STK4011 and STK9011
Autumn 2016

Hypothesis testing

Covers (most of) the following material from chapter 8:

- Section 8.1
- Sections 8.2.1 and 8.2.3
- Section 8.3.1
- Section 8.3.2 (until definition 8.3.16)

Ørnulf Borgan
Department of Mathematics
University of Oslo

Basic concepts

Assume that we have random variables

$\mathbf{X} = (X_1, X_2, \ldots, X_n)$ with joint pmf or pdf

$f(\mathbf{x} | \theta) = f(x_1, \ldots, x_n | \theta)$ where $\theta \in \Theta$

We want to test the null hypothesis $H_0 : \theta \in \Theta_0$
versus the alternative hypothesis $H_1 : \theta \in \Theta_0^c$

A hypothesis test is a procedure that specifies:

- for which values of $\mathbf{X}$ we reject H_0 (accept H_1)
- for which values of $\mathbf{X}$ we do not reject H_0 (accept H_0)

Usually a test is specified in terms of a test statistic $W = W(\mathbf{X})$

Likelihood ratio tests

Let $X_1, X_2, \ldots, X_n$ be a random sample from the population $f(x | \theta)$, so $X_1, X_2, \ldots, X_n$ are iid and their pmf or pdf is $f(x | \theta)$, where θ may be a vector

Then the likelihood is given by

$$L(\theta | \mathbf{x}) = L(\theta | x_1, \ldots, x_k) = \prod_{i=1}^{n} f(x_i | \theta)$$

The likelihood ratio test statistic for testing $H_0 : \theta \in \Theta_0$ versus $H_1 : \theta \in \Theta_0^c$ is

$$\lambda(\mathbf{x}) = \frac{\sup_{\Theta_0} L(\theta | \mathbf{x})}{\sup_{\Theta} L(\theta | \mathbf{x})}$$

The likelihood ratio test (LRT) has rejection region of the form $\{\mathbf{x} : \lambda(\mathbf{x}) \le c\}$

Let $\hat{\theta}$ be the unrestricted maximum likelihood estimator of θ, i.e. the value of θ that maximizes the likelihood when $\theta \in \Theta$

Let $\hat{\theta}_0$ be the maximum likelihood estimator of θ under the null hypothesis, i.e. the value of θ that maximizes the likelihood when $\theta \in \Theta_0$

Then the LRT statistics takes the form

$$\lambda(\mathbf{x}) = \frac{L(\hat{\theta}_0 | \mathbf{x})}{L(\hat{\theta} | \mathbf{x})}$$

Example 8.2.2 (normal LRT)

Let $X_1, X_2, \ldots, X_n$ be iid $n(\theta, 1)$

We will test $H_0 : \theta = \theta_0$ versus $H_1 : \theta \ne \theta_0$

Example 8.2.3 (exponential LRT)

Let $X_1, X_2, ..., X_n$ be iid with pdf ($-\infty < \theta < \infty$)

$$f(x|\theta) = \begin{cases} e^{-(x-\theta)} & x \ge \theta \\ 0 & x < \theta \end{cases}$$

We will test $H_0: \theta \le \theta_0$ versus $H_1: \theta > \theta_0$

Theorem 8.2.4

If $T(\mathbf{X})$ is a sufficient statistic for θ and $\lambda^*(t)$ and $\lambda(\mathbf{x})$ are the LRT statistics based on T and $\mathbf{X}$, respectively, then $\lambda^*(T(\mathbf{x})) = \lambda(\mathbf{x})$ for all $\mathbf{x}$ in the sample space

Example 8.2.5 (LRT and sufficiency)

Let $X_1, X_2, ..., X_n$ be iid $n(\theta, 1)$

We will test $H_0: \theta = \theta_0$ versus $H_1: \theta \ne \theta_0$

Example 8.2.6 (normal LRT, unknown variance)

Let $X_1, X_2, ..., X_n$ be iid $n(\mu, \sigma^2)$

We will test $H_0: \mu \le \mu_0$ versus $H_1: \mu > \mu_0$

Union-intersection tests

Assume the null hypothesis may be expressed as

$$H_0: \theta \in \bigcap_{\gamma \in \Gamma} \Theta_\gamma$$

where Γ is an index set (finite or infinite)

Suppose there are tests available for testing

$$H_{0\gamma}: \theta \in \Theta_\gamma \quad \text{versus} \quad H_{1\gamma}: \theta \in \Theta_\gamma^c$$

The rejection region for the test of $H_{0\gamma}: \theta \in \Theta_\gamma$ is

$$\{\mathbf{x}: T_\gamma(\mathbf{x}) \in R_\gamma\}$$

Then the union-intersection test has rejection region

$$\bigcup_{\gamma \in \Gamma} \{\mathbf{x}: T_\gamma(\mathbf{x}) \in R_\gamma\}$$

In particular, if the test for $H_{0\gamma}: \theta \in \Theta_\gamma$ has rejection region

$$\{\mathbf{x}: T_\gamma(\mathbf{x}) > c\}$$

then the union-intersection test has rejection region

$$\bigcup_{\gamma \in \Gamma} \{\mathbf{x}: T_\gamma(\mathbf{x}) > c\} = \{\mathbf{x}: \sup_{\gamma \in \Gamma} T_\gamma(\mathbf{x}) > c\}$$

Example 8.2.8 (normal union-intersection test)

Let $X_1, X_2, ..., X_n$ be iid $n(\mu, \sigma^2)$

We will test $H_0: \mu = \mu_0$ versus $H_1: \mu \ne \mu_0$

Intersection-union tests

Assume the null hypothesis may be expressed as

$$H_0 : \theta \in \bigcup_{\gamma\in\Gamma} \Theta_\gamma$$

where Γ is an index set (finite or infinite)

Suppose that $\{\mathbf{x}: T_\gamma(\mathbf{x}) \in R_\gamma\}$ is the rejection region for a test of $H_{0\gamma} : \theta \in \Theta_\gamma$ versus $H_{1\gamma} : \theta \in \Theta_\gamma^c$

Then the intersection-union test has rejection region

$$\bigcap_{\gamma\in\Gamma} \{\mathbf{x}: T_\gamma(\mathbf{x}) \in R_\gamma\}$$

If the test of $H_{0\gamma} : \theta \in \Theta_\gamma$ has rejection region $\{\mathbf{x}: T_\gamma(\mathbf{x}) \geq c\}$ the rejection region of the intersection-union test becomes $\{\mathbf{x}: \inf_{\gamma\in\Gamma} T_\gamma(\mathbf{x}) \geq c\}$

Error probabilities and power

We will test $H_0 : \theta \in \Theta_0$ versus $H_1 : \theta \in \Theta_0^c$

We may make two types of error:

		Decision: Accept H_0	Decision: Reject H_0
Truth	H_0	Correct decision	Type I Error
	H_1	Type II Error	Correct decision

Let R be the rejection region of the test, so we reject $H_0 : \theta \in \Theta_0$ if $\mathbf{X} \in R$

Probability of Type I error: $P_\theta(\mathbf{X} \in R), \ \theta \in \Theta_0$

Probability of Type II error:

$$P_\theta(\mathbf{X} \in R^c) = 1 - P_\theta(\mathbf{X} \in R), \ \theta \in \Theta_0^c$$

Power function: $\beta(\theta) = P_\theta(\mathbf{X} \in R)$

Example 8.3.3 (normal power function)

Let $X_1, X_2, \ldots, X_n$ be iid $n(\theta, \sigma^2)$ with σ^2 known

We will test $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$

Casella & Berger distinguish between the size and level of a test:

- a test with power function $\beta(\theta)$ is a size α test if $\sup_{\theta\in\Theta_0} \beta(\theta) = \alpha$
- a test with power function $\beta(\theta)$ is a level α test if $\sup_{\theta\in\Theta_0} \beta(\theta) \leq \alpha$

Example 8.3.7 (size of normal LRT, modified)

Let $X_1, X_2, \ldots, X_n$ be iid $n(\theta, \sigma^2)$ with σ^2 known

We will test $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$

A test with power function $\beta(\theta)$ is unbiased if $\beta(\theta') \geq \beta(\theta'')$ for all $\theta' \in \Theta_0^c$ and $\theta'' \in \Theta_0$

Most powerful tests

Definition 8.3.11

Let C be a class of tests for testing $H_0: \theta \in \Theta_0$ versus $H_1: \theta \in \Theta_0^c$. A test in the class C, with power function $\beta(\theta)$, is a uniformly most powerful (UMP) class C test if $\beta(\theta) \geq \beta'(\theta)$ for every $\theta \in \Theta_0^c$ and every $\beta'(\theta)$ that is a power function of a test in class C

We will use this definition when C is the class of all level α tests

Theorem 8.3.12 (Neyman-Pearson Lemma)

Consider testing $H_0: \theta = \theta_0$ versus $H_1: \theta = \theta_1$, where the pdf or pmf corresponding to θ_i is $f(\mathbf{x} | \theta_i);\ i = 0,1$, using a test with rejection region R that satisfies

$$\mathbf{x} \in R \text{ if } f(\mathbf{x}|\theta_1) > k\, f(\mathbf{x}|\theta_0)$$
$$\mathbf{x} \in R^c \text{ if } f(\mathbf{x}|\theta_1) < k\, f(\mathbf{x}|\theta_0) \qquad (8.3.1)$$

for some $k \geq 0$, and

$$\alpha = P_{\theta_0}(\mathbf{X} \in R) \qquad (8.3.2)$$

Then

a) Any test that satisfies (8.3.1) and (8.3.2) is a UMP level α test

b) If there exists a test satisfying (8.3.1) and (8.3.2) with $k > 0$, then every UMP level α test is a size α test [i.e. satisfies (8.3.2)] and every UMP level α test satisfies (8.3.1) except perhaps on a set A satisfying $P_{\theta_0}(\mathbf{X} \in A) = P_{\theta_1}(\mathbf{X} \in A) = 0$

Corollary 8.3.13

Consider the hypothesis testing problem of Theorem 8.3.12. Suppose that $T = T(\mathbf{X})$ is a sufficient statistic for θ and let $g(t|\theta_i)$ be the pdf or pmf of T corresponding to $\theta_i;\ i = 0,1$. Then any test based on T with rejection region S is a UMP level α test if it satisfies

$$t \in S \text{ if } g(t|\theta_1) > k\, g(t|\theta_0)$$
$$t \in S^c \text{ if } g(t|\theta_1) < k\, g(t|\theta_0)$$

for some $k \geq 0$, where

$$\alpha = P_{\theta_0}(T \in S)$$

Example 8.3.15 (UMP normal test)

Let $X_1, X_2, ..., X_n$ be iid $n(\theta, \sigma^2)$ with σ^2 known

We will find the UMP test for testing test

$H_0 : \theta = \theta_0$ versus $H_1 : \theta = \theta_1$ where $\theta_0 > \theta_1$