

Lecture 15 — November 12

Lecturer: Lester Mackey

Scribe: Pengfei Gao, Qijia Jiang



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

15.1 Beyond UMP Testing

We began our study of hypothesis testing and strategies for how to find a uniformly most powerful (UMP) test in the following cases:

- simple null versus simple alternative using the NP lemma
- one-sided tests using monotone likelihood ratios
- a general strategy, where possible, in situations featuring composite nulls and composite alternatives.

We have seen that UMP tests need not exist, such as when $p_\theta \sim \mathcal{N}(\theta, \sigma^2)$ and we want to test $\theta = \theta_0$ against $\theta \neq \theta_0$. When uniform optimality is not achievable, we have a variety of alternative optimality strategies, involving the constraining or collapsing of the risk function, at our disposal. These strategies parallel the approaches we took in the estimation setting.

15.1.1 Collapse the power function: Maximize the average power

One alternative to maximizing the power function uniformly is to maximize the average power under some prior distribution.

Let $X \sim \mathcal{P}_\theta$, $H_0 : \theta \in \Omega_0$ versus $H_1 : \theta \in \Omega_1$, and let Λ be a probability distribution over Ω_1 . We can choose ϕ to maximize average power

$$\int_{\Omega_1} \mathbb{E}_\theta \phi(X) d\Lambda(\theta) = \int_{\Omega_1} \int_{\mathcal{X}} \phi(x) p_\theta(x) d\mu(x) d\Lambda(\theta) = \int_{\mathcal{X}} \phi(x) \int_{\Omega_1} p_\theta(x) d\Lambda(\theta) d\mu(x). \quad (15.1)$$

If we define the marginal mixture distribution $g(x) = \int p_\theta(x) d\Lambda(\theta)$ then our problem has been reduced to testing H_0 versus g .

15.1.2 Constrain by enforcing unbiasedness

Definition 1 (Unbiasedness). Let $\alpha \in [0, 1]$. A test ϕ is **unbiased** level- α if

$$\forall \theta_1 \in \Omega_1 \quad \mathbb{E}_{\theta_1} \phi(X) \geq \alpha \quad \text{and} \quad \forall \theta_0 \in \Omega_0 \quad \mathbb{E}_{\theta_0} \phi(X) \leq \alpha.$$

Unbiasedness enforces the appealing property that the probability of rejection is greater under any alternative distribution than it is under any null distribution. A uniformly most powerful test is always unbiased if it exists.

While this notion of unbiasedness differs from the definition we encountered when discussing point estimation, we can check that this is actually a special case of risk unbiasedness when the loss function L is such that $L(\theta_0, \text{reject}) = 1 - \alpha$ and $L(\theta_1, \text{accept}) = \alpha$.

15.1.3 Constrain by enforcing invariance

Let $X_1, \dots, X_n \stackrel{\text{iid}}{\sim} \mathcal{N}(\theta, \sigma^2)$ for σ, θ both unknown, and test $H_0 : \theta = 0$ versus $H_1 : \theta \neq 0$. For $i \in \{1, \dots, n\}$, let $X'_i = cX_i$ with $c > 0$. Then $\mathbb{E}(X'_i) = \theta' = c\theta$. Since testing $\theta = 0$ is equivalent to testing $\theta' = 0$, it is natural to impose the **invariance constraint**

$$\forall c > 0 \quad \phi(X) = \phi(cX). \quad (15.2)$$

Such a test is unaffected by arbitrary rescaling of the data (which might occur when changing units from centimeters to meters for example). There are cases when a UMP does not exist but a UMP test among the invariant ϕ exists (a topic for next week).

15.1.4 Collapse the power function: Maximize worst case power

We could alternatively consider the problem of maximizing the worst case power of a test. In this case, we will maximize the minimum power over $\theta_1 \in \Omega_1$ subject to the standard constraint that our size is no larger than our level α . A test of this form is called **maximin**.

15.1.5 Constrain by enforcing monotonicity

Let X, Y be independent, $X \sim \mathcal{N}(\theta_X, 1)$ and $Y \sim \mathcal{N}(\theta_Y, 1)$ for θ_X, θ_Y unknown, and test $H_0 : \theta_X \leq 0, \theta_Y \leq 0$.

A **monotonicity restriction** implies that if ϕ rejects upon observing (x, y) , then it should also reject for (x', y') where $x' > x$ and $y' > y$.

15.2 Uniformly most powerful unbiased tests

Today, we will focus on finding uniformly most powerful unbiased (UMPU) tests in settings in which UMP tests do not exist. These tests often exist for testing $\theta_1 \leq \tilde{\theta}$ vs $\theta_1 > \tilde{\theta}$ in the presence of nuisance parameters $(\theta_2, \dots, \theta_{k+1})$ and for testing $\theta = \tilde{\theta}$ vs $\theta \neq \tilde{\theta}$.

15.2.1 General Setting

Let us test $H_0 : \theta \in \Omega_0$ vs $H_1 : \theta \in \Omega_1$. Typically, we take Ω_0, Ω_1 to be subsets of a Euclidean space, and we introduce ω the **common boundary** between Ω_0 and Ω_1 :

$$\omega = \bar{\Omega}_0 \cap \bar{\Omega}_1.$$

That is, ω is the intersection of the closures of Ω_0 and Ω_1 (closed under limits).

Example 1. If we are testing $H_0: \theta = \tilde{\theta}$, $H_1: \theta \neq \tilde{\theta}$, then $\omega = \Omega_0 = \{\tilde{\theta}\}$.

Example 2. If we are testing $H_0: \theta_1 \leq \tilde{\theta}$ vs $H_1: \theta_1 > \tilde{\theta}$ in the presence of nuisance parameters $(\theta_2, \dots, \theta_{k+1})$, then $\omega = \{\theta = (\theta_1, \dots, \theta_{k+1}) \in \mathbb{R}^{k+1} : \theta_1 = \tilde{\theta}\}$.

Generally, if the power function $\theta \mapsto \beta_\phi(\theta)$ is continuous in θ (as is the case for any canonical form exponential family on the natural parameter space), then ϕ unbiased and of level α implies that $\beta_\phi(\theta) = \alpha$ for all $\theta \in \omega$. We have a name for tests that match the level on the boundary.

Definition 2 (α -similarity). A test ϕ satisfying $\mathbb{E}_\theta \phi(X) = \alpha$ for all $\theta \in \omega$ is called **α -similar** on ω .

The following lemma tells us we can find a UMPU test by looking only at α -similar tests.

Lemma 1 (TSH 4.1.1). If $\theta \mapsto \beta_\phi(\theta)$ is continuous (in θ) on Ω for all ϕ , and ϕ_0 is a UMP test amongst α -similar level- α tests, then ϕ_0 is UMPU at level α .

Proof. Firstly, because ϕ_0 is UMP α -similar tests, it is at least as powerful as $\phi_\alpha(X) \equiv \alpha$, and the power of ϕ_0 on Ω_1 is therefore $\geq \alpha$. Hence, ϕ_0 is unbiased.

Secondly, an unbiased level- α test must, by definition, have expectation value $\leq \alpha$ for $\theta \in \Omega_0$ and $\geq \alpha$ for $\theta \in \Omega_1$. By continuity such a test must have expectation α on the common boundary. Therefore, the set of unbiased level- α tests is a subset of α -similar level- α tests, amongst which ϕ_0 is most powerful. Hence, ϕ_0 is also as powerful as any unbiased level- α test. ϕ_0 is UMPU. \square

15.2.2 Two-sided Testing without Nuisance Parameters

Let us test $H_0: \theta = \theta_0$ vs. $H_1: \theta \neq \theta_0$, when X is distributed according some member of the one-dimensional exponential family

$$p_\theta(x) = h(x) \exp(\theta T(x) - A(\theta))$$

We have seen that no UMP test exists in the normal case. Our goal here is to find a UMPU test.

Since we are working with an exponential family, the power function is continuous, and, by Lemma 1, it suffices to find a UMP level α test amongst α -similar tests. Since $\omega = \Omega_0$, any UMP α -similar test ϕ has

$$\beta_\phi(\theta_0) = \mathbb{E}_{\theta_0} \phi(X) = \alpha, \quad (15.3)$$

and

$$\beta_\phi(\theta_0) \leq \beta_\phi(\theta) \text{ for all } \theta \in \mathbb{R} \quad (15.4)$$

since $\phi_\alpha(x) \equiv \alpha$ is also α -similar.

Since θ_0 minimizes β_ϕ , and β_ϕ is differentiable with derivative $\beta'_\phi(\theta) = \int \phi(x) \frac{d}{d\theta} p_\theta(x) d\mu(x)$,¹ we have the constraint

$$0 = \beta'_\phi(\theta_0) = \int \phi(x) \frac{d}{d\theta} p_{\theta_0}(x) d\mu(x) \quad (15.5)$$

for any UMP α -similar test. Hence, it suffices to find a UMP test satisfying (15.3) and (15.5). We have learned to find UMP tests under a single level constraint, but how do we find a UMP test under multiple constraints? We develop the tools in the next section.

15.3 Method of Undetermined Multipliers (MoUM)

To maximize power subject to multiple constraints, we generalize the Neyman-Pearson lemma.

Lemma 2 (TSH Lemma 3.6.1). Suppose F_1, \dots, F_{m+1} are real-valued functions defined on a common domain U . We will maximize $F_{m+1}(u)$ subject to constraints of the form

$$F_i(u) = c_i \text{ for } i = 1, \dots, m$$

where c_1, \dots, c_m are known constants. To do this, it suffices to find u_0 that satisfies the constraints and maximizes

$$F_{m+1}(u) - \sum_{i=1}^m k_i F_i(u) \quad (15.6)$$

for any choice of the **undetermined multipliers** k_1, \dots, k_m .

In practice, we maximize $F_{m+1} - \sum_{i=1}^m k_i F_i$ for arbitrary k_i 's, and then choose any solution that satisfies the constraints.

Proof. If u satisfies constraints and u_0 optimizes $F_{m+1} - \sum_{i=1}^m k_i F_i$, then

$$F_{m+1}(u) - \sum_{i=1}^m k_i F_i(u) \leq F_{m+1}(u_0) - \sum_{i=1}^m k_i F_i(u_0).$$

Because u and u_0 satisfy constraints:

$$\sum_{i=1}^m k_i F_i(u) = \sum_{i=1}^m k_i c_i = \sum_{i=1}^m k_i F_i(u_0).$$

This implies that $F_{m+1}(u) \leq F_{m+1}(u_0)$, so u_0 is maximal. \square

¹Here, $\frac{d}{d\theta} p_\theta(x)$ is the derivative of $(x, \theta) \mapsto p_\theta(x)$ with respect to the second variable and taken at the point (x, θ) . Reference: TSH Thm 2.7.1

15.3.1 MoUM for Test Functions

Now, we will apply MoUM to the case where U is space of test functions ϕ :

$$F_i(\phi) = \int \phi(x) f_i(x) d\mu(x).$$

Our goal is to maximize $\int \phi(x) f_{m+1}(x) d\mu(x)$ subject to $\int \phi(x) f_i(x) d\mu(x) = c_i$.

First, maximize

$$F_{m+1}(\phi) - \sum_i k_i F_i(\phi) = \int \phi(x) \left(f_{m+1}(x) - \sum_{i=1}^m k_i f_i(x) \right) d\mu(x).$$

Any solution has the form

$$\phi(x) = \begin{cases} 1 & \text{if } f_{m+1}(x) > \sum_i k_i f_i(x) \\ 0 & \text{if } f_{m+1}(x) < \sum_i k_i f_i(x) \end{cases}.$$

Eventually, we will choose k_i 's to ensure that all constraints are satisfied.

15.3.2 Application of MoUM to our 2-sided testing problem

In this setting, $H_0 : \theta = \theta_0$. We will fix a simple alternative $\theta = \theta' \neq \theta_0$ and hope that our best test has no θ' dependence. We would like to maximize power $\int \phi(x) p_{\theta'}(x) d\mu(x)$ subject to

$$\int \phi(x) p_{\theta_0}(x) d\mu(x) = \alpha \quad (15.7)$$

$$\int \phi(x) \frac{d}{d\theta} p_{\theta_0}(x) d\mu(x) = 0. \quad (15.8)$$

For a 1-parameter exponential family, we have

$$p_{\theta}(x) = h(x) e^{\theta T(x) - A(\theta)} \quad \text{and} \quad (15.9)$$

$$\frac{d}{d\theta} p_{\theta}(x) = h(x) e^{\theta T(x) - A(\theta)} (T(x) - A'(\theta)) = p_{\theta}(x) (T(x) - \mathbb{E}_{\theta}[T(X)]). \quad (15.10)$$

Applying the reasoning from the previous section, we find that a most powerful test has rejection region defined by

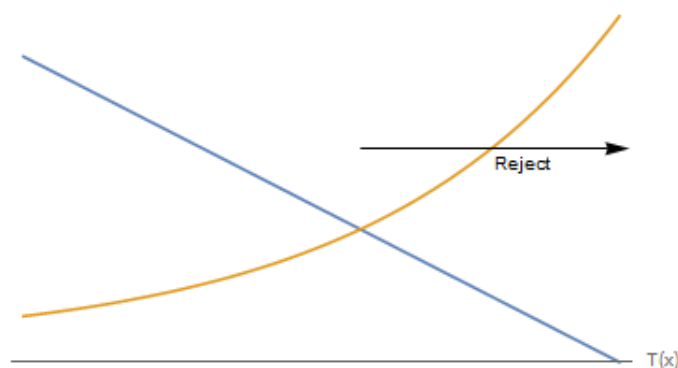
$$p_{\theta'}(x) > k_1 p_{\theta_0}(x) + k_2 \frac{d}{d\theta} p_{\theta_0}(x)$$

for some values of k_1 and k_2 , which is equivalent to

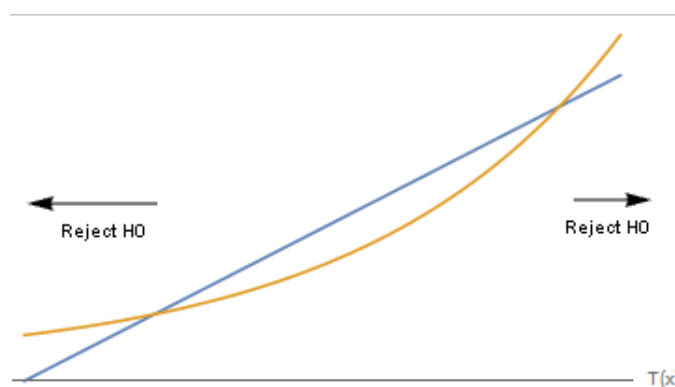
$$\frac{e^{(\theta' - \theta_0)T(x)}}{k'_1 + k'_2 T(x)} > \text{const}$$

with some rearranging.

Now consider the set of values of $T(x)$ satisfying this constraint. Because the constraint is that an exponential function exceeds a linear function, the set of values of $T(x)$ satisfying this constraint is either a one-sided interval



or all points outside a closed interval.



The first possibility will not give rise to an unbiased test, because the result would be a one-sided test with monotone power functions. Therefore any optimal ϕ is of the form

$$\phi(x) = \begin{cases} 1 & \text{if } T(x) > C_1 \text{ or } T(x) < C_2 \\ \gamma_i & \text{if } T(x) = C_i \\ 0 & \text{otherwise} \end{cases}.$$

A simplification is possible if $T(x)$ is symmetrically distributed under θ_0 . Then the optimal test rejects whenever $|T(x)| > \text{const}$. Such tests are called **equitailed tests**.

Example 3. Suppose $X_1, \dots, X_n \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$ with $H_0 : \sigma = \sigma_0$ and $H_1 : \sigma \neq \sigma_0$. The optimal test has an acceptance region of the form

$$C_1 \leq \frac{\sum_i X_i^2}{\sigma_0^2} \leq C_2$$

The middle expression is a sufficient statistic that is χ_n^2 distributed under H_0 . How do we choose C_1 and C_2 ? Let f_n be density χ_n^2 . The level constraint is:

$$P(C_1 \leq T(X) \leq C_2) = \int_{C_1}^{C_2} f_n(y) dy = 1 - \alpha,$$

and the derivative constraint (15.8) with substitution (15.10) gives

$$\mathbb{E}_{\theta_0}[T(X)\phi(X)] = \mathbb{E}_{\theta_0}[T(x)]\mathbb{E}_{\theta_0}[\phi(X)].$$

On the left side $\phi(x) = 1$ on the complement of (C_1, C_2) , and on the right, the mean of a χ_n^2 is n , and the level is α . Thus,

$$\int_{(C_1, C_2)^c} y f_n(y) = n\alpha.$$

For χ_n^2 distributions, $y f_n(y) = n f_{n+1}(y)$, and the derivative constraint ultimately becomes

$$\int_{C_1}^{C_2} f_{n+2}(y) = 1 - \alpha.$$

We have two integral equations, and we can solve them for the unknown boundaries C_1, C_2 .