STA 332/MATH 343: Homework 6 – Solutions

Duke University, Fall 2026

Due: Thursday, October 29 at 11:59 PM on Gradescope

Hypothesis Testing

When asked for an implementable test, specify as simple a rejection statistic as possible and its cutoff.

1. **Definitions**
   (a) Define the power function, size, and level of a test.
   (b) Express the Type I and Type II error probabilities at a given θ in terms of the power function.
   (c) Define a uniformly most powerful (UMP) level-α test.
   (d) Define the monotone likelihood ratio (MLR) property.

   (a) For a test ψ taking values in $\{0, 1\}$, where 1 means rejection, the power function is

   $$\beta_\psi(\theta) = \mathbb{P}_\theta(\psi(\underline{X}) = 1).$$

   The size is $\sup_{\theta \in \Theta_0} \beta_\psi(\theta)$. The test has level α if

   $$\sup_{\theta \in \Theta_0} \beta_\psi(\theta) \leq \alpha.$$

   Thus a level-α test need not have size exactly α.

   (b) If $\theta \in \Theta_0$, the Type I error probability is

   $$\mathbb{P}_\theta(\text{reject } H_0) = \beta_\psi(\theta).$$

   If $\theta \in \Theta_1$, the Type II error probability is

   $$\mathbb{P}_\theta(\text{do not reject } H_0) = 1 - \beta_\psi(\theta).$$

   (c) A test ψ^* is UMP of level α if it has level α and, for every other level-α test ψ,

   $$\beta_{\psi^*}(\theta) \geq \beta_\psi(\theta) \qquad \text{for every } \theta \in \Theta_1.$$

   The same test must maximize power at every parameter value in the alternative.

   (d) A family has monotone likelihood ratio in a statistic T if, for every $\theta_1 > \theta_0$, the ratio

   $$\frac{L_{\underline{x}}(\theta_1)}{L_{\underline{x}}(\theta_0)}$$

   is a nondecreasing function of $T(\underline{x})$, with the usual extended-ratio convention where the denominator is zero (points where both likelihoods vanish may be ignored). If the ratios are instead nonincreasing in T for every $\theta_1 > \theta_0$, the family has nondecreasing MLR in $-T$.

2. **Testing a uniform endpoint.** Let $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} \mathsf{Unif}[0, \theta]$, where $\theta > 0$. We wish to test

$$H_0 : \theta \leq 1 \qquad \text{versus} \qquad H_1 : \theta > 1.$$

Let $M = \max\{X_1, \ldots, X_n\}$ and fix an $\alpha \in (0, 1)$.

(a) Consider the test that rejects when $M > c_\alpha$. Find c_α so that the test has size exactly α.

(b) Compute the power function for all $\theta > 0$.

(c) Now consider the test which rejects when $M > 1$. What is its size? Compute its power function for all $\theta > 0$.

(d) Plot the power functions for the two tests on the same axes and compare. Is one test uniformly more powerful then the other?

(a) By independence, for $0 \leq m \leq \theta$,

$$\mathbb{P}_\theta(M \leq m) = \prod_{i=1}^{n} \mathbb{P}_\theta(X_i \leq m) = \left(\frac{m}{\theta}\right)^n.$$

For a fixed cutoff $c \in (0, 1)$, the rejection probability is zero when $\theta \leq c$ and equals $1 - (c/\theta)^n$ when $\theta > c$. It is therefore nondecreasing in θ, so the supremum over the null is attained at $\theta = 1$:

$$\sup_{0 < \theta \leq 1} \mathbb{P}_\theta(M > c) = \mathbb{P}_1(M > c) = 1 - c^n.$$

Setting this equal to α gives

$$\boxed{c_\alpha = (1 - \alpha)^{1/n}.}$$

(b) The power function of the first test is

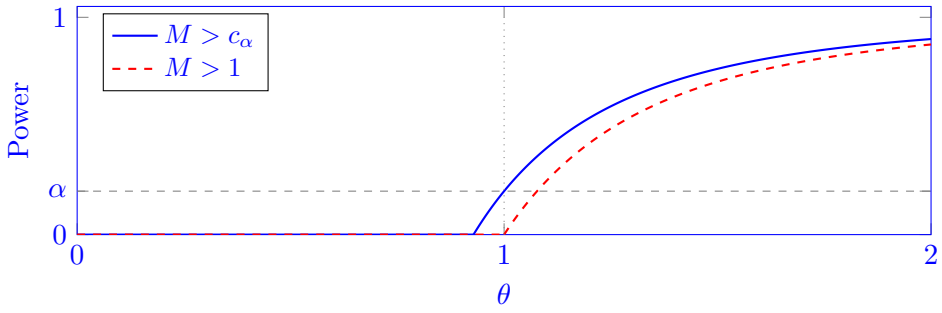
$$\beta_\alpha(\theta) = \begin{cases} 0, & 0 < \theta \leq c_\alpha, \\ 1 - \dfrac{1 - \alpha}{\theta^n}, & \theta > c_\alpha. \end{cases}$$

In particular, $\beta_\alpha(1) = \alpha$.

(c) Under every null value $0 < \theta \leq 1$, we have $M \leq 1$ with probability one. Thus the second test has size zero. Its power function is

$$\beta_0(\theta) = \begin{cases} 0, & 0 < \theta \leq 1, \\ 1 - \dfrac{1}{\theta^n}, & \theta > 1. \end{cases}$$

(d) The plot below illustrates the two power functions for $n = 3$ and $\alpha = 0.2$. The comparison holds for every $n \geq 1$ and $0 < \alpha < 1$.



Since $c_\alpha < 1$, the rejection region $\{M > 1\}$ is contained in $\{M > c_\alpha\}$. Hence $\beta_\alpha(\theta) \geq \beta_0(\theta)$ for every $\theta > 0$. In particular, at every alternative $\theta > 1$,

$$\beta_\alpha(\theta) - \beta_0(\theta) = \frac{\alpha}{\theta^n} > 0.$$

Thus the first test is uniformly more powerful than the second. It also has a larger size (α rather than zero); both tests have level α.

3. **Power of a two-sided normal test.** Let $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\theta, 1)$, where $\theta \in \mathbb{R}$. For

$$H_0 : \theta = 0 \qquad \text{versus} \qquad H_1 : \theta \neq 0,$$

consider the test

$$\psi(\underline{X}) = \mathbb{1}_{\{|\sqrt{n}\,\bar{X}_n| > z_{\alpha/2}\}}.$$

Here, we use the definition $z_\alpha = \Phi^{-1}(1-\alpha)$, where Φ is the standard normal cdf.

(a) Verify that this test has level α. Does it have size exactly α?

(b) Find its power function $\beta(\theta)$ in terms of Φ.

(c) For each fixed $\theta \in \mathbb{R}$, compute $\lim_{n\to\infty} \beta(\theta)$.

(a) Under H_0, $\sqrt{n}\,\bar{X}_n \sim \mathcal{N}(0,1)$ exactly. By symmetry,

$$\mathbb{P}_0(\psi_n = 1) = \mathbb{P}(Z > z_{\alpha/2}) + \mathbb{P}(Z < -z_{\alpha/2}) = \alpha/2 + \alpha/2 = \alpha.$$

Thus the size is exactly α, and the test has level α.

(b) Under θ, $\sqrt{n}\,\bar{X}_n$ has the same distribution as $Z + \sqrt{n}\,\theta$, where $Z \sim \mathcal{N}(0,1)$. With $z = z_{\alpha/2}$,

$$\beta_n(\theta) = 1 - \Phi(z - \sqrt{n}\,\theta) + \Phi(-z - \sqrt{n}\,\theta).$$

(c) If $\theta = 0$, the power equals α for every n. If $\theta > 0$, the first term $1 - \Phi(z - \sqrt{n}\,\theta)$ tends to one and the last term tends to zero. If $\theta < 0$, the first term tends to zero and the last term tends to one. Consequently,

$$\lim_{n\to\infty} \beta_n(\theta) = \begin{cases} \alpha, & \theta = 0, \\ 1, & \theta \neq 0. \end{cases}$$

Under the null the rejection probability remains controlled at α; at each fixed nonzero alternative, the Type II error probability tends to zero. No normal approximation is needed.

4. **Testing a normal variance.** Let $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2)$, where μ is known and $\sigma^2 > 0$ is unknown. Let $T = \sum_{i=1}^n (X_i - \mu)^2$.

   (a) By standardizing, show that
   $$\frac{T}{\sigma^2} = \sum_{i=1}^n Z_i^2$$
   for suitable i.i.d. standard Gaussian random variables $Z_1, \ldots, Z_n$.

   A sum of squares of n i.i.d. standard Gaussian random variables is called a *chi-squared random variable with n degrees of freedom*, denoted χ^2_n. Write F_n for its cdf and $\chi^2_{n,\alpha} = F_n^{-1}(\alpha)$ for its *lower* α-quantile (in contrast to our convention $z_\alpha = \Phi^{-1}(1-\alpha)$ for Gaussian quantiles).

   (b) For $\sigma_1^2 > \sigma_0^2 > 0$, use Neyman–Pearson to find the implementable MP level-α test for $H_0 : \sigma^2 = \sigma_0^2$ versus $H_1 : \sigma^2 = \sigma_1^2$.

   (c) Show that the family has MLR in the sufficient statistic T. Use Karlin–Rubin to find the UMP level-α test for $H_0 : \sigma^2 \le \sigma_0^2$ versus $H_1 : \sigma^2 > \sigma_0^2$.

   (d) Find the power function for the test in part (c), for all $\sigma^2 > 0$. Verify from the power function that the size of the test is α.

   (a) Let $\sigma = \sqrt{\sigma^2} > 0$ and define $Z_i = (X_i - \mu)/\sigma$. Standardizing a normal random variable gives $Z_i \sim \mathcal{N}(0,1)$, and the Z_i remain independent because each is a function of a separate independent X_i. Therefore,
   $$\frac{T}{\sigma^2} = \sum_{i=1}^n \frac{(X_i - \mu)^2}{\sigma^2} = \sum_{i=1}^n Z_i^2 \sim \chi^2_n.$$
   Consequently, $\mathbb{P}_{\sigma^2}(T/\sigma^2 \le t) = F_n(t)$.

   (b) The likelihood and ratio are
   $$L_{\underline{x}}(\sigma^2) = (2\pi\sigma^2)^{-n/2} e^{-T(\underline{x})/(2\sigma^2)},$$
   $$\frac{L_{\underline{x}}(\sigma_1^2)}{L_{\underline{x}}(\sigma_0^2)} = \left(\frac{\sigma_0^2}{\sigma_1^2}\right)^{n/2} \exp\left\{\frac{1}{2}\left(\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right) T(\underline{x})\right\}.$$
   The coefficient of T is positive, so Neyman–Pearson gives an upper-tail rule. Under σ_0^2, $T/\sigma_0^2 \sim \chi^2_n$, hence the test rejects when
   $$T > \sigma_0^2 \chi^2_{n,1-\alpha}.$$
   Continuity gives rejection probability exactly α at σ_0^2.

   (c) Factorization shows that T is sufficient. The ratio in (b) is increasing in T for every $\sigma_1^2 > \sigma_0^2$, proving MLR. Karlin–Rubin gives the same rejection rule for the composite hypotheses.

   (d) For every $\sigma^2 > 0$,
   $$\beta(\sigma^2) = 1 - F_n\left(\frac{\sigma_0^2}{\sigma^2}\chi^2_{n,1-\alpha}\right).$$
   This is increasing in σ^2, since the argument of F_n decreases as σ^2 increases. Thus
   $$\sup_{0<\sigma^2\le\sigma_0^2} \beta(\sigma^2) = \beta(\sigma_0^2) = \alpha.$$

5. **A logistic location model.** Suppose we are given a single sample X_1 with density

$$f_\theta(x) = \frac{e^{x-\theta}}{(1+e^{x-\theta})^2}, \qquad x \in \mathbb{R}, \quad \theta \in \mathbb{R}.$$

You may use $\mathbb{P}_\theta(X > c) = 1/(1+e^{c-\theta})$.

(a) Prove that $h(x) = (1+e^x)/(1+e^{x-1})$ is increasing in x.

(b) Find the implementable MP level-α test for $H_0 : \theta = 0$ versus $H_1 : \theta = 1$.

(c) Find the power function of this test and verify that $\beta(0) = \alpha$.

(a) Differentiating gives

$$h'(x) = \frac{e^x(1-e^{-1})}{(1+e^{x-1})^2} > 0.$$

(b) For the single observation $\underline{x} = (x)$, the likelihood ratio is

$$\frac{L_{\underline{x}}(1)}{L_{\underline{x}}(0)} = \frac{f_1(x)}{f_0(x)} = e^{-1}\left(\frac{1+e^x}{1+e^{x-1}}\right)^2 = e^{-1}h(x)^2.$$

Since h is positive and increasing, the ratio is increasing. The MP test rejects when $X > c_\alpha$, where

$$\frac{1}{1+e^{c_\alpha}} = \alpha, \qquad c_\alpha = \log\left(\frac{1-\alpha}{\alpha}\right).$$

(c) Substitution into the supplied survival function gives

$$\beta(\theta) = \frac{1}{1+((1-\alpha)/\alpha)e^{-\theta}} = \frac{\alpha e^\theta}{1-\alpha+\alpha e^\theta}.$$

The size is $\beta(0) = \alpha$.

6. **A lower-tail Karlin–Rubin test.** Let $X_1, \dots, X_n$ be i.i.d. with density

$$f_\theta(x) = \frac{2\sqrt{\theta}}{\sqrt{\pi}} e^{-\theta x^2}, \qquad x > 0, \quad \theta > 0,$$

and let $T = \sum_{i=1}^n X_i^2$.

(a) Show that T is sufficient. For $\theta_1 > \theta_0$, compute the likelihood ratio and determine whether it is increasing or decreasing in T. Identify a statistic in which the family has nondecreasing MLR.

(b) By comparing densities, show that $\sqrt{2\theta}\, X_i$ has the same distribution as $|Z|$, where $Z \sim \mathcal{N}(0,1)$. Deduce the distribution of $2\theta T$.

(c) Use Karlin–Rubin to find the implementable UMP level-α test for

$$H_0 : \theta \le \theta_0 \qquad \text{versus} \qquad H_1 : \theta > \theta_0.$$

Formulate the test using the χ^2 quantiles introduced in Problem 4.

(d) Find the power function and verify directly that the test has size α. Explain why small values of T favor the alternative.

(a) On $(0,\infty)^n$, the likelihood is

$$L_{\underline{x}}(\theta) = \left(\frac{2}{\sqrt{\pi}}\right)^n \theta^{n/2} e^{-\theta T(\underline{x})}.$$

Factorization gives sufficiency of T. For $\theta_1 > \theta_0$,

$$\frac{L_{\underline{x}}(\theta_1)}{L_{\underline{x}}(\theta_0)} = \left(\frac{\theta_1}{\theta_0}\right)^{n/2} e^{-(\theta_1-\theta_0)T(\underline{x})}.$$

This ratio is decreasing in T, so the family has nondecreasing MLR in the sufficient statistic $-T$.

(b) With $Y = \sqrt{2\theta}\, X_i$, the change-of-variables formula gives

$$f_Y(y) = \frac{2\sqrt{\theta}}{\sqrt{\pi}} e^{-y^2/2} \frac{1}{\sqrt{2\theta}} = \sqrt{\frac{2}{\pi}} e^{-y^2/2}, \qquad y > 0.$$

This is the density of $|Z|$. Thus $2\theta X_i^2$ has the same distribution as Z^2. Independence yields $2\theta T \sim \chi^2_n$.

(c) Karlin–Rubin applied to $-T$ gives a lower-tail rule in T. To obtain size α at θ_0, reject when

$$T < \frac{\chi^2_{n,\alpha}}{2\theta_0}.$$

(d) The power is

$$\beta(\theta) = F_n\left(\frac{\theta}{\theta_0}\chi^2_{n,\alpha}\right).$$

It is increasing in θ, so $\sup_{0<\theta\le\theta_0} \beta(\theta) = \beta(\theta_0) = \alpha$. The representation $X_i \overset{d}{=} |Z|/\sqrt{2\theta}$ shows that larger θ produces smaller observations and hence smaller T.