Session 10

Recall...

A typical ROC appears as follows:

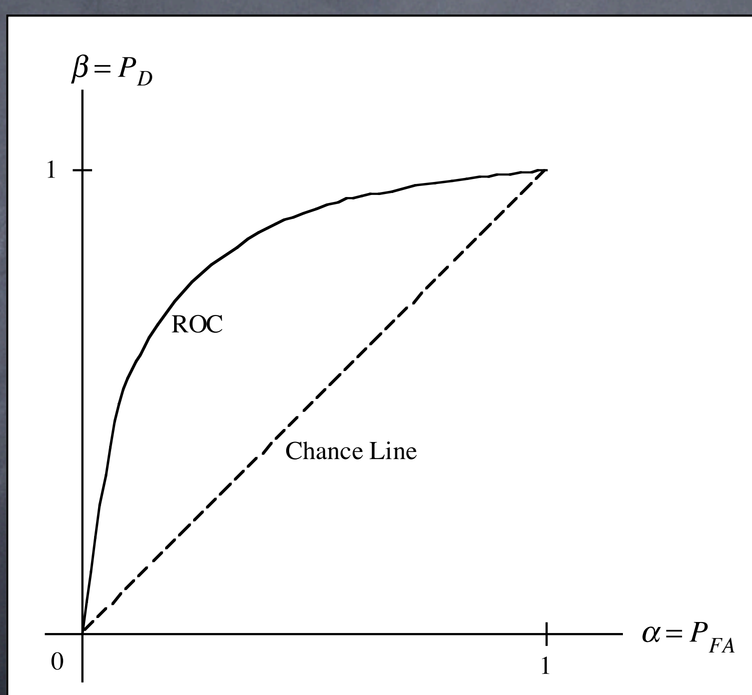

Recall...

Properties of the ROCs of Likelihood Ratio Tests

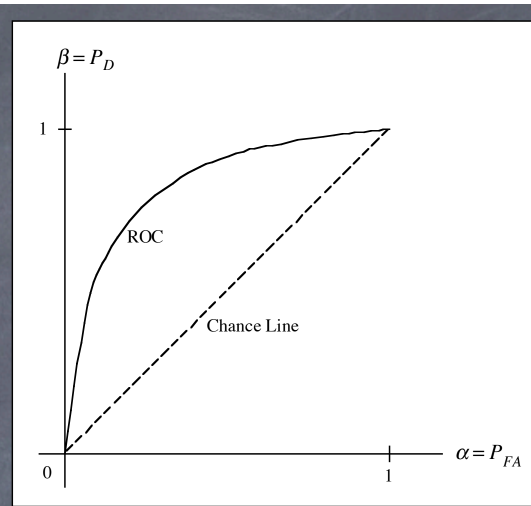

1. All continuous likelihood ratio tests have ROCs that are convex downward.

2. Points on the *chance line* ($\beta = \alpha$) can be achieved without observing any data by picking hypothesis H_1 at random with probability α. (i.e., flipping a biased coin with probability α of coming up "heads").

Recall...

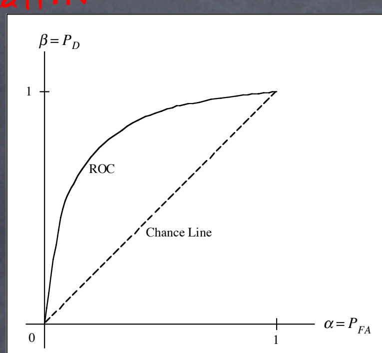

Properties of the ROCs of Likelihood Ratio Tests (Continued)

3. All continuous likelihood ratio tests have ROC's that are above the chance line. This is a consequence of property 1, because the points $(\alpha, \beta) = (0, 0)$ and $(\alpha, \beta) = (1, 1)$ are contained on all ROC's.

4. The slope of the ROC at any particular point is given by the threshold achieving that operating point, i.e.,

$$\frac{d\beta}{d\alpha} = \frac{d\beta/dk}{d\alpha/dk} = \frac{f_{L,\theta_1}(\ell)}{f_{L,\theta_0}(\ell)} = k \geq 0$$

Proof of Property 1

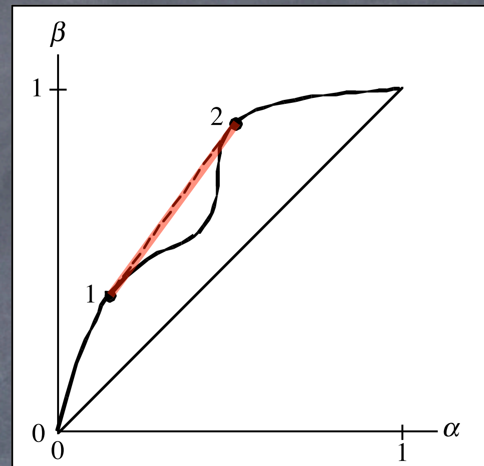

Suppose that the ROC of a LRT is non-concave as shown.

Threshold k_1 achieves point 1

Threshold k_2 achieves point 2

A randomized test that selects k_1 with probability p and and k_2 with probability $1-p$ can achieve any point on the line connecting *point 1* and *point 2*.

But this line is above the LRT's ROC—contradicting the optimality of the LRT. Hence the ROC must be concave downward. ∎

Testing Composite Hypotheses

Recall that when H_0 is composite

$$\alpha = \sup_{\underline{\theta} \in \Theta_0} \left\{ \mathrm{E}_{\underline{\theta}}[\phi(\underline{X})] \right\}$$

and when H_1 is composite

$$\beta(\underline{\theta}) = \mathrm{E}_{\underline{\theta}}[\phi(\underline{X})], \quad \underline{\theta} \in \Theta_1$$

Ideally, we would like to be able to design a test $\phi(\underline{X})$ of size α with the largest possible $\beta(\underline{\theta})$ for each pair $(\underline{\theta}_0, \underline{\theta}_1) \in \Theta_0 \times \Theta_1$.

10.6

Uniformly Most Powerful Tests

Definition: A test $\phi(\underline{X})$ of H_0 versus H_1 is *uniformly most powerful* of size α if it has size α and its power $\beta(\underline{\theta}_1)$ is uniformly greater than the power of any test $\phi^{'}(\underline{X})$ with size $\alpha^{'} \leq \alpha$ testing Θ_0 versus $\{\underline{\theta}_1\}$ for any $\underline{\theta}_1 \in \Theta_1$.

Note that in this definition, the test $\phi^{'}(\underline{X})$ must have size $\alpha^{'} \leq \alpha$ (typically $\alpha^{'} = \alpha$).

Furthermore, $\phi(\underline{X})$ has power $\beta \geq \beta^{'}$.

$\phi^{'}(\underline{X})$ is designed assuming knowledge of the particular $\underline{\theta}_1$ in effect, whereas $\phi(\underline{X})$ is designed only with the knowledge that $\underline{\theta}_1 \in \Theta_1$.

10.7

UMP Tests (Continued)

A UMP test cannot be a function of the actual $\underline{\theta}_1 \in \Theta_1$ in effect.

A UMP test must perform as well as any test designed with knowledge of $\underline{\theta}_1$.

When you ask for a UMP test, you are asking for a lot.

In many detection problems, UMP tests do not exist.

In some detection problems, UMP tests do exist.

UMP Tests (Continued)

Property of a UMP test: A UMP test exists if and only if the likelihood ratio test of H_0 versus H_1 can be completely defined without knowledge of the $\underline{\theta}_1 \in \Theta_1$ in effect.

The "if" part of this property is obvious.

The "only if" part can be seen as follows: If we cannot use an LRT, we must come up with another test. This would be inferior to an LRT with knowledge of $\underline{\theta}_1 \in \Theta_1$.
Hence the test cannot be UMP.

Example A

Let $\underline{X} = (X_1, \ldots, X_N)$ be a vector of i.i.d. Gaussian random variables.

Under hypothesis H_0,

$$f_{\theta_0}(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left[\frac{-(x_n - \theta_0)^2}{2}\right], \quad n = 1, \ldots, N.$$

Under hypothesis H_1,

$$f_{\theta_1}(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left[\frac{-(x_n - \theta_1)^2}{2}\right] \quad n = 1, \ldots, N.$$

Assume that $\theta_1 > \theta_0$.

10.10

Thus we have two simple hypotheses:

$$H_0 : \Theta_0 = \{\theta_0\}$$

and

$$H_1 : \Theta_1 = \{\theta_1\}.$$

The likelihood ratio $L(\underline{X})$ is

$$\begin{aligned} L(\underline{X}) &= \frac{f_{\theta_1}(\underline{X})}{f_{\theta_0}(\underline{X})} \\ &= \frac{f_{\theta_1}(X_1) f_{\theta_1}(X_2) \cdots f_{\theta_1}(X_N)}{f_{\theta_0}(X_1) f_{\theta_0}(X_2) \cdots f_{\theta_0}(X_N)} \\ &= \exp\left[-\frac{1}{2}\sum_{n=1}^{N}(X_n - \theta_1)^2 + \frac{1}{2}\sum_{n=1}^{N}(X_n - \theta_1)^2 \right], \\ &= \exp\left[(\theta_1 - \theta_0)\sum_{n=1}^{N} X_n - \frac{1}{2}N(\theta_1^2 - \theta_0^2) \right]. \end{aligned}$$

10.11

The loglikelihood ratio is given by

$$\begin{aligned} \ell(\underline{X}) &= \ln L(\underline{X}) \\ &= (\theta_1 - \theta_0)\sum_{n=1}^{N} X_n - \frac{1}{2}N(\theta_1^2 - \theta_0^2). \end{aligned}$$

Define the statistic $T(\underline{X})$ as

$$T(\underline{X}) = \frac{1}{N}\sum_{n=1}^{N} X_n.$$

We can rewrite the likelihood ratio test as expressed as

$$N(\theta_1 - \theta_0)T(\underline{X}) - \frac{1}{2}N(\theta_1^2 - \theta_0^2) \underset{H_0}{\overset{H_1}{\gtrless}} \ln k,$$

or equivalently

$$T(\underline{X}) \underset{H_0}{\overset{H_1}{\gtrless}} k' = \frac{\ln k + N(\theta_1^2 - \theta_0^2)/2}{N(\theta_1 - \theta_0)}.$$

The sample mean $T(\underline{X})$ is a sufficient statistic for this test, as it contains all of the information in $\underline{X} = (X_1, \ldots, X_N)$. We can write the test as

$$\phi(\underline{X}) = \begin{cases} 1, & \text{for } T(\underline{X}) > k^{'}, \\ 0, & \text{for } T(\underline{X}) \leq k^{'}. \end{cases}$$

To set the threshold $k^{'}$ for a size α test, we note that under H_0, $T(\underline{X})$ is a Gaussian random variable with mean θ_0 and variance $1/N$. So

$$\begin{aligned} \alpha &= \mathrm{E}_{\theta_0}[\phi(\underline{X})] \\ &= P_{\theta_0}(\{T(\underline{X}) > k^{'}\}) \\ &= 1 - \Phi\left(\frac{k^{'} - \theta_0}{1/\sqrt{N}}\right) \\ &= 1 - \Phi\left(\sqrt{N}(k^{'} - \theta_0)\right), \end{aligned}$$

from which it follows that

$$k^{'} = \theta_0 + \frac{1}{\sqrt{N}}\Phi^{-1}(1-\alpha).$$

Under H_1, $T(\underline{X})$ is a Gaussian random variable with mean θ_1 and variance $1/N$. Thus the power of the test is

$$\begin{aligned} \beta &= \mathrm{E}_{\theta_1}[\phi(\underline{X})] \\ &= P_{\theta_1}(\{T(\underline{X}) > k^{'}\}) \\ &= 1 - \Phi\left(\frac{k^{'} - \theta_1}{1/\sqrt{N}}\right) \\ &= 1 - \Phi\left(\sqrt{N}(k^{'} - \theta_1)\right), \\ &= 1 - \Phi\left(\sqrt{N}\left(\theta_0 + \frac{1}{\sqrt{N}}\Phi^{-1}(1-\alpha) - \theta_1\right)\right), \\ &= 1 - \Phi\left(\Phi^{-1}(1-\alpha) - \sqrt{N}(\theta_1 - \theta_0)\right). \end{aligned}$$

Example B

Let $\underline{X} = (X_1, \ldots, X_N)$ be a random sample of i.i.d. Gaussian random variables. Under hypothesis H_0, the pdf of X_n is

$$f_{\theta_0}(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left[\frac{-(x_n - \theta_0)^2}{2}\right], \quad n = 1, \ldots, N,$$

and under hypothesis H_1, the pdf of X_n is

$$f_{\theta_1}(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left[\frac{-(x_n - \theta)^2}{2}\right] \quad n = 1, \ldots, N,$$

where we will assume that $\theta \in (\theta_0, \infty)$, i.e., under hypothesis H_1, θ can take on any value greater than θ_0. Thus we have $\Theta_0 = \{\theta_0\}$ and $\Theta_1 = (\theta_0, \infty)$. So H_0 is a simple hypothesis, and H_1 is a composite hypothesis.

Is there a UMP test of size α in this case? Yes, because for any $\theta \in \Theta_1$, the test in Example A is the most powerful test of size α for testing H_0 versus the simple hypothesis $H_1' : \Theta_1 = \{\theta\}$ for any particular $\theta > \theta_0$. This test is not a function of the particular $\theta \in \Theta_1$ in effect, so the test is UMP, however the power $\beta(\theta)$ is definitely a function of θ:

$$\beta(\theta) = 1 - \Phi(\Phi^{-1}(1-\alpha) - \sqrt{N}(\theta - \theta_0)), \quad , \forall \theta > \theta_0.$$

Example C

Let $\underline{X} = (X_1, \ldots, X_N)$ be a random sample of i.i.d. Gaussian random variables. Under hypothesis H_0, the pdf of X_n is

$$f_{\theta_0}(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left[\frac{-(x_n - \theta_0)^2}{2}\right], \quad n = 1, \ldots, N,$$

and under hypothesis H_1, the pdf of X_n is

$$f_{\theta_1}(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left[\frac{-(x_n - \theta)^2}{2}\right] \quad n = 1, \ldots, N,$$

where we will assume that $\theta \in (-\infty, \theta_0) \cup (\theta_0, \infty)$, i.e., under hypothesis H_1, θ can take on any value other than θ_0. Thus we have $\Theta_0 = \{\theta_0\}$ and $\Theta_1 = \{\theta \in \mathbf{R} : \theta \neq \theta_0\} = (-\infty, \theta_0) \cup (\theta_0, \infty)$. So H_0 is a simple hypothesis, and H_1 is a composite hypothesis.

The unique size α test that achieves maximum power for $\theta_1 > \theta_0$ is

$$\phi_+(\underline{X}) = \begin{cases} 1, & \text{for } T(\underline{X}) > \theta_0 + \frac{1}{N}\Phi^{-1}(1-\alpha), \\ 0, & \text{elsewhere.} \end{cases}$$

By symmetry, the most powerful test of size α for $\theta_1 < \theta_0$ is

$$\phi_-(\underline{X}) = \begin{cases} 1, & \text{for } T(\underline{X}) < \theta_0 - \frac{1}{N}\Phi^{-1}(1-\alpha), \\ 0, & \text{elsewhere.} \end{cases}$$

So we have two different size α tests for testing $H_0 : \{\theta = \theta_0\}$ versus $H_1 : \{\theta \neq \theta_0\}$. $\phi_+(\underline{X})$ is most powerful when $\theta_1 > \theta_0$, whereas $\phi_-(\underline{X})$ is mot powerful when $\theta_1 < \theta_0$. Neither test is uniformly most powerful. We must know $\theta_1 \in \Theta_1$ in order to know which of these two tests to apply. As we noted earlier, a UMP test cannot require such knowledge of θ_1, so a UMP test does not exist in this case.

10.18

The Karlin-Rubin Test

Herman Rubin
Purdue Prof. of Statistics

Karlin-Rubin Theorem: Let X be a scalar random variable with pdf parameterized by a scalar parameter θ. Assume that the likelihood ratio

$$L_{\Theta_1}(x) = \frac{f_{\theta_1}(x)}{f_{\theta_0}(x)} \tag{1}$$

is a non-decreasing function of x for every pair (θ_0, θ_1) such that $\theta_1 > \theta_0$. Then the threshold test

$$\phi(x) = \begin{cases} 1, & \text{for } x > x_0, \\ \gamma, & \text{for } x = x_0, \\ 0, & \text{for } x < x_0, \end{cases} \tag{2}$$

such that

$$\mathrm{E}_0[\phi(X)] = P_{\theta_0}(\{X > x_0\}) + \gamma P_{\theta_0}(\{X = x_0\}) = \alpha$$

is the UMP test of size α for testing $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$.

Proof: See Scharf

10.19

Example of Use of Karlin-Rubin Theorem

Example: *The number of photons detected by an optical detector in time interval $[0, T]$ is a Poisson random variable N with mean θ:*

$$p_\theta(n) = P_\theta(\{N = n\}) = \frac{\theta^k e^{-\theta}}{n!}, \quad n = 0, 1, 2, \ldots,$$

where $\theta > 0$. The likelihood ratio is a non-decreasing function of n for all $\theta_1 > \theta_0$:

$$L_{\theta_1}(n) = \frac{p_{\theta_1}(n)}{p_{\theta_0}(n)} = \left(\frac{\theta_1}{\theta_0}\right)^n e^{-(\theta_1 - \theta_0)}.$$

Thus the conditions for the Karlin-Rubin theorem hold, and thus it follows that the UMP test of size α for testing $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$ is a threshold test of the form

$$\phi(N) = \begin{cases} 1, & \text{for } N > n_0, \\ \gamma, & \text{for } N = n_0, \\ 0, & \text{for } N < n_0, \end{cases}$$

where n_0 and γ are selected to yield a size alpha test.

10.20

The Generalized Likelihood Ratio Test

- Sometimes we can't find a UMP test
  - Doesn't exist
  - Difficult to construct
- We then need an alternative:
  - Locally Most Powerful (LMP) tests
  - Generalized Likelihood ratio test (GLRT)

10.21

The Generalized Likelihood Ratio Test

The *generalized likelihood ratio test* (GLRT) gets around the composite hypothesis testing problem by effectively turning it into a test between two simple hypotheses.

These simple hypotheses are selected to be the most likely value of

$$\theta_0 \in \Theta_0 \text{ under } H_0$$

and

$$\theta_1 \in \Theta_1 \text{ under } H_1$$

given the observed data.

The Generalized Likelihood Ratio Test (Cont.)

- Using these two simple hypotheses, a likelihood ratio test is implemented to test between them.
- The composite hypothesis corresponding to the simple hypothesis declared by the simple likelihood test is the composite hypothesis declared by the GLRT.
- While the GLRT is not optimal in any particular sense, it seems like a reasonable approach to dealing with the composite hypothesis testing problem.
- In many cases where a UMP test does exist, the GLRT exhibits nearly optimal behavior.

Generalized Likelihood Ratio Test (GLRT)

Consider two composite hypotheses $H_0 : \underline{\theta} \in \Theta_0$ and $H_1 : \underline{\theta} \in \Theta_1$. The *Generalized Likelihood Ratio Test* (GLRT) consists of the following procedure:

1. Assume H_0 is true and estimate the value of θ from the observed data using a *maximum likelihood estimate* (MLE):
$$\hat{\underline{\theta}}_0 = \arg\max_{\underline{\theta} \in \Theta_0} f_{\underline{\theta}}(\underline{X}).$$
2. Assume H_1 is true and estimate the value of θ from the observed data using a (MLE):
$$\hat{\underline{\theta}}_1 = \arg\max_{\underline{\theta} \in \Theta_1} f_{\underline{\theta}}(\underline{X}).$$
3. Replace the original problem of testing the composite hypotheses H_0 versus H_1 with the problem of testing the simple hypotheses $\hat{H}_0 : \hat{\Theta}_0 = \{\hat{\theta}_0\}$ versus $\hat{H}_1 : \hat{\Theta}_1 = \{\hat{\theta}_1\}$. If $\hat{H}_0$ is decided in the simple hypothesis problem, then H_0 is decided as the composite hypothesis. If $\hat{H}_1$ is decided in the simple hypothesis problem, then H_1 is decided as the composite hypothesis.

10.24

GLRT-continued

When we carry out this procedure, we get a GLRT of the form

$$L_g(\underline{X}) = \frac{\max\limits_{\underline{\theta}\in\Theta_1} f_{\underline{\theta}}(\underline{X})}{\max\limits_{\underline{\theta}\in\Theta_0} f_{\underline{\theta}}(\underline{X})} \underset{H_0}{\overset{H_1}{\gtrless}} L_0,$$

which yields a statistical test of the form

$$\phi(\underline{X}) = \begin{cases} 1, & \text{for } L_g(\underline{X}) > L_0, \\ \gamma, & \text{for } L_g(\underline{X}) = L_0, \\ 0, & \text{for } L_g(\underline{X}) < L_0, \end{cases}$$

where, in principle, L_0 and γ are selected to yield a size α test.

10.25

L_0 and γ

$$\boxed{\phi(\underline{X}) = \begin{cases} 1, & \text{for } L_g(\underline{X}) > L_0, \\ \gamma, & \text{for } L_g(\underline{X}) = L_0, \\ 0, & \text{for } L_g(\underline{X}) < L_0. \end{cases}}$$

Selecting L_0 and γ to yield a size α test is difficult.

The reason is that the size of the test is still defined as

$$\alpha = \sup_{\underline{\theta}\in\Theta_0} \mathrm{E}_{\underline{\theta}}\left[\phi(\underline{X})\right].$$

We do not use $\mathrm{E}_{\underline{\hat{\theta}}_0}\left[\phi(\underline{X})\right]$ as the size of the test.

Example: *Suppose we wish to test the hypotheses H_0 versus H_1 that the random sample $\underline{X} = (X_1, \ldots, X_N)$ comes from a density*

$$f_\theta(x_n) = \frac{1}{\sqrt{2\pi}} \exp\left\{ \frac{-(x_n - \theta)^2}{2} \right\},$$

where under $H_0 : \Theta_0 = [-1, 1]$, and under $H_1 : \Theta_0 = \{\theta \in \mathbf{R} : |\theta| > 1\}$.

We note that in general,

$$f_\theta(\underline{X}) = \frac{1}{(2\pi)^{N/2}} \exp\left\{ -\frac{1}{2} \sum_{n=1}^{N} (x_n - \theta)^2 \right\}, \tag{1}$$

from which it follows that the (unconstrained) maximum likelihood estimate $\hat{\theta}_{ML}$ of θ can be found by solving

$$\frac{\partial}{\partial \theta} f_\theta(\underline{X}) = 0$$

for θ, yielding

$$\hat{\theta}_{ML} = \frac{1}{N} \sum_{n=1}^{N} X_n.$$

For this case, we can easily see that

$$\hat{\theta}_0 = \arg \max_{\underline{\theta} \in \Theta_0} f_{\underline{\theta}}(\underline{X}) = \begin{cases} \hat{\theta}_{ML}, & \text{for } \hat{\theta}_{ML} \in [-1, 1], \\ -1, & \text{for } \hat{\theta}_{ML} < -1, \\ 1, & \text{for } \hat{\theta}_{ML} > 1, \end{cases}$$

and

$$\hat{\theta}_1 = \arg \max_{\underline{\theta} \in \Theta_1} f_{\underline{\theta}}(\underline{X}) = \begin{cases} \hat{\theta}_{ML}, & \text{for } \hat{\theta}_{ML} \notin [-1, 1], \\ -1, & \text{for } \hat{\theta}_{ML} \in [-1, 0], \\ 1, & \text{for } \hat{\theta}_{ML} \in (0, 1]. \end{cases}$$

Using this $\hat{\theta}_0$ and $\hat{\theta}_1$ we can now construct the GLRT

$$L_g(\underline{X}) = \frac{f_{\hat{\theta}_1}(\underline{X})}{f_{\hat{\theta}_0}(\underline{X})} \underset{H_0}{\overset{H_1}{\gtrless}} L_0$$

with an appropriately chosen threshold L_0.

Recall: $\alpha = \sup_{\underline{\theta} \in \Theta_0} \mathrm{E}_{\underline{\theta}} [\phi(\underline{X})]; \quad \alpha \neq \mathrm{E}_{\hat{\underline{\theta}}_0} [\phi(\underline{X})].$