STATS 300A: Theory of Statistics **Fall 2015**

Lecture 14 — November 10

Lecturer: Lester Mackey *Scribe: Jun Yan, Matteo Sesia*

Warning: *These notes may contain factual and/or typographic errors.*

14.1 Overview

14.1.1 Hypothesis Testing Optimality Goal

Recall that the hypothesis testing problem can be formulated as $H_0 : \ \theta \in \Omega_0$ vs. $H_1 : \theta \in \Omega_1$. Here our goal is to find a uniformly most powerful (UMP) level-α test ϕ which

$$\begin{aligned} \text{maximizes the power function} \quad & \mathbb{E}_{\theta_1}\phi(x) \\ \text{subject to} \quad & \mathbb{E}_{\theta_0}\phi(x) \leq \alpha, \end{aligned}$$

for every $\theta_0 \in \Omega_0$ and $\theta_1 \in \Omega_1$. In other words, ϕ maximizes the power over the alternative space while keeping the size of ϕ less than the required level α over the entire null set.

14.1.2 Strategies for Finding UMPs

Although the existence of a UMP test is not generally guaranteed, there are some general purpose strategies to find a UMP test when one exists. One well-studied strategy contains the following three steps:

1. **Reduce the composite alternative to a simple alternative**: If H_1 is composite, fix $\theta_1 \in \Omega_1$, and test the null hypothesis against the simple alternative $\theta = \theta_1$. (Hope that doesn't depend on θ_1.)

2. **Collapse the composite null to a simple null**: If H_0 is composite, collapse the null hypothesis to a simple one by averaging over the null space Ω_0. We will discuss this strategy in today's lecture.

3. **Apply Neyman Pearson lemma**: Find the MP LRT for testing the resulting simple null versus the resulting simple alternative using the NP lemma. Note that if the resulting test does not depend on θ_1, then it will be UMP for the H_0 vs H_1.

14.2 Optimal Tests for Composite Nulls

In previous lectures, our focus was on hypothesis testing problems with a simple null. Here we introduce a new strategy to deal with cases with a composite null.

14.2.1 The Model

Consider the case with a simple alternative:

$$\begin{aligned} H_0: &\quad X \sim f_\theta, \quad \theta \in \Omega_0 \\ H_1: &\quad X \sim g, \end{aligned}$$

where g is known. We now impose a prior distribution Λ on Ω_0. So we consider the new hypothesis

$$H_\Lambda: \ X \sim h_\Lambda(x) = \int_{\Omega_0} f_\theta(x)\, d\Lambda(\theta),$$

where $h_\Lambda(x)$ is the marginal distribution of X induced by Λ. In order to reduce the problem to a simple versus simple case, let us test H_Λ against H_1. Notice that the MP test given by the NP lemma should be checked to work for the original composite null. This task can be achieved by picking Λ to be *the least favorable distribution* which will be defined later.

In the more general case of a composite null vs. a composite alternative, once an MP test for the composite null vs. simple alternative is found, we can check whether it works for every θ_1 in the alternative parameter space. If so, the resulting test is UMP for the composite vs. composite case.

14.2.2 Least Favourable Distribution

Let β_Λ be the power of the MP level-α test ϕ_Λ for testing H_Λ vs. g.

Definition 1 (Least favorable Distribution)**.** Λ is a least favorable distribution if $\beta_\Lambda \le \beta_{\Lambda'}$ for any prior Λ'.

Hence, Λ will be the least favorable distribution if the MP test under Λ has smaller power than the MP test under any other prior distribution. The following theorem can help us to deal with the case of composite null by using the notion of least favorable distribution, which tells that if we choose Λ in the right way, we can get the MP.

Theorem 1 (TSH 3.8.1)**.** Suppose ϕ_Λ is a MP level-α test for testing H_Λ against g. If ϕ_Λ is level-α for the original hypothesis H_0 (i.e., $\mathbb{E}_{\theta_0}\phi_\Lambda(x) \le \alpha, \ \forall\ \theta_0 \in \Omega_0$), then

1. The test ϕ_Λ is MP for original $H_0: \ \theta \in \Omega_0$ vs. g.
2. The distribution Λ is least favorable.

Proof. 1. Let ϕ^* be any other level-α test of $H_0: \ \theta \in \Omega_0$ vs. g. Then ϕ^* is also a level-α test for H_Λ vs. g, because

$$\mathbb{E}_\theta \phi^*(X) = \int \phi^*(x) f_\theta(x)\, d\mu(x) \le \alpha, \ \forall\ \theta \in \Omega_0,$$

which implies that

$$\int \phi^*(x) h_\Lambda(x)\, d\mu(x) = \int\int \phi^*(x) f_\theta(x)\, d\mu(x) d\Lambda(\theta) \le \int \alpha d\Lambda(\theta) = \alpha.$$

Since ϕ_Λ is MP for H_Λ vs. g, we have

$$\int \phi^*(x)g(x)\,d\mu(x) \le \int \phi_\Lambda(x)g(x)\,d\mu(x),$$

Hence ϕ_Λ is a MP test for H_0 vs. g, because ϕ_Λ is also level α.

2. Let Λ' be any distribution on Ω_0. Since $\mathbb{E}_\theta \phi_\Lambda(x) \le \alpha$, $\forall\ \theta \in \Omega_0$, we know that ϕ_Λ must be level-α for $H_{\Lambda'}$ vs. g. Thus $\beta_\Lambda \le \beta_{\Lambda'}$, so Λ is the least favorable distribution. □

14.2.3 Examples

Example 1 (Testing in the presence of nuisance parameters)**.** Let $X_1, \ldots, X_n$ be i.i.d. $\mathcal{N}(\theta, \sigma^2)$, where both θ, σ^2 are unknown. We consider testing $H_0 :\ \sigma \le \sigma_0$ against $H_1 :\ \sigma > \sigma_0$. To find a UMP test, we follow the previously mentioned strategy:

1. First we fix a simple alternative (θ_1, σ_1) for some arbitrary θ_1 and $\sigma_1 > \sigma_0$.

2. Second, we choose a prior distribution Λ to collapse our null hypothesis over. Intuitively, the least favorable prior should make the alternative hypothesis hard to distinguish. Hence, a rule of thumb consists in concentrating Λ on the boundary between H_1 and H_0 (i.e. the line $\{\sigma = \sigma_0\}$). Thus Λ will be a probability distribution over $\theta \in \mathbb{R}$ for the fixed $\sigma = \sigma_0$.

   Another useful observation is that, given any test function $\phi(x)$ and a sufficient statistic T, there exists a test function η that has the same power as ϕ but depends on x only through T:

   $$\eta(T(x)) = \mathbb{E}[\phi(x)|T(x)].$$

   Hence, we can restrict our attention to the sufficient statistics (Y, U), where $Y = \bar{X}$ and $U = \sum_{i=1}^n (X_i - \bar{X})^2$. We know that $Y \sim \mathcal{N}(\theta, \sigma^2/n)$, $U \sim \sigma^2\chi^2_{n-1}$, and Y is independent of U by Basu's theorem.

   Thus, for Λ supported on $\sigma = \sigma_0$, we obtain the joint density of (Y, U) under H_Λ as

   $$c_0 u^{\frac{n-3}{2}} \exp\left(-\frac{u}{2\sigma_0^2}\right) \int \exp\left(-\frac{n}{2\sigma_0^2}(y-\theta)^2\right) d\Lambda(\theta)$$

   and the joint density under alternative hypothesis (θ_1, σ_1) as

   $$c_1 u^{\frac{n-3}{2}} \exp\left(-\frac{u}{2\sigma_1^2}\right) \exp\left(-\frac{n}{2\sigma_1^2}(y-\theta_1)^2\right).$$

   From the above observations, we see that the choice of Λ only affects the distribution of Y. To achieve minimal maximum power against the alternative (i.e., to be least favorable), we need to choose Λ such that the two distributions become as close as possible. Under the alternative hypothesis, $Y \sim \mathcal{N}\left(\theta_1, \frac{\sigma_1^2}{n}\right)$. Under H_Λ, the distribution of Y is in a convolution form, i.e., $Y = Z + \Theta$ for $Z \sim \mathcal{N}\left(0, \frac{\sigma_0^2}{n}\right)$, $\Theta \sim \Lambda$, where Z and Θ are independent. Hence, if we choose $\Theta \sim \mathcal{N}\left(\theta_1, \frac{\sigma_1^2 - \sigma_0^2}{n}\right)$, Y will have the same distribution under the null and the alternative, which is $\mathcal{N}\left(\theta_1, \frac{\sigma_1^2}{n}\right)$. Under this choice of prior, the LRT rejects for large values of $\exp\left(-\frac{u}{2\sigma_1^2} + \frac{u}{2\sigma_0^2}\right)$, i.e., it rejects for large values of u (since $\sigma_1 > \sigma_0$). So the MP test rejects H_Λ if $\sum_{i=1}^n (X_i - \bar{X})^2$ lies above some threshold determined by the size constraint. In particular, it rejects if $\sum_{i=1}^n (X_i - \bar{X})^2 > \sigma_0^2 C_{n-1,1-\alpha}$, where $C_{n-1,1-\alpha}$ is the $(1-\alpha)^{th}$ quantile of χ^2_{n-1}.

3. Next we check if the MP test is level-α for the composite null. For any (θ, σ) with $\sigma \leq \sigma_0$, the probability of rejection is:

$$\mathbb{P}_{\theta,\sigma}\left(\frac{\sum_{i=1}^n (X_i - \bar{X})^2}{\sigma^2} > \frac{\sigma_0^2\ C_{n-1,1-\alpha}}{\sigma^2}\right) = \mathbb{P}\left(\chi^2_{n-1} > \frac{\sigma_0^2}{\sigma^2} C_{n-1,1-\alpha}\right) \leq \alpha,$$

while equality holds iff $\sigma = \sigma_0$. Hence, it follows from Theorem 1 that our test is MP for testing the original null H_0 vs. $\mathcal{N}(\theta_1, \sigma_1)$.

4. Finally, the MP level-α test for testing the composite null H_0 vs. an arbitrarily chosen (θ_1, σ_1) does not depend on the choice of (θ_1, σ_1). Hence it is UMP for testing the original composite null vs. the composite alternative.

Example 2 (Nonparametric Quality Checking)**.** Identical light bulbs have lifetime $X_1, \ldots, X_n$ with an arbitrary distribution $\mathbb{P}$ over $\mathbb{R}$. Let u be a fixed threshold for a satisfactory lifetime and $\mathbb{P}(X \leq u)$ be the probability of a given light bulb being unsatisfactory. Given the data of sample lifetimes we may be interested in testing whether the probability of having an unsatisfactory light bulb is too large:

$$H_0 : \mathbb{P}(X \leq u) \geq p_0 \quad \text{vs.} \quad H_1 : \mathbb{P}(X \leq u) < p_0.$$

Here p_0 is a fixed quality parameter.

0. Before we start our search for the UMP test, let us reparametrize the distribution $\mathbb{P}$ as follows. Let $\mathbb{P}^-$ and $\mathbb{P}^+$ be the conditional distributions of $X|X \leq u$ and $X|X > u$ respectively, and let $p = \mathbb{P}(X \leq u)$. Then, $\mathbb{P}$ has a one-to-one correspondence with $(\mathbb{P}^+, \mathbb{P}^-, p)$. For any fixed $\mathbb{P}$, let p^- and p^+ be the conditional densities of $\mathbb{P}^-$ and $\mathbb{P}^+$ with respect to some measure μ (existence of the densities and base measure can be justified, e.g. by Radon-Nikodym theorem in measure theory). The joint density of $X_1, \ldots, X_n$ at values $x_1, \ldots, x_n$ when $x_{i_1}, \ldots, x_{i_m} \leq u < x_{j_1}, \ldots, x_{j_{n-m}}$ is then given by

$$p^m \left(\prod_{j=1}^m p^-(x_{i_j})\right) (1-p)^{n-m} \left(\prod_{k=1}^{n-m} p^+(x_{j_k})\right).$$

1. As before, we fix a simple alternative $(\mathbb{P}^-, \mathbb{P}^+, p_1)$ where $p_1 < p_0$.

2. We next choose a proper prior. We guess that Λ mostly concentrates on the boundary point (p^+, p^-, p_0). If so, for testing H_Λ vs. the simple alternative, the LRT rejects for large values of

$$\frac{\binom{n}{m} p_1^m (1-p_1)^{n-m}}{\binom{n}{m} p_0^m (1-p_0)^{n-m}},$$

which is equivalent to testing $Bin(n, p_0)$ vs. $Bin(n, p_1)$. Thus, the MP test, which rejects for small values of $m = \#\{i : X_i \le u\}$[1], is given by

$$\phi_\Lambda(x) = \begin{cases} 1, & \text{if } m < k \\ \gamma, & \text{if } m = k \\ 0, & \text{if } m > k, \end{cases}$$

where k and γ are both determined by the level constraint $\mathbb{E}_{p_0}\phi_\Lambda(x) = \alpha$.

3. Now we check if ϕ_Λ is level-α for our composite null H_0. Note that the power function of ϕ_Λ depends on $\mathbb{P}$ only through $p = \mathbb{P}(X \le u)$. Given that this family has MLR in m, the power function would be monotone. So for any $p > p_0$, the rejection probability under the null is still smaller than α. Hence, ϕ_Λ is the MP test for testing the composite null H_0 against the simple alternative $H_1 : (\mathbb{P}^-, \mathbb{P}^+, p_1)$.

4. Finally, ϕ_Λ has no dependence on the choice of alternative hypothesis. Therefore, ϕ_Λ is UMP for testing the composite null H_0 against the composite alternative H_1.

[1] This test is called *sign test* since it only depends on $\text{sign}(X_i - u)$.