

ECE 7251: Signal Detection and Estimation

Spring 2002

Prof. Aaron Lanterman
Georgia Institute of Technology

Lecture 29, 3/22/02:
Generalized Likelihood Ratio Tests
and Model Order Selection Criteria

The Setup

- Usual parametric data model $p(y; \mathbf{q})$
- In previous lecture on LMP tests, we assumed special structures like:

$$H_0 : \mathbf{q} = \mathbf{q}_0, H_1 : \mathbf{q} > \mathbf{q}_0$$

$$\text{or } H_0 : \mathbf{q} = \mathbf{q}_0, H_1 : \mathbf{q} \neq \mathbf{q}_0$$

- What to do if we have a more general structure like:

$$H_0 : \mathbf{q} \in \mathcal{S}_0, H_1 : \mathbf{q} \in \mathcal{S}_1$$

- Often, we do something a bit ad hoc!

The GLRT

- Find parameter estimates $\hat{\mathbf{q}}_0$ and $\hat{\mathbf{q}}_1$ under H_0 and H_1
- Substituting estimates into likelihood ratio yields a generalized loglikelihood ratio test:

$$\Lambda_{GLR}(y) = \frac{p(y; \hat{\mathbf{q}}_1)_{H_1}}{p(y; \hat{\mathbf{q}}_0)_{H_0}} \geq 1$$

- If convenient, use ML estimates:

$$\frac{\max_{\mathbf{q} \in \mathcal{S}_1} p(y; \mathbf{q})_{H_1}}{\max_{\mathbf{q} \in \mathcal{S}_0} p(y; \mathbf{q})_{H_0}} \geq 1$$

Two-Sided Gaussian Mean Example

$$y_i \in \mathcal{N}(\mathbf{q}, \mathbf{s}^2), H_0 : \mathbf{q} = 0, H_1 : \mathbf{q} \neq 0$$

$$\begin{aligned} \ln \frac{p(y; \hat{\mathbf{q}})}{p(y; 0)} &= - \sum_{i=1}^n \frac{\left(y_i - \frac{1}{n} \sum_{j=1}^n y_j \right)^2}{2\mathbf{s}^2} + \sum_{i=1}^n \frac{y_i^2}{2\mathbf{s}^2} \\ &= \frac{2 \sum_{i=1}^n y_i \frac{1}{n} \sum_{j=1}^n y_j - n \left(\frac{1}{n} \sum_{i=1}^n y_i \right)^2}{2\mathbf{s}^2} \end{aligned}$$


Two-Sided Gaussian Example Con't

$$\frac{2n \left(\frac{1}{n} \sum_{i=1}^n y_i \right)^2 - n \left(\frac{1}{n} \sum_{i=1}^n y_i \right)^2}{2\mathbf{s}^2} = \frac{n}{2\mathbf{s}^2} \bar{y}^2 \geq 1 \begin{matrix} H_1 \\ H_0 \end{matrix}$$

Same as the LMPU test from last lecture! $\begin{cases} H_1 \\ H_0 \end{cases} \left\{ \begin{array}{l} | \bar{y} | \geq \mathbf{g} \end{array} \right.$

- Chapter 9 of Hero derives and analyzes the GLRT for every conceivable Gaussian problem – a fantastic reference!

Gaussian Performance Comparison



We take a performance hit from not knowing the true mean

(Graph from p. 95 of Van Trees Vol. 1)

Some Gaussian Examples

- Single population:
 - Tests on mean, with unknown variance yield “T tests”
 - Statistic has a Student T distribution
 - Asymptotically Gaussian
- Two populations:
 - Tests on equality of variances, with unknown means yields a “Fisher F test”
 - Statistic has a Fisher F distribution
 - Asymptotically Chi Square
- See Chapter 9 of Hero

Asymptotics to the Rescue

- Suppose $n \rightarrow \infty$. Since the ML estimates are asymptotically consistent, the GLRT is asymptotically UMP
- If the GLRT is hard to analyze directly, sometimes asymptotic results can help
- Assume a partition

$$\mathbf{q} = (\mathbf{j}_1, \dots, \mathbf{j}_p, \underbrace{\mathbf{x}_1, \dots, \mathbf{x}_q}_{\text{(nuisance parameters)}})$$

Asymptotics Con't

- Consider GLRT for a two sided problem
 $H_0: \mathbf{j} = \mathbf{j}_0, H_1: \mathbf{j} \neq \mathbf{j}_0$
 where \mathbf{x} is unknown, but we don't care what it is
- When the density $p(y; \mathbf{q})$ is smooth under H_0 , it can be shown that for large n

$$2 \ln \Lambda_{GLR}(Y) = 2 \ln \frac{p(Y; \hat{\mathbf{q}})}{p(Y; \mathbf{q}_0)} \sim \mathbf{c}_p \left. \begin{array}{l} \text{(Chi-square} \\ \text{with } p \\ \text{degrees of} \\ \text{freedom)} \end{array} \right\}$$
- Recall $E[\mathbf{c}_p] = p, \text{ var}(\mathbf{c}_p) = 2p$

A Strange Link to Bayesianland

- Remember if we had a prior $p(\mathbf{q})$, we could handle composite hypothesis tests by integrating and reducing things to a simple hypothesis test

$$p(y) = \int_{R^p} p(y | \mathbf{q}) p(\mathbf{q}) d\mathbf{q}$$

- If $p(\mathbf{q})$ varies slowly compared to $p(y | \mathbf{q})$ around the MAP estimate, we can approx.

$$p(y) \approx p(\hat{\mathbf{q}}) \int_{R^p} \exp[L(\mathbf{q})] d\mathbf{q}$$

Laplace's Approximation

- Do a Taylor series expansion

$$\int_{R^p} \exp\left[L(\hat{\mathbf{q}}_{ML}) - \frac{(\mathbf{q} - \hat{\mathbf{q}}_{ML})^T F(y; \mathbf{q})(\mathbf{q} - \hat{\mathbf{q}}_{ML})^T}{2} \right] d\mathbf{q}$$

$$= e^{L(\hat{\mathbf{q}}_{ML})} \int_{R^p} \exp\left[-\frac{(\mathbf{q} - \hat{\mathbf{q}}_{ML})^T F(y; \hat{\mathbf{q}}_{ML})(\mathbf{q} - \hat{\mathbf{q}}_{ML})^T}{2} \right] d\mathbf{q}$$

where $F(y; \hat{\mathbf{q}}_{ML}) = \left[-\frac{d^2 L(\mathbf{q})}{d\mathbf{q}_r d\mathbf{q}_c} \Big|_{\mathbf{q}=\hat{\mathbf{q}}_{ML}} \right] \left. \begin{array}{l} \text{Empirical} \\ \text{Fisher info} \end{array} \right\}$

Laplace's Approximation Con't

- Recognize quadratic form of the Gaussian:

$$\int_{R^p} \exp\left[-\frac{(\mathbf{q} - \hat{\mathbf{q}}_{ML})^T F(y; \hat{\mathbf{q}}_{ML})(\mathbf{q} - \hat{\mathbf{q}}_{ML})^T}{2} \right] d\mathbf{q}$$

$$= \frac{(2\pi)^{p/2}}{\sqrt{\det F(y; \hat{\mathbf{q}}_{ML})}}$$

- So $p(y) = p(\hat{\mathbf{q}}_{ML}) p(y | \hat{\mathbf{q}}_{ML}) \frac{(2\pi)^{p/2}}{\sqrt{\det F(y; \hat{\mathbf{q}}_{ML})}}$

Large Sample Sizes

- Consider the logdensity:

$$\ln p(y) \approx \ln p(\hat{\mathbf{q}}_{ML}) + \ln p(y|\hat{\mathbf{q}}_{ML}) + \frac{p}{2} \ln 2\mathbf{p} - \frac{1}{2} \ln \det F(y; \hat{\mathbf{q}}_{ML})$$
- Suppose we have n i.i.d. samples. By the law of large numbers:

$$\begin{aligned} \ln \det F(y|\hat{\mathbf{q}}_{ML}) &\approx \ln \det F(\hat{\mathbf{q}}_{ML}) = \ln \det nF_1(\hat{\mathbf{q}}_{ML}) \\ &= \ln \det [nI \cdot F_1(\hat{\mathbf{q}}_{ML})] = \ln \det [nI] + \ln \det F_1(\hat{\mathbf{q}}_{ML}) \\ &= \ln n^p + \ln \det F_1(\hat{\mathbf{q}}_{ML}) = p \ln n + \ln \det F_1(\hat{\mathbf{q}}_{ML}) \end{aligned}$$

Schwarz's Result

- As n gets big

$$\begin{aligned} \ln p(y) &\approx \ln p(\hat{\mathbf{q}}_{ML}) + L(\hat{\mathbf{q}}_{ML}) \\ &\quad + \frac{p}{2} \ln 2\mathbf{p} - \frac{1}{2} p \ln n - \frac{1}{2} \ln \det F_1(\hat{\mathbf{q}}_{ML}) \\ &\approx L(\hat{\mathbf{q}}_{ML}) - \frac{p}{2} \ln n \end{aligned}$$
- Called Bayesian Information Criterion (BIC) or Schwarz Information Criterion (SIC)
- Often used in model selection; second term is a penalty on model complexity

Minimum Description Length

- BIC is related to Rissanen's Minimum Description Length criterion; $(p/2) \ln(n)$ is viewed as the optimum number of "nats" (like bits, but different base) used to encode the ML parameter estimate with limited precision
- Data is encoded with a string of length

$$\text{description length} = \underbrace{-L(\hat{\mathbf{q}}_{ML})}_{\substack{\text{nats used to encode} \\ \text{data given ML est.}}} + \frac{p}{2} \ln n$$
- Choose model which describes the data using the smallest number of bits

References

- A.R. Barron, J. Rissanen, B. Yu, The Minimum Description Length Principle in Coding and Modeling, *IEEE Trans. Info. Theory*, Vol. 44, No. 6, Oct. 1998, pp. 2743-2760.
- A.D. Lanterman, *Schwarz, Wallace, and Rissanen: Intertwining Themes in Theories of Model Order Estimation*, *International Statistical Review*, Vol. 69, No. 2, August 2001, pp. 185-212.
- Special Issues:
 - Statistics and Computing (Vol. 10, No. 1, 2000)
 - The Computer Journal (Vol. 42, No. 4, 1999)