Testing for a Signal

Dr Wolfgang Rolke

Dr Angel Lopez

Likelihood Ratio Test

- One of the most commonly used methods in Statistics
- Previously used by us to find limits in the presence of nuisance parameters
- For the problem of setting limits already known in HEP as the method of MINOS errors in MINUIT

Basic Idea

Let's consider the following general problem: we have data X from a distribution with density $f(x; \theta)$ where θ is a vector of parameters with $\theta \in \Theta$ and Θ is the entire parameter space. We wish to test

$$H_0 : \theta \in \Theta_0 \text{ vs. } Ha : \theta \in \Theta^c{}_0$$

where Θ_0 is some subset of Θ.

Now the likelihood function is defined by

$$L(\theta|x) = f(x; \theta)$$

and the likelihood ratio test statistic is defined by

$$\lambda(x) = \sup_{\Theta 0} L(\theta|x)/\sup_{\Theta} L(\theta|x)$$

Now small values of $\lambda(x)$ lead to us to reject the null hypothesis because they indicate that the data can be much better explained under the alternative hypothesis, and we have a test of the form

$$\text{Reject } H_0 \text{ if } \lambda(x)<c$$

where c depends on the desired type I error probability α

How to find c for a given α? A well-known theorem says that, if there is enough data, $L(x)=-2\log \lambda(x)$ has a chi-square distribution

Improving the Test

In HEP we compare the number of events x in the signal region with the number of events y in the sidebands. If there is a large excess of signal events we claim a discovery. Probabilistically, though, there will also occasionally be fewer events in the signal region than are expected from the sidebands, which also leads to a large L(x,y), but which of course does not indicate a true signal. The test therefore is

$$\text{Reject } H_0 \text{ if } x>y \text{ and } L(x,y)>c$$

A simple Example

- $X \sim P(\mu+b)$
- $Y \sim P(\tau b)$

Then

$$L(x,y)=2(x\log x+y\log y-(x+y)\log((x+y)/(1+\tau))-y\log\tau)$$

And we reject the null hypothesis if

$x>y/\tau$ and $L(x,y)>q\chi^2(1-2\alpha)$

Extensions: Multiple channels

The test can be extended to multiple channels in two ways:

(1) FullLrt: Declare a signal if there is some excess in many channels

null distribution is a linear combination of chi-squares

(2) MxLrt: Declare a signal if there is a strong excess in any one channel

null distribution via extreme value theory

Both have strengths and weaknesses and we will consider both

Extensions: Marked Poisson

- Often we have additional information on the events such as their mass. Then we can include this in the likelihood as well

→ marked Poisson as shown by Kyle Cranmer

Questions

- How much data is enough for the chi-square approximation to work?
- Discoveries in HEP are made at the 5σ level – does it work this far out in the tails?
- Can we use multiple channels?
- How about marked Poisson?

- $X \sim \text{Pois}(e\mu + b)$
- $Y \sim \text{Pois}(\tau\, b)$
- $b = 5$ to $b = 150$
- $\tau = 1$
- $\sigma = 3, 4$ and 5
- Exact calculation

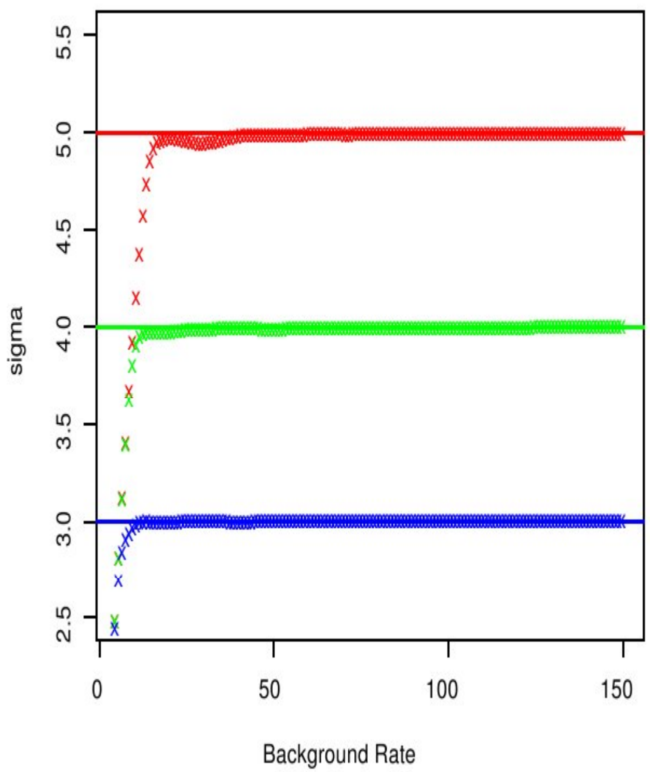

Fig. 1.

- $b = 100$
- $\tau = 1$
- μ from 0 to 150
- $\alpha = 3\sigma$, 4σ and 5σ
- Exact calculation

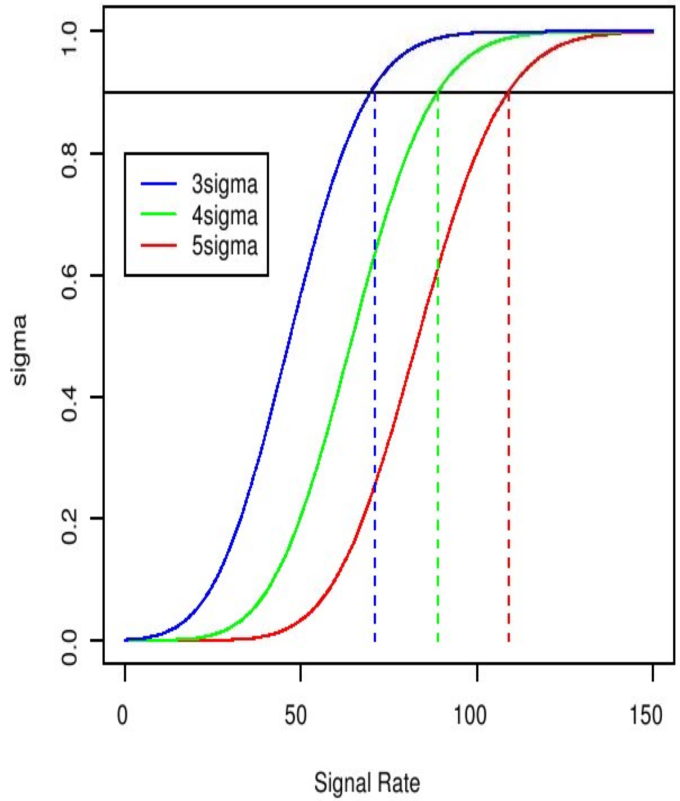

Fig. 2.

- 2 channels
- No signals
- $b_1 = b_2 = 10 : 100$
- $\tau_1 = \tau_2 = 1$
- Exact calculation

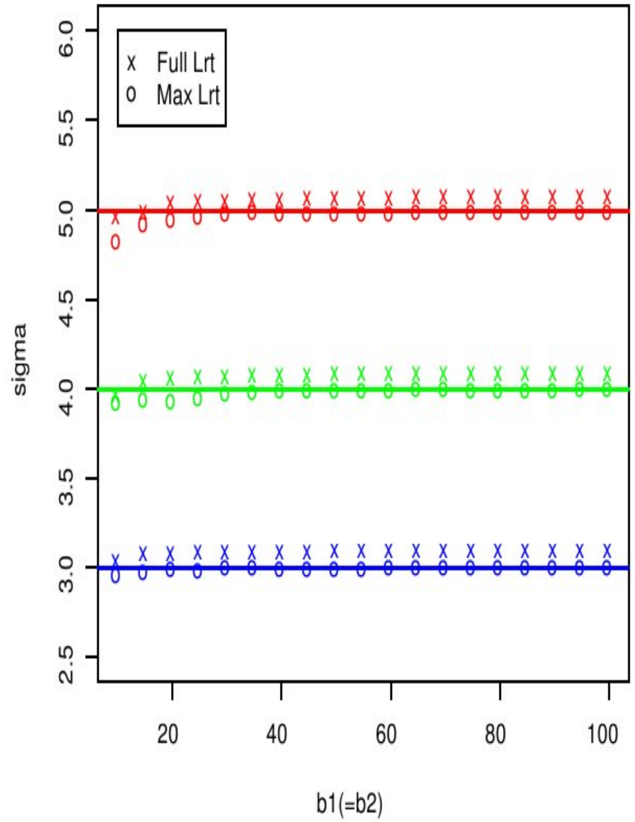

Fig. 4.

- $b_1 = .. = b_k = 50$
- all $\tau_i = 1$
- $k = 1, .., 10$
- Mini MC

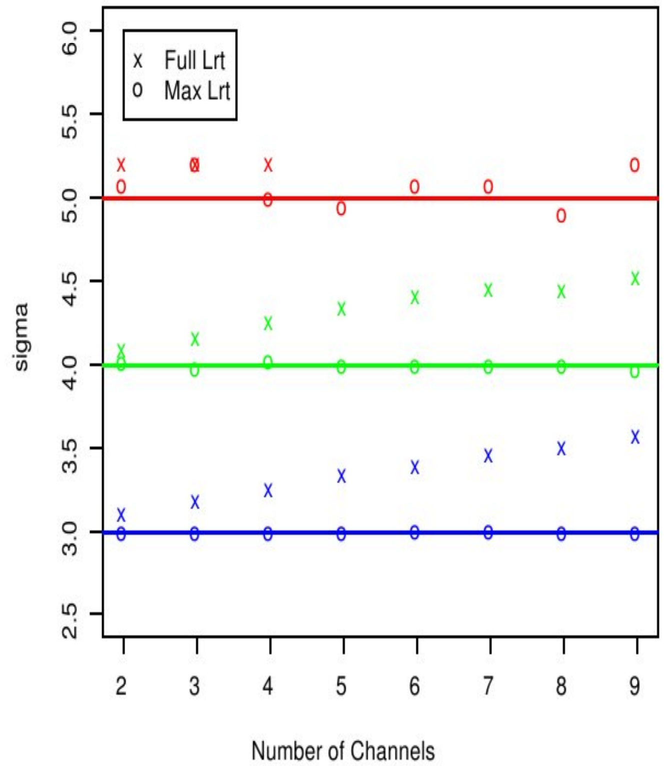

Fig. 5.

- 10 channels
- $b_1 = .. = b_{10} = 50$
- all $\tau_i = 1$
- $\mu_1 = .. = \mu_{10} = \mu$ going from 0 to 50.
- Mini MC

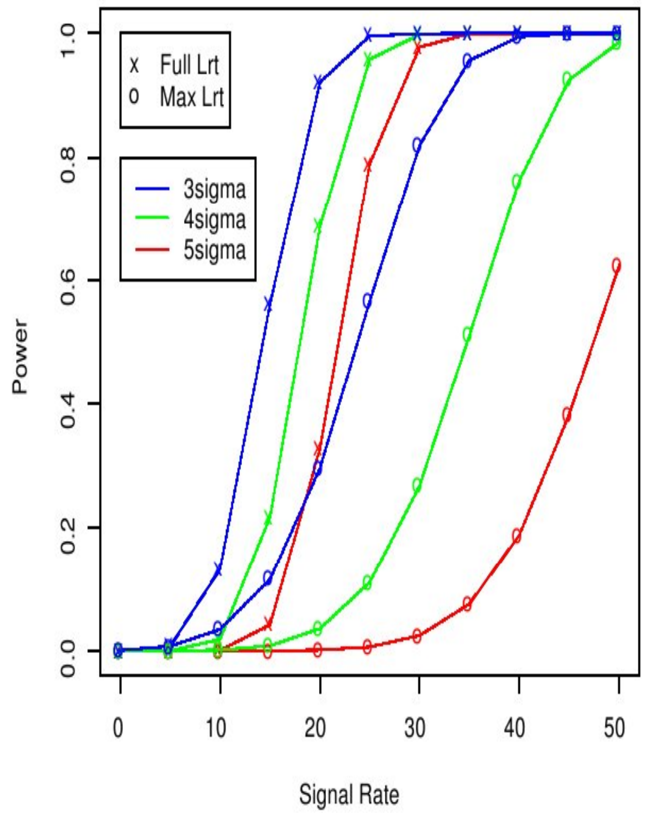

Fig. 6.

- 10 channels
- $b_1 = .. = b_{10} = 50$
- all $\tau_i = 1$
- $\mu_2 = .. = \mu_{10} = 0$
- μ_1 going from 0 to 100.
- Mini MC

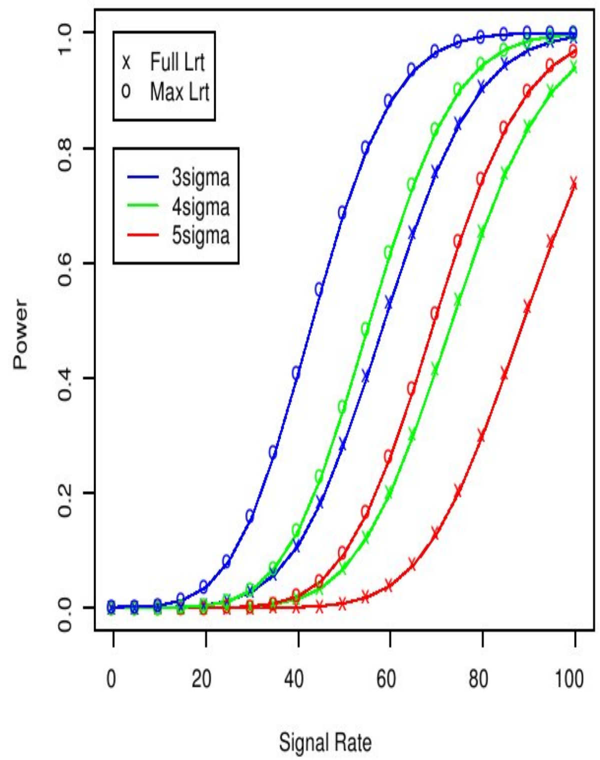

Fig. 7.

Summary

- Test works

  (i.e., it achieves nominal type I error probability, even at 5σ)
- Extends to multiple channels
- Extends to marked Poisson
- Calculation is very fast
- Implementation of different models is easy (using MINUIT)
- For a detailed discussion check out

  physics/0606006