Uniformly Most Powerful Tests

November 16, 2023

Outline

1. Definition and examples

2. The monotone likelihood ratio case

Definition

Definition (UMP test)

A level α test ϕ^* is uniformly most powerful (UMP) for $H : \theta \in \Theta_0$ versus $K : \theta \in \Theta_1$ if

$$\beta(\theta, \phi^*) \geq \beta(\theta, \phi), \ \ \theta \in \Theta_1.$$

for any other test ϕ at the level α.

- We have seen that in the case of simple hypotheses the most powerful (which is always UMP, why?) always exists (Neyman-Pearson).
- Are there any other examples?

Definition

Definition (UMP test)

A level α test ϕ^* is uniformly most powerful (UMP) for $H : \theta \in \Theta_0$ versus $K : \theta \in \Theta_1$ if

$$\beta(\theta, \phi^*) \geq \beta(\theta, \phi), \;\; \theta \in \Theta_1.$$

for any other test ϕ at the level α.

- We have seen that in the case of simple hypotheses the most powerful (which is always UMP, why?) always exists (Neyman-Pearson).
- Are there any other examples?
- We will see that there are classes of problems for which such tests exist.
- We will also note that this property is rather limited to one-dimensional parameter families.

1D- normal distribution example

- Consider $X = (X_1, ..., X_n)$ a sample of n random variables distributed as $N(\mu, 1)$.

1D- normal distribution example

- Consider $X = (X_1, ..., X_n)$ a sample of n random variables distributed as $N(\mu, 1)$.
- We want test $H : \mu = 0$ versus $K : \mu = \nu$, where ν is known.

1D- normal distribution example

- Consider $X = (X_1, ..., X_n)$ a sample of n random variables distributed as $N(\mu, 1)$.
- We want test $H : \mu = 0$ versus $K : \mu = \nu$, where ν is known.
- The likelihood ratio is

$$\frac{L(0)}{L(\nu)} = \exp\left(\nu \sum_{i=1}^{n} X_i - n\frac{\nu^2}{2\sigma^2}\right)$$

- Note that any strictly increasing function of an optimal statistic is optimal because the two statistics generate the same of critical regions (as we have argued in the previous lecture).
- Therefore, $\sqrt{n}\bar{X}$ is optimal by the Neyman-Pearson lemma.
- Clearly, the critical value for this test must be the quantile $z_{1-\alpha}$ of the standard normal distribution.
- We note that the test does not depend on the specific value of ν.
- Thus by a simple logical argument it is also the most powerful test for the problem $H : \mu = 0$ versus $K : \mu > 0$

1D- normal distribution example

- Consider $X = (X_1, ..., X_n)$ a sample of n random variables distributed as $N(\mu, 1)$.
- We want test $H : \mu = 0$ versus $K : \mu = \nu$, where ν is known.
- The likelihood ratio is

$$\frac{L(0)}{L(\nu)} = \exp\left(\nu \sum_{i=1}^{n} X_i - n\frac{\nu^2}{2\sigma^2}\right)$$

- Note that any strictly increasing function of an optimal statistic is optimal because the two statistics generate the same of critical regions (as we have argued in the previous lecture).
- Therefore, $\sqrt{n}\bar{X}$ is optimal by the Neyman-Pearson lemma.
- Clearly, the critical value for this test must be the quantile $z_{1-\alpha}$ of the standard normal distribution.
- We note that the test does not depend on the specific value of ν.
- Thus by a simple logical argument it is also the most powerful test for the problem $H : \mu = 0$ versus $K : \mu > 0$
- Can you deliver details of such logical argument?

Using the power to chose the sample size

- We evaluated the power of the test from the previous slide

$$\Phi(z_\alpha + \nu\sqrt{n})$$

- By the Neyman–Pearson lemma this is the largest power available with a level α test.
- If we want the probability of detecting a signal ν to be at least a preassigned value β (say, .90 or .95), then we solve

$$\Phi(z_\alpha + \nu\sqrt{n}) = \beta$$

for n and find that we need to take

$$n = (\nu)^{-2}(z_{1-\alpha} + z_\beta)^2/\nu^2$$

- This is the smallest possible n for any size α test.

Multivariate normal – Fisher's discrimination function

Consider a more general problem of testing for the means and covariances of a multivariate normal distributions.

$$\boldsymbol{\theta}_0 = (\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) \text{ vs. } \boldsymbol{\theta}_1 = (\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1),$$

The likelihood ratio is

$$\frac{L(\boldsymbol{\theta}_0)}{L(\boldsymbol{\theta}_1)} = \frac{det^{1/2}(\boldsymbol{\Sigma}_0) \exp\left(-(\mathbf{x} - \boldsymbol{\mu}_1)^\top \boldsymbol{\Sigma}_1^{-1} (\mathbf{x} - \boldsymbol{\mu}_1)\right)}{det^{1/2}(\boldsymbol{\Sigma}_1) \exp\left(-(\mathbf{x} - \boldsymbol{\mu}_0)^\top \boldsymbol{\Sigma}_0^{-1} (\mathbf{x} - \boldsymbol{\mu}_0)\right)}$$

Thus the N-P test is equivalent to rejecting when one gets large value of

$$(\mathbf{x} - \boldsymbol{\mu}_0)^\top \boldsymbol{\Sigma}_0^{-1} (\mathbf{x} - \boldsymbol{\mu}_0) - (\mathbf{x} - \boldsymbol{\mu}_1)^\top \boldsymbol{\Sigma}_1^{-1} (\mathbf{x} - \boldsymbol{\mu}_1).$$

The special important case of $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_0$ reduces to the *Fisher discrimination function* being large

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_0)^\top \boldsymbol{\Sigma}_0^{-1} \mathbf{x}$$

The lack of the multivariate UMP test

- We observe in the previous example the N-P test statistics depends intrinsically on θ_1 even in the case when the covariances are equal and known.
- If this is the case than there is no UMP test. Why?
- However in the one-dimensional case, which is the special case of this one, we have seen that the test is the UMP.
- We see some contradiction.
- The reason is that the one-dimensional case is special since the unknown parameter $\nu = \mu_1$ entered as a scaling which cancels in the definition of the region.
- To see it, consider more generally $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_0 + \lambda\boldsymbol{\Delta}_0$ with all $\boldsymbol{\mu}_0, \boldsymbol{\Delta}_0, \boldsymbol{\Sigma}_0$ known then the N-P test reduces
$$\boldsymbol{\Delta}_0^\top \boldsymbol{\Sigma}_0^{-1}(\mathbf{x} - \boldsymbol{\mu}_0) \geq c$$
with $c = z_{1-\alpha}(\boldsymbol{\Delta}_0^\top \boldsymbol{\Sigma}_0^{-1}\boldsymbol{\Delta}_0)^{1/2}$. See Problems 4.2.8, 4.2.9 in the assignment.
- But in the general situation when $\boldsymbol{\mu}_1$ is truly multidimensional the N-P test is depending on the values of $\boldsymbol{\theta}_1$ and thus there is no UMP test.

Outline

MLR families of the likelihood

- Typically the MP test of $H : \theta = \theta_0$ versus $K : \theta = \theta_1$ depends on θ_1 and the test is not UMP.
- However, we have seen models where, in the case of a real parameter, there is a statistic T such that the test with critical region $\{x : T(x) \geq c\}$ is UMP.
- This is part of a general phenomena we now describe.

Definition (Definition 4.3.2)

The family of models is said to be a monotone likelihood ratio (MLR) family in T if for $\theta_1 < \theta_2$ the distributions P_{θ_1} and P_{θ_2} are distinct and there exists a statistic $T(x)$ such that the likelihood ratio

$$\frac{L(\theta_2)}{L(\theta_1)} = \frac{p(x, \theta_2)}{p(x, \theta_1)}$$

is an increasing function of $T(x)$.

The main results

Theorem (Theorem 4.3.1)

For a MLR family in T, $\beta(\theta)$ is increasing in θ and if $\beta(\theta_0) = \alpha$, then the corresponding test it UMP level α for $H : \theta \leq \theta_0$ vs. $K : \theta > \theta_0$.

Corollary (Corollary 4.3.1)

For a MLR family in T, if the distribution function F_0 of $T(X)$ under $X \sim P_{\theta_0}$ is continuous and $t_{1-\alpha}$ is the $1-\alpha$-quantile of this distribution, then the test that rejects H if and only if $T(x) \geq t_{1-\alpha}$ is UMP level α for testing $H : \theta \leq \theta_0$ vs. $K : \theta > \theta_0$.

- The proofs of both are simple logical exercises left for your own enjoyment.
- Examples 4.3.1, 4.3.2 are recommended to be followed .
- For the one-parameter exponential family model

$$p(x, \theta) = h(x) \exp(\eta(\theta) T(x) - B(\theta)).$$

If $\eta(\theta)$ is strictly increasing, then this family is MLR.