : \mu = 0$ against $H_1 : \mu = 1$.

   (a) On the same diagram, sketch the p.d.f. of the $N(\mu, 1/n)$ under both H_0 and H_1. Shade the regions under the curves corresponding to α and β.

   (b) Construct a hypothesis test (using a p-value or the likelihood ratio test) with size α.

   (c) What is the power β as a function of n and α?

   (d) Find β when $n = 10$.

   (e) Find β when $n = 100$.

   (f) How large does n have to be for β to be at least 90%?

3. (*) Simple null, composite hypothesis:

   (a) Let $X_1, \ldots, X_n \sim N(\theta, 1)$. Find a p-value for testing $H_0 : \theta = 0$ against $H_1 : \theta < 0$. What is the critical region in terms of $X_1, \ldots, X_n$? You are working at the 5% significance level, say.

   (b) Let $X_1, \ldots, X_n \sim N(\theta, 1)$. Find a p-value for testing $H_0 : \theta = 0$ against $H_1 : \theta \neq 0$.

   (c) Let $X_1, \ldots, X_n \sim N(0, \sigma^2)$. Find a p-value for testing $H_0 : \sigma^2 = 1$ against $H_1 : \sigma^2 < 1$.

4. (*) Let $X_1, \ldots X_n \sim f(\cdot \mid \theta)$ be a collection of i.i.d.r.v. Let $L(\theta \mid x)$ denote the likelihood,

$$L(\theta|x) = \prod_{i=1}^{n} f(x_i \mid \theta).$$

   Assume f is sufficiently regular...
   Wilks' theorm states that when testing a composite null hypothesis $H_0 : \theta \in \Omega_0$ against a composite alternative hypothesis $H_1 : \theta \in \Omega_1 = \Omega \setminus \Omega_0$, the distribution of

$$-2 \log \frac{L(\hat{\theta}_0 \mid X)}{L(\hat{\theta} \mid X)}$$

   converges in distribution as $n \to \infty$to $\chi^2_{\dim \Omega - \dim \Omega_0}$ where

$$\hat{\theta} = \arg\max_{\theta \in \Omega} L(\theta \mid X), \qquad \hat{\theta}_0 = \arg\max_{\theta \in \Omega_0} L(\theta \mid X).$$

(a) The probability distribution when rolling a die is described by 6 numbers, adding to one. You have two dice, and you roll each die n times. Denote the outcomes $X_1, \ldots, X_n$ for die 1, and $Y_1, \ldots, Y_n$ for die 2.

i	1	2	3	4	5	6
X	130	116	109	197	223	225
Y	100	108	128	209	210	245

Test H_0 : the dice have the same distribution against H_1 : the distributions differ.

(b) Here is some data on US voters gender and political affiliation

Gender\Party	Democrat	Independent	Republican
Male	703	327	468
Female	484	239	477

Is gender and political affiliation linked?

(c) [Not part of the Homework] Look up and use Pearson's χ^2 test statistic to answer part (a) and (b). [Pearson's χ^2 test statistics is an approximation to the likelihood ratio test statistic.]

5. Problems with Hypothesis Testing

(a) The county of Midsomer has a population of 10,001. Or it did; there has been a murder. If we pick someone from the population uniformly at random, there is a one in 10,000 chance they are the murderer.
A genetic test has been devised to test someone's DNA with a DNA sample from the murder weapon. Relatives have similar DNA, so the machine will show a red light, if the test subject is the murderer, or any of the murderer's 99 closest blood relatives. For the other 9,900 inhabitants of Midsomer, the machine will show a green light.
Consider the hypothesis test being applied to a randomly chosen citizen of Midsomer. Consider H_0 : "not a murderer" against H_1 : "a murderer".

i. The machine show a red light. Should you (i) reject H_0, or (ii) fail to reject H_0 / accept H_1, at a 5% significance level? Interpret your test result.
ii. What is the probability the test subject is the killer?
iii. How, in the general context of hypothesis testing, can you reduce the number of false positives a test produces?

(b) Suppose $X_1, \ldots, X_{100} \sim N(\theta, 1)$. Consider a test $H_0 : \theta = 0$ against $H_1 : \theta = 100$. Calculate a critical region for a hypothesis test at a 5% significance level. What is the power of the test? You observe $\bar{X} = 10$. Do you reject H_0 or fail to do so? Compare the values of the likelihood function at $\theta = 0$ and $\theta = 100$. What do you think of the conclusion of the test?

6. Simpson's paradox:

(a) Load the admission statistics for UCB: `data(UCBAdmissions)`. This a "famous" dataset. It count the number of applications to 6 different UC Berkeley departments in 1973.

(b) First calculate some summary statistics:

i. the total number of people accepted/rejected,
ii. the number of men and women in the dataset,
iii. and the number of applications per department.

(c) Produce a 2x2 table showing the numbers of men/women whose applications where accepted/rejected. Do the proportions seem to differ? (Do a χ^2-test.)

(d) Repeat part (c) 6 times for each of the 6 departments. Which of the individual departments display the same pattern of behavior seen in part (c). Can you explain this?

(e) Produce a 2x6 table showing, separated by gender, the fraction of people applying to the six departments.

(f) Produce a 2x6 table showing the fraction of acceptances for each of the six departments.