STA732
Statistical Inference

Lecture 17: Uniformly Most Powerful Tests

Yuansi Chen

Spring 2023

Duke University

https://www2.stat.duke.edu/courses/Spring23/sta732.01/

Recap from Lecture 16

- Reviewed hypothesis testing
- Introduced the Neyman-Pearson Paradigm
- Proved Neyman-Pearson Lemma in simple vs simple testing: likelihood ratio test exists and is optimal (most powerful at level α)

Goal of Lecture 17

1. Uniformly most powerful tests
2. One-param exponential family and UMP one-sided tests
3. Monotone likelihood ratios and UMP one-sided tests
4. p-values
5. Duality between testing and interval estimation

Chap. 12.3-12.7 of Keener or Chap. 3 of Lehmann and Romano

Uniformly most powerful tests (UMP)

Def. Uniformly most powerful tests

A test ϕ^* with level α is called uniformly most powerful (UMP) if

$$\mathbb{E}_\theta \phi^* \geq \mathbb{E}_\theta \phi, \quad \forall \theta \in \Omega_1,$$

for all ϕ with level at most α.

In other words,

$$\beta_{\phi^*}(\theta) \geq \beta_\phi(\theta), \quad \forall \theta \in \Omega_1$$
$$\text{s.t. } \beta_\phi(\theta) \leq \alpha, \quad \forall \theta \in \Omega_0$$

UMP typically only exist for 1 sided testing in certain 1 parameter families

Discussion on the formulation

- Compare the above formulation to a loss-based formulation
- Compare to Bayes factor

Identifiable model

Def. Identifiable family

A family $\mathscr{P} = \{P_\theta, \theta \in \Omega\}$ is identifiable if

$$\theta_1 \neq \theta_2 \Rightarrow P_{\theta_1} \neq P_{\theta_2}.$$

One-param exponential family and UMP one-sided tests

One-param exponential family in simple vs simple

$X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} p_\eta(x) = e^{\eta T(x) - A(\eta)} h(x)$

$H_0 : \eta = \eta_0$ vs $H_1 : \eta = \eta_1 > \eta_0$.

One-param exponential family in simple vs one-sided composite

What can we say about LRT in the previous slide? for the testing

$$H_0 : \eta = \eta_0 \text{ vs } H_1 : \eta > \eta_0.$$

H_1 is an example of one-sided alternative, since the parameter values only lie on one side of the parameter η_0.

The LRT ϕ^* is also UMP for the simple vs one-sided composite test!

Monotone likelihood ratios and UMP one-sided Tests

Monotone likelihood ratio (MLR)

Def. Families with monotone likelihood ratio

We say that the family of densities $p_\theta : \theta \in \mathbb{R}$ has monotone likelihood ratio in $T(x)$ if

1. $\theta \neq \theta'$ implies $p_\theta = p_{\theta'}$ (identifiability)
2. $\theta < \theta'$ implies $\frac{p_{\theta'}(x)}{p_\theta(x)}$ is a nondecreasing function of $T(x)$ (monotonicity)

Example of MLR family: double exponential

$$p_\theta(x) = \frac{1}{2} \exp(-|x - \theta|)$$

Not an example of MLR family: Cauchy location

$$p_\theta(x) = \frac{1}{\pi} \frac{1}{1 + (x - \theta)^2}$$

Show that it is not MLR with $T(x) = x$.

MLR and UMP

Thm. 12.9 in Keener

Suppose the family of densities has monotone likelihood ratios. The testing problem is $H_0 : \theta \le \theta_0$ vs $H_1 : \theta > \theta_0$. Set the level $\alpha \in (0, 1)$. Let

$$\phi^*(x) = \begin{cases} 0 & T(x) < c \\ \gamma & T(x) = c \\ 1 & T(x) > c \end{cases}$$

with c, γ chosen such that $\mathbb{E}_{\theta_0} \phi^*(x) = \alpha$. Then

1. ϕ^* is a UMP level-α test
2. $\beta(\theta) = \mathbb{E}_\theta \phi^*(X)$ is non-decreasing in θ, strictly increasing whenever $\beta(\theta) \in (0, 1)$
3. If $\theta_1 < \theta_0$ then ϕ^* minimizes $\mathbb{E}_{\theta_1} \phi(X)$ among all tests with $\mathbb{E}_{\theta_0} \phi(X) = \alpha$

Picture and proof: Intuitively, ϕ^* is a LRT for $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$ for any pair $\theta_0 < \theta_1$ (with sig. level depends on θ_0). The arbitrary choice of θ_0, θ_1 allows us to extend the simple vs. simple to composite vs. composite

Example

Suppose our data $X_1, \ldots, X_n$ are i.i.d. from unifrom distribution on $(0, \theta)$.

$$p_\theta(x) = \begin{cases} \frac{1}{\theta^n}, & M(x) > 0 \text{ and } T(x) < \theta \\ 0 & \text{otherwise} \end{cases}$$

where $M(x) = \min\{x_1, \ldots, x_n\} > 0$ and $T(x) = \max\{x_1, \ldots, x_n\} < \theta$. Show that

- It has monotone likelihood ratios
- For $H_0 : \theta \leq 1,$ vs $H_1 : \theta > 1$, calculate the power function of the UMP test

p-values

Motivation

- Testing at a fixed level α as described is one of two standard (non-Bayesian) approaches to the evaluation of hypotheses
- The other way is based on p-values
    - The LRT determines k in $\frac{p_1(x)}{p_0(x)} > k$ as a function of α
    - Instead of determine whether the hypothesis is accepted or rejected at the given significance level, we can determine the smallest significance level at which the hypothesis would be rejected for the given observation

“Informal definition” of p-value

Suppose $\phi(x)$ rejects for large values of $T(x)$. Then

$$\hat{p}(x) = \mathbb{P}_{H_0}\left(T(X) \geq T(x)\right).$$

The chance that you observe something more extreme than what you have observed

Example

$X \sim \mathcal{N}(\theta, 1), H_0 : \theta = 0$ vs $H_1 : \theta > 0$. Let

$$\phi_\alpha(x) = \mathbf{1}_{x > z_\alpha},$$

the one sided test rejects for large value of $T(X) = X$.
Then the p-value is

$$\hat{p}(x) = \mathbb{P}_0(X > x) = 1 - \Phi(x).$$

Formal definition

Assume we have a test ϕ_α for each significance level, $\sup_{\theta \in \Omega_0} \mathbb{E}_\theta \phi_\alpha \leq \alpha$.
Assume tests are monotone in α:
if $\alpha_1 \leq \alpha_2$, then $\phi_{\alpha_1}(x) \leq \phi_{\alpha_2}(x)$
Then

$$\hat{p}(x) = \inf \{\alpha : \phi_\alpha(x) = 1\}$$

In the non-randomized test case,

$$\hat{p}(x) = \inf \{\alpha : x \in R_\alpha\}$$

Duality between testing and interval estimation

Confidence region

Def. Confidence region

We say a random set $C(X)$ is a $1-\alpha$ confidence region for a parameter $g(\theta)$ if

$$\mathbb{P}_\theta(g(\theta) \in C(X)) \geq 1-\alpha, \forall \theta \in \Omega.$$

Get a confidence region from a test

Suppose we have a level-α test $\phi(x;a)$ for the testing

$$H_0 : g(\theta) = a \text{ vs } H_1 : g(\theta) \neq a, \forall a \in g(\Omega)$$

Let $C(x) = \{a : \phi(x;a) < 1\}$, which is the all non-rejected values of θ.

Then we have

$$\mathbb{P}_\theta(g(\theta) \notin C(X)) = \mathbb{P}_\theta(\phi(X; g(\theta)) = 1) \leq \alpha, \forall \theta$$

We get a $1 - \alpha$ confidence region!

Get a test from a confidence region

Let $\phi(X) = \mathbf{1}_{a \notin C(X)}$. Then for any θ such that $g(\theta) = a$, we have

$$\mathbb{E}_\theta \phi(X) = \mathbb{P}_\theta(g(\theta) \notin C(X)) \leq \alpha.$$

We get a α-level test!

Example

Suppose $X \sim \text{Exp}(\theta) = \frac{1}{\theta} e^{-\frac{x}{\theta}}, x > 0, \theta > 0$
with CDF $\mathbb{P}_\theta(X \leq x) = 1 - e^{-\frac{x}{\theta}}$.
Invert the LRT for $H_0 : \theta \leq \theta_0$, to get a confidence region (or lower confidence bound)

Summary

- Introduced UMP: uniformly most powerful tests
- UMP is possible, for one sided tests, for exponential family and family with monotone likelihood ratios
- p-value as the second standard (non-Bayesian) approach to evaluate hypothesis testing
- There is duality between testing and interval estimation

What is next?

- Two-Sided Hypotheses
- Least Favorable Distributions

Thank you