STAT 513 fa 2020 Lec 06 slides

Building tests of hypotheses

Karl B. Gregory

University of South Carolina

These slides are an instructional aid; their sole purpose is to display, during the lecture, definitions, plots, results, etc. which take too much time to write by hand on the blackboard. They are not intended to explain or expound on any material.

Let $X_1, \dots, X_n$ be iid with pdf or pmf $f(x;\theta)$.

- The *likelihood function* for θ is given by

$$L(\theta; X_1, \dots, X_n) = \prod_{i=1}^{n} f(X_i; \theta).$$

- The *log-likelihood function* is

$$\ell(\theta; X_1, \dots, X_n) = \sum_{i=1}^{n} \log f(X_i; \theta).$$

- The *maximum likelihood estimator (MLE)* of θ is given by

$$\hat{\theta} = \text{argmax}_{\theta \in \Theta} L(\theta; X_1, \dots, X_n) = \text{argmax}_{\theta \in \Theta} \ell(\theta; X_1, \dots, X_n)$$

So $\hat{\theta}$ is like the value of θ that makes the prob. of the observed data the highest.

Review: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, with σ^2 known.

1. Write down the likelihood function for μ.
2. Give the log-likelihood.
3. Find the MLE for μ.

Likelihood ratio

For $X_1, \dots, X_n$ having likelihood $L(\theta; X_1, \dots, X_n)$ and for some hypotheses

$$H_0\colon \theta \in \Theta_0 \text{ versus } H_1\colon \theta \in \Theta_1,$$

the *likelihood ratio (LR)* is defined as

$$\text{LR}(X_1, \dots, X_n) = \frac{\sup_{\theta \in \Theta_0} L(\theta; X_1, \dots, X_n)}{\sup_{\theta \in \Theta} L(\theta; X_1, \dots, X_n)}.$$

The likelihood ratio must take values in the interval _______.

A ______ (larger/smaller) likelihood ratio casts ______ (more/less) doubt on H_0.

Likelihood ratio test

A *likelihood ratio test (LRT)* is a test of the form

$$\text{Reject } H_0 \text{ iff } \text{LR}(X_1, \ldots, X_n) < c$$

for some $c \in [0, 1]$.

A ______ value of c gives the test ______ power and ______ size.

The critical value c can be chosen to give the test a desired size.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, σ^2 known. For

$$H_0\colon \mu = \mu_0 \text{ versus } H_1\colon \mu \neq \mu_0$$

1. Give the LR.
2. Calibrate the LRT to have size α for any $\alpha \in (0, 1)$.

Help me compute the LR! 😭

We can write the likelihood ratio as

$$\mathrm{LR}(X_1,\dots,X_n) = \frac{\sup_{\theta\in\Theta_0} L(\theta;X_1,\dots,X_n)}{\sup_{\theta\in\Theta} L(\theta;X_1,\dots,X_n)} = \frac{L(\hat\theta_0;X_1,\dots,X_n)}{L(\hat\theta;X_1,\dots,X_n)},$$

where

$$\hat\theta_0 = \mathrm{argmax}_{\theta\in\Theta_0}\,\ell(\theta;X_1,\dots,X_n)$$
$$\hat\theta = \mathrm{argmax}_{\theta\in\Theta}\,\ell(\theta;X_1,\dots,X_n)$$

So we just need to find these and plug them in:

- $\hat\theta_0$ is a *restricted maximum likelihood estimator*; best estimator in null space.
- $\hat\theta$ is the MLE.

Note: If the null is a simple hypothesis, i.e. H_0: $\theta = \theta_0$, then $\hat\theta_0 = \theta_0$.

Also: If $\hat\theta \in \Theta_0$, then $\hat\theta = \hat\theta_0$, so that $\mathrm{LR}(X_1,\dots,X_n) = 1$.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, σ^2 known. For

$$H_0\colon \mu \leq \mu_0 \text{ versus } H_1\colon \mu > \mu_0$$

1. Find the restricted MLE $\hat{\mu}_0$.
2. Give the LR.
3. Calibrate the LRT to have size α for any $\alpha \in (0, 1)$.

Exercise: Let $Y \sim \text{Geometric}(p)$ distribution, $p \in [0, 1]$ unknown. Suppose we wish to test

$$H_0\colon p \leq p_0 \text{ versus } H_1\colon p > p_0$$

for some $p_0 \in [0, 1]$.

1. Give the likelihood.
2. Give the log-likelihood.
3. Find the MLE $\hat{p}$ of p.
4. Find the restricted MLE $\hat{p}_0$.
5. Find an expression for the likelihood ratio.
6. For any $\alpha \in (0, 1)$, calibrate the rejection region of the likelihood ratio test so that it has size at most α.

Why do we like likelihood ratio tests?

Neyman-Pearson Lemma

Let $X_1, \dots, X_n$ have the likelihood $L(\theta; X_1, \dots, X_n)$, and suppose we wish to test

$$H_0\colon \theta = \theta_0 \text{ versus } H_1\colon \theta = \theta_1.$$

Then the test

$$\text{Reject } H_0 \text{ iff } \frac{L(\theta_0; X_1, \dots, X_n)}{L(\theta_1; X_1, \dots, X_n)} < c, \quad \text{some } c \in [0, 1]$$

is the most powerful test among tests of same hyps with the same or smaller size.

We want powerful tests with controlled size: N-P Lemma points us toward LRTs.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Poisson}(\lambda)$, λ unknown, and consider

$$H_0\colon \lambda = \lambda_0 \text{ versus } H_1\colon \lambda = \lambda_1,$$

where $\lambda_0 < \lambda_1$.

1. For any $\alpha \in (0, 1)$ find a test with size at most α.
2. For $\lambda_0 = 2$ and $\lambda_1 = 3$ and a sample of size 5, give the most powerful test with size at most 0.05.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Exponential}(\beta)$, β unknown, and consider

$$H_0\colon \beta = \beta_0 \text{ versus } H_1\colon \beta = \beta_1,$$

where $\beta_1 < \beta_0$.

1. For any $\alpha \in (0, 1)$, find the most powerful test of size α.
2. For $\beta_0 = 3$ and $\beta_1 = 2$ and $n = 5$, give the most powerful test with size 0.01.

Uniformly most powerful test

A test of

$$H_0\colon \theta \in \Theta_0 \text{ versus } H_1\colon \theta \in \Theta_1$$

with power function $\gamma(\theta)$ is called *uniformly most powerful* if

$$\gamma(\theta) > \gamma'(\theta) \text{ for all } \theta \in \Theta_1,$$

where γ' is the power function of any other same-sized test of the same hyps.

We can sometimes show that the LRT is the uniformly most powerful test.

The Neyman-Pearson Lemma provides the basis for proving those results.

For two-sided hypotheses there is no uniformly most powerful test!

Example: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$ based on which it is of interest to test

$$H_0\colon \mu = \mu_0 \text{ versus } H_1\colon \mu \neq \mu_0.$$

Letting

$$T_n = \sqrt{n}(\bar{X}_n - \mu_0)/S_n,$$

each of the following tests is a size-α test of the two-sided set of hypotheses:

$$\text{Reject } H_0 \text{ iff } T_n > t_{n-1,\alpha}$$
$$\text{Reject } H_0 \text{ iff } T_n < -t_{n-1,\alpha}$$
$$\text{Reject } H_0 \text{ iff } |T_n| > t_{n-1,\alpha/2}$$

The power curves of these tests are shown on the next slide.

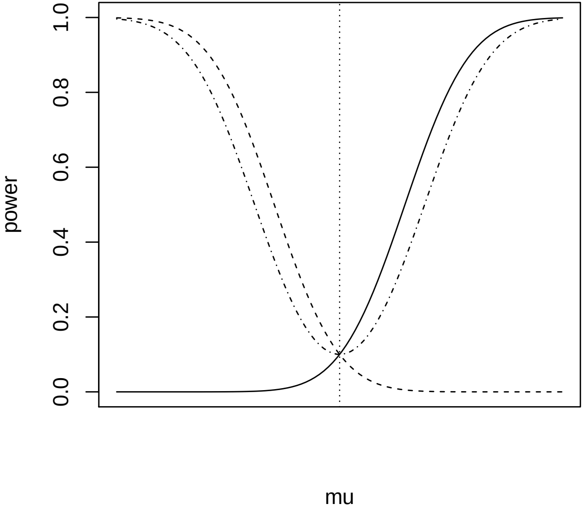
1.0
0.8
0.6
0.4
0.2
0.0
power
mu

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, μ and σ^2 unknown, and consider

$$H_0\colon \mu \leq \mu_0 \text{ versus } H_1\colon \mu > \mu_0.$$

Show that the test

$$\text{Reject } H_0 \text{ iff } \sqrt{n}(\bar{X}_n - \mu_0)/S_n > t_{n-1,\alpha}$$

is the LRT with size α. Steps:

1. Find the MLEs $\hat{\mu}$ and $\hat{\sigma}^2$.
2. Find the restricted MLEs $\hat{\mu}_0$ and $\hat{\sigma}_0^2$ under $H_0\colon \mu \leq \mu_0$.
3. Show that $\text{LR} < c$ is equivalent to $\sqrt{n}(\bar{X}_n - \mu_0)/S_n > c_1$.

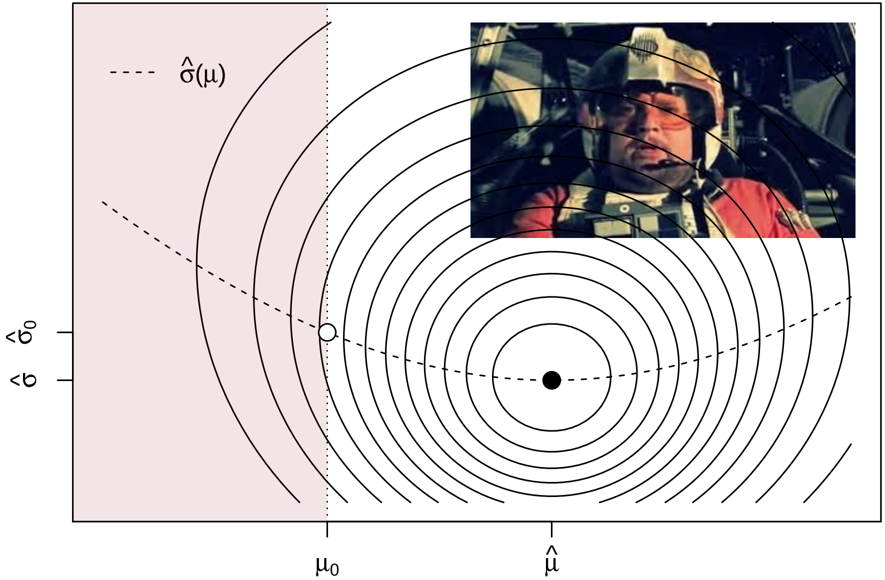
$\hat{\sigma}(\mu)$
$\hat{\sigma}_0$
$\hat{\sigma}$
μ_0
$\hat{\mu}$

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, μ and σ^2 unknown, and consider

$$H_0\colon \sigma^2 = \sigma_0^2 \text{ versus } H_1\colon \sigma^2 \neq \sigma_0^2.$$

Show that the likelihood ratio test has the form

$$\text{Reject } H_0 \text{ iff } \frac{(n-1)S_n^2}{\sigma_0^2} < c_1 \text{ or } \frac{(n-1)S_n^2}{\sigma_0^2} > c_2.$$

Steps:

1. Find the MLEs $\hat{\mu}$ and $\hat{\sigma}^2$.
2. Find the restricted MLEs $\hat{\mu}_0$ and $\hat{\sigma}_0^2$ under $H_0\colon \sigma^2 = \sigma_0^2$.
3. Show that $\text{LR} < c$ is equivalent to above.

Wilk's Theorem

Let $X_1, \dots, X_n$ be a rs with likelihood $L(\theta; X_1, \dots, X_n)$, $\theta \in \Theta$, and consider

$$H_0\colon \theta \in \Theta_0 \text{ versus } H_1\colon \theta \in \Theta_1,$$

where $\dim(\Theta) = d$ and $\dim(\Theta_0) = d_0 < d$. Then under H_0, we have

$$-2 \log \frac{\sup_{\theta \in \Theta_0} L(\theta; X_1, \dots, X_n)}{\sup_{\theta \in \Theta} L(\theta; X_1, \dots, X_n)} \rightarrow \chi^2_{d-d_0} \text{ in distribution}$$

as $n \rightarrow \infty$, provided some conditions (beyond scope of this course) hold.

Wilk's Theorem allows us to define the following test:

Definition (Asymptotic likelihood ratio test)

The size-α *asymptotic likelihood ratio test* is

$$\text{Reject } H_0 \text{ iff } -2\log\frac{\sup_{\theta\in\Theta_0} L(\theta; X_1,\dots,X_n)}{\sup_{\theta\in\Theta} L(\theta; X_1,\dots,X_n)} > \chi^2_{d-d_0,\alpha},$$

and this test has size approximately equal to α for large n.

- $d_0 = \dim(\Theta_0)$ is the number of parameters left unspecified by H_0.
- $d = \dim(\Theta)$ is the total number of unknown parameters.

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, μ and σ^2 unknown, and consider

$$H_0\colon \mu = \mu_0 \text{ versus } H_1\colon \mu \neq \mu_0.$$

1. Find the size-α asymptotic LRT.
2. Run a simulation to compare the power curve of the asymptotic LRT to that of the (non-asymptotic) LRT

$$\text{Reject } H_0 \text{ iff } |\sqrt{n}(\bar{X}_n - \mu_0)/S_n| > t_{n-1,\alpha/2}.$$

   Use the settings $n = 20$ and $\alpha = 0.05$.

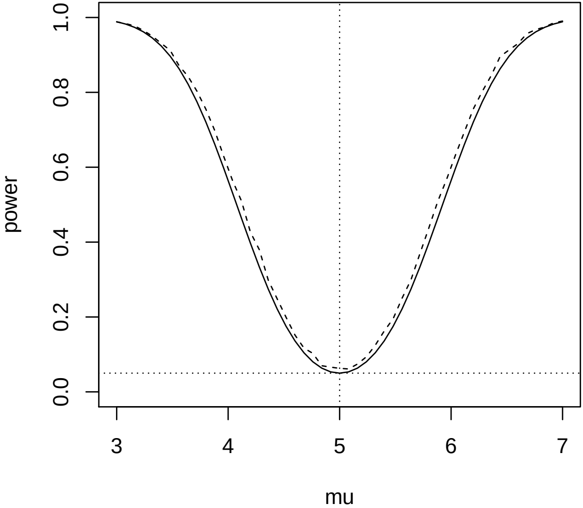
power
0.0
0.2
0.4
0.6
0.8
1.0
3
4
5
6
7
mu

Exercise: Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Poisson}(\lambda)$, λ unknown, and consider testing

$$H_0\colon \lambda = \lambda_0 \text{ versus } H_1\colon \lambda \neq \lambda_0.$$

1. Find an expression for the likelihood ratio.
2. For any $\alpha \in (0,1)$, find a test which has size approaching α as $n \to \infty$.
3. Let $\lambda_0 = 3$ and run a simulation to get power curves for the test under the sample sizes $n = 5, 10, 20, 40$ using $\alpha = 0.05$.
4. Problems with the test? For what n does it have the desired size?
5. Find the p-value of the asymptotic likelihood ratio test of

   $$H_0\colon \lambda = 3 \text{ versus } H_1\colon \lambda \neq 3$$

   associated with a sample of size $n = 25$ with sample mean equal to 3.5.

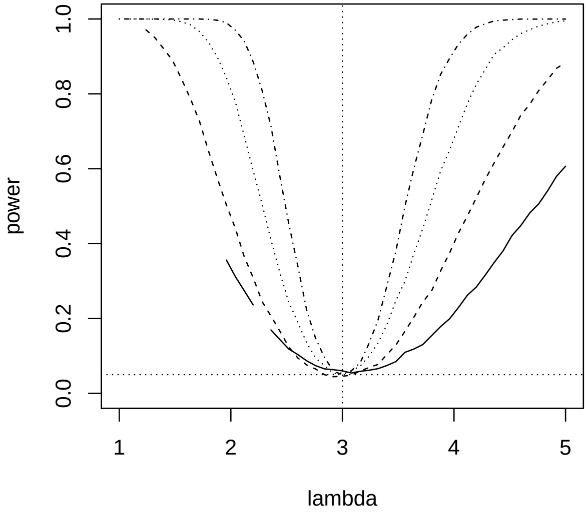
power
0.0
0.2
0.4
0.6
0.8
1.0
1
2
3
4
5
lambda

Exercise: Let $X_{k1}, \dots, X_{kn_k} \overset{\text{ind}}{\sim} \text{Exponential}(\lambda_k)$, $k = 1, 2$ and consider

$$H_0\colon \lambda_1 = \lambda_2 \text{ versus } H_1\colon \lambda_1 \neq \lambda_2.$$

- Find an expression for the likelihood ratio.
- For any $\alpha \in (0, 1)$, find a test which has size approaching α as $n \to \infty$.
- Plot power curves of the test with $\alpha = 0.05$ under

$$(n_1, n_2) = (10, 20), (20, 40), (50, 100), (100, 200)$$

  when $\lambda_1 = 3$ with λ_2 varying from 1 to 5.
- Find the p-value of the asymptotic LRT associated with observing $\bar{X}_1 = 4.1$ and $\bar{X}_2 = 3.2$ when $n_1 = 30$ and $n_2 = 34$.

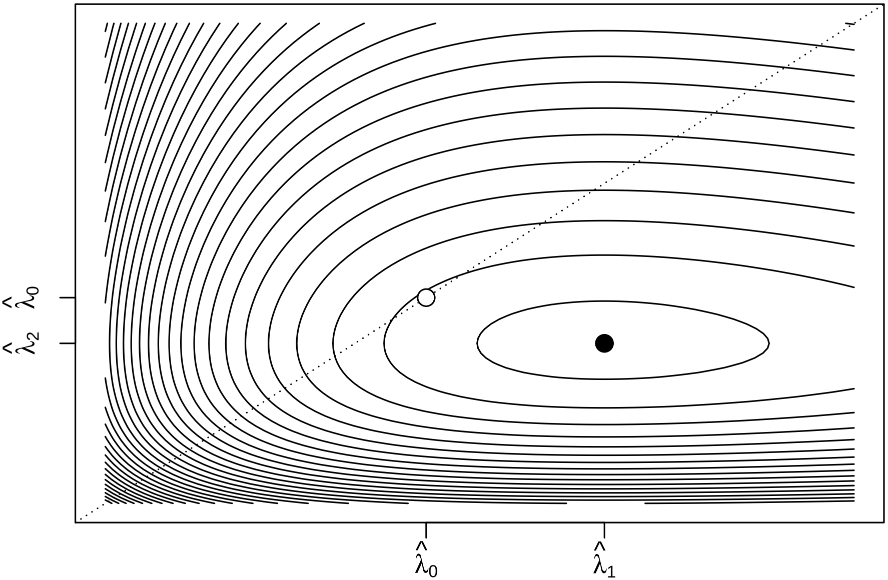
$\hat{\lambda}_0$
$\hat{\lambda}_2$
$\hat{\lambda}_0$
$\hat{\lambda}_1$

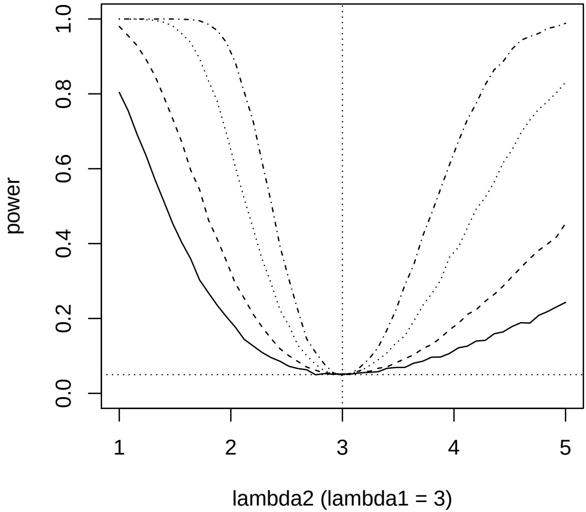
power
0.0
0.2
0.4
0.6
0.8
1.0
1
2
3
4
5
lambda2 (lambda1 = 3)