STAT 210A: Theoretical Statistics	**Fall 2007**

Lecture 19 — November 1st

Lecturer: M. Wainwright *Scribe: Ori Stitelman*

This is the danger environment.
Warning: These scribe notes have only been mildly proofread.

19.1 Outline

1. Uniformly Most Powerful Tests (UMP)
2. Generalized Likelihood Ratio

Reading: Keener, Chapter 14

19.2 Uniformly Most Powerful Tests (UMP)

Motivation: Frequently have composite hypothesis tests and want to test whether in one set versus the other set:

$$H_0 : \theta \in \Omega_0$$
$$H_1 : \theta \in \Omega_1$$
$$\Omega_1 \cap \Omega_1 = 0$$

Example: Testing whether a coin is fair versus not fair is an example of a composite hypothesis test where $X \sim Bin(n, \theta)$ and:

$$\Omega_0 = \{1/2\}$$
$$\Omega_1 = [0, 1] \setminus \{1/2\}$$

Definition: A test δ of size α is uniformly most powerful (UMP) if $\forall$ other tests ϕ of size α we have $\beta_\delta(\theta) \geq \beta_\phi(\theta) \, \forall \theta \in \Omega_1$. Note: These tests do not always exist.

Example: Continuing the coin flipping example above with the following hypotheses:

$$\widetilde{H}_0 : \theta = 1/2$$
$$\widetilde{H}_1 : \theta = \theta_1$$

We are in search of a test that for any choice of θ_1 the test is most powerful. By Neyman Pearson presented in last lecture the Likelihood Ratio Test (LRT) is most powerful:

$$T(x) = \frac{(1/2)^x (1/2)^{n-x}}{\theta_1^x (1-\theta_1)^{n-x}} = \left(\frac{1-\theta_1}{\theta_1}\right)^x \frac{(1/2)^n}{(1-\theta_1)^n}$$

Setting a threshold on $T(x)$ is a most powerful test for level α but there is no uniformly most powereful test for $\{1/2\}$ versus $[0,1] \setminus \{1/2\}$. This can be seen by letting $\theta_1 = 3/4$ and observing as x increases one is more likely to accept $\widetilde{H}_1$ and then letting $\theta_1 = 1/4$ and as x increases one is more likely to accept $\widetilde{H}_0$. Thus there is no UMP. This suggests that some special structure is needed for an UMP to exist.

Definition: A family of densities/PMFs $\{p(\cdot;\theta) | \theta \in \Omega\}$ has monotone likelihood ratio (MLR) in $T(x)$ s.t. $\forall \theta_0 < \theta_1$. Where $\Omega \subseteq \Re$. Then LRT:

$$L(x) = \frac{p(x;\theta_1)}{p(x;\theta_0)}$$

is a non-decreasing function of T(x).

Note: If have MLR and one sided test then have structure for UMP to exist.

Example: LRT from coin flipping example above is MLR.

Example: (1-D Exponential Family)
Say $p(x;\theta) = h(x)\, exp\{\eta(\theta) T(x) - A(\theta)\}$ Then the likelihood ratio is:

$$L(x) = \frac{p(x;\theta_1)}{p(x;\theta_0)} = exp\{(\eta(\theta_1) - \eta(\theta_0)) T(x) - A(\theta_1) + A(\theta_0)\}$$

and if $\eta()$ is monotone then the family is MLR. $\eta()$ is monotone if $\theta_0 < \theta_1 \Rightarrow \eta(\theta_0) \leq \eta(\theta_0)$. This case covers the bernoulli/binomial example shown above:

$$\eta(\theta) = \log\left(\frac{\theta}{1-\theta}\right)$$

Which is a monotone function of θ. Also normal location models (where $\eta(\theta) = \theta$), poisson, and exponential distributions are of this form.

Example: Distributions which are not in the exponential family can also have MLR. For example let $X_1, \ldots, X_n \sim Uni[0,\theta]$. Then:

$$p(X_1, \ldots, X_n) = \left(\frac{1}{\theta}\right)^n \mathbf{1}\left(X_{(n)} \leq \theta\right)$$

for $x_i \geq 0$ and the likelihood ratio is:

$$L(x) = \frac{p(\vec{x};\theta_1)}{p(\vec{x};\theta_0)} = \left(\frac{\theta_0}{\theta_1}\right)^n \frac{\mathbf{1}\left(X_{(n)} \leq \theta_1\right)}{\mathbf{1}\left(X_{(n)} \leq \theta_0\right)}$$

for $\theta_1 > \theta_0$, this is monotonic in $T(\vec{x}) = X_{(n)}$ for all $(X_1, \ldots, X_n)$ s.t. $X_{(n)} \leq \theta_1$.

Up until now we have discussed the structure of the likelihood ratio for a UMP to exist. Now lets examine the structure needed on the testing regions.

One Sided Tests: The following are examples of a one sided tests:

$$\begin{aligned} H_0 &: \theta = \theta_0 \\ H_1 &: \theta > \theta_0 \end{aligned}$$

$$\begin{aligned} H_0 &: \theta = \theta_0 \\ H_1 &: \theta < \theta_0 \end{aligned}$$

Theorem 19.1. *(UMP for one sided problems) Say family P is MLR in T(x), and consider test:*

$$\begin{aligned} H_0 &: \theta \leq \theta_0 \\ H_1 &: \theta > \theta_0 \end{aligned}$$

Then $\exists$ uniformly most powerful test of size α. With

$$\delta(x) = \begin{cases} 1 & \textit{if } T(x) > C \\ \gamma & \textit{if } T(x) = C \\ 0 & \textit{if } T(x) < C \end{cases}$$

For suitable constants C, and γ s.t. $\beta_\delta(\theta_0) = \alpha$. Where γ is a randomized result.

Thus the power at the split point equals α which insures the power of the test equals α. Furthermore, due to monotonicity, $\delta(x)$ is a likelihood ratio test in disguise.

Proof: For any $\theta_1 > \theta_0$, let's find a UMP test for θ_1 versus θ_0, by Neyman-Pearson can choose constant C, γ independent of θ_1 and θ_0 s.t. a test has the following form:

$$\delta(x) = \begin{cases} 1 & \text{if T(x)} > \text{C} \\ \gamma & \text{if T(x)} = \text{C} \\ 0 & \text{if T(x)} < \text{C} \end{cases}$$

with $\beta_\delta(\theta_0) = \alpha$. By MLR, can equivalently write as LRT in $\delta'(X)$:

$$\delta'(x) \equiv \delta(x) = \begin{cases} 1 & \text{if L(x)} > \text{C'} \\ \gamma' & \text{if L(x)} = \text{C'} \\ 0 & \text{if L(x)} < \text{C'} \end{cases}$$

where C' is a function of θ_1 and θ_0 and $\beta_{\delta'}(\theta_0) = \alpha$. Thus, this is an UMP test, since for any $\theta_1 > \theta_0$, it is a LRT with level α. $\square$

19.3 Generalized Likelihood Ratio Tests

Considering the following composite hypothesis test:

$$H_0 : \theta \in \Omega_0$$
$$H_1 : \theta \in \Omega_1$$
$$\Omega_1 \cap \Omega_1 = 0$$

A natural extension of an LRT is the generalized LRT:

$$L(x) = \frac{\sup_{\theta \in \Omega_1 \cup \Omega_0} [p(x;\theta)]}{\sup_{\theta \in \Omega_0} [p(x;\theta)]} \geq 1$$

Where the numerator is typically unconstrained maximum likelihood and the denominator constrained maximum likelihood. The rationale for the GLR can be explained by the following plot:

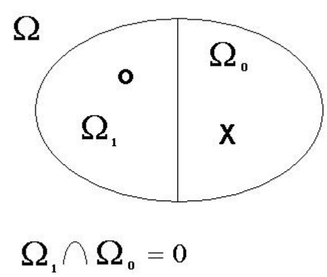

So if $\theta \in \Omega_0$ the overall maximum likelihood estimate is more likely to be in Ω_0, say the X in the plot above, and thus L(X) $\approx$ 1. Alternatively if $\theta \in \Omega_1$ the overall maximum likelihood estimate is more likely to be in Ω_1, say the circle in the above plot, and thus the L(X)>> 1. Furthermore, A GLR test often yields a UMP test when a UMP test exists.
Example: Say $X_i \sim N(\mu, \sigma^2)$ for i= 1,...,n IID where both parameters are unknown. Where $\Omega = \{(\mu, \sigma^2) \,|\mu \in \Re, \sigma^2 > 0\}$ and want to test $\Omega_0 = \{\mu = 0, \sigma^2 = 1\}$ versus $\Omega_1 = \{\Omega \backslash \Omega_0\}$. So the test is whether the data are produced from a standard normal distribution or some other normal distribution.

The denominator of the GLR is as follows:

$$\sup_{\theta \in \Omega_0} [p(x;\theta)] = \frac{1}{(2\pi)^{n/2}} \exp\left\{ \frac{-1}{2} \sum_{i=1}^{n} (x_i)^2 \right\}$$

And The numerator:

$$\sup_{\theta\in\Omega_1\cup\Omega_0}\left[p\left(x;\theta\right)\right]=\frac{1}{\left(2\pi\hat{\sigma}\right)^{n/2}}\exp\left\{\frac{-1}{2\hat{\sigma}^2}\sum_{i=1}^{n}\left(x_i-\hat{\mu}\right)^2\right\}$$

Where $\hat{\mu}$ and $\hat{\sigma}^2$ are the maximum likelihood estimates.

$$\hat{\mu}=\frac{1}{n}\sum_{i=1}^{n}x_i$$
$$\hat{\sigma}^2=\frac{1}{n}\sum_{i=1}^{n}\left(x_i-\hat{\mu}\right)^2$$

Thus, the likelihood ratio is:

$$L\left(x\right)=\left(\hat{\sigma}^2\right)^{-n/2}\exp\left\{\frac{1}{2}\sum_{i=1}^{n}x_i^2-\frac{1}{2\hat{\sigma}^2}\sum_{i=1}^{n}\left(x_i-\hat{\mu}\right)^2\right\}$$

and the GLRT is a test based on thresholding the above quantity. In practice this would be done by choosing a t_α s.t. $P_{\theta_0}\left(L\left(x\right)>t_\alpha\right)=\alpha$.

After a bit of algebra:

$$+2\log L\left(x\right)=n\log\left(\hat{\sigma}^2\right)-n+\sum_{i=1}^{n}x_i^2$$

Next lecture we will look at thresholding based on the above quantity and asymptotically arriving at chi-squared distribution.