STAT 703/J703
January 25th, 2007
-Lecture 4-

Instructor: Brian Habing
Department of Statistics
LeConte 203
Telephone: 803-777-3578
E-mail: habing@stat.sc.edu

8.5 – Maximum Likelihood

The idea behind maximum likelihood estimation is to find the parameters that seem most likely to have resulted in the observed statistic.

In the case of more than one observation, the likelihood is:

$\text{lik}(\theta|x_1,\ldots x_n)=f(x_1,\ldots x_n|\theta)=\Pi f(x_i|\theta)$

It is unpleasant to find the maximum of a product though…

An easier function to work with is the <u>log likelihood</u>

L(q)=log(lik(θ))

$=\log(\Pi f(x_i|\theta))$

$=\Sigma \log(f(x_i|\theta))$

Consider a sample from a Poisson distribution….

Unlike the Poisson distribution, a Gamma distribution has two parameters to deal with…

```
ngamloglike<-function(pars,data){
 a<-pars[1]
 l<-pars[2]
 n<-length(data)
 -1*
 (n*a*log(l)+(a1)*sum(log(data))-
  l*sum(data)-n*lgamma(a))
  }
optim(c(6,6),neggloglike,data=a)
```


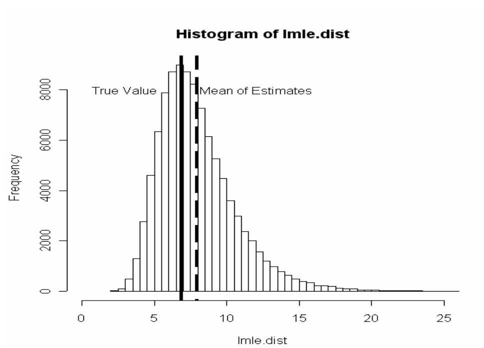

MLE Example 1 (Section 8.5.1)

Consider a multinomial experiment, where each of the *n* observations has probability p_i of falling in cell i=1…m.

$$f(x_1,\ldots x_m \mid p_1\ldots,p_m)=\frac{n!}{\prod_{i=1}^{m} x_i!}\prod_{i=1}^{m} p_i^{x_i}$$

Often there is some restriction on the probabilities however. Consider the case in genetics where a gene occurs with frequence θ in the population.

Say we observe 342 ++, 500 +- and 187 --. What is our estimate of θ?

MLE example 2 – Logistic Regression

In linear regression we try to predict one continuous variable from another.

Under a few basic assumptions this results in fairly easy parameter estimation and use of t and F distributions.

Consider the case of attempting to predict a 0,1 variable from a continuous variable.

$$P[Y_i = 1 | x_i] = \frac{1}{1 + e^{-(\alpha + \beta x_i)}}$$