Detection

Sensor Fusion

Fredrik Gustafsson
fredrik.gustafsson@liu.se

Gustaf Hendeby
gustaf.hendeby@liu.se

Linköping University

The Detection Problem

Basic question for the linear or nonlinear model:

- H_1: Is there any term Hx or $h_k(x)$ at all?
- H_0: Or is the sensor just measuring noise?

This is a statistical hypothesis test.

We will start with simplest case, and then increase complexity.

Simplest Test: No alternate model

Formulate the hypotheses:

$$H_0 : \mathbf{y} = \mathbf{e},$$
$$H_1 : \mathbf{y} \neq \mathbf{e},$$

where $\mathbf{e} \in \mathcal{N}(0, \mathbf{R})$. That is, H_1 does not specify what the measurement is for the alternate hypothesis.
Define the test statistic

$$T(\mathbf{y}) = \mathbf{y}^T \mathbf{R}^{-1} \mathbf{y} \sim \chi^2_{\dim(\mathbf{y})},$$

which from the Gaussian distribution is χ^2 distributed if H_0 is true.
We can then pick a threshold h from the χ^2 distribution and use the test

$$T(\mathbf{y}) \underset{H_0}{\overset{H_1}{\gtrless}} h$$

Test with alternate model

Suppose now we have a completely specified alternate model

$$H_0 : \mathbf{y} = \mathbf{e},$$
$$H_1 : \mathbf{y} = \mathbf{H}x^0 + \mathbf{e},$$

where $\mathbf{e} \in \mathcal{N}(0, \mathbf{R})$. When both hypotheses have a well-known distribution, the (two times) log likelihood ratio (LR) should be used as test statistic:

$$T(\mathbf{y}) = 2\log\frac{p(\mathbf{y}|H_1)}{p(\mathbf{y}|H_0)} = 2\log\frac{p_{\mathbf{e}}\big(\mathbf{y} - \mathbf{H}\hat{x}^0\big)}{p_{\mathbf{e}}\big(\mathbf{y}\big)} = \cdots = \big(\mathbf{y} - \mathbf{H}\hat{x}^0\big)^T \mathbf{R}^{-1}\big(\mathbf{y} - \mathbf{H}\hat{x}^0\big)$$

The distribution for $T(\mathbf{y})$ under H_0 and H_1, respectively, can be computed and used for theoretical evaluation.

Optimality

Neyman-Pearson's Lemma

The likelihood ratio is the optimal test statistic for two hypotheses with well-specified distributions.

- Optimal here means that it gives the best power (most detections) of all tests with the same false alarm rate.
- The threshold h decides the compromise between power and false alarms.
- A ROC plot shows the power versus false alarm rate when h is varied. (ROC stands for Receiver Operating Characteristics for historical reasons)

Test with alternate linear model (1/2)

Consider now the more realistic case where we do not know the true parameter x^0 in the model,

$$H_0 : \mathbf{y} = \mathbf{e},$$
$$H_1(x) : \mathbf{y} = \mathbf{H}x + \mathbf{e},$$

where $\mathbf{e} \in \mathcal{N}(0, \mathbf{R})$. We can always try to estimate it and use it in the log likelihood ratio. In the general case, we maximize the LR wrt x, which is the same as plugging in the ML estimate in the LR, and get the Generalized Likelihood Ratio (GLR) test

$$T(\mathbf{y}) = \log \frac{\max_x p(\mathbf{y}|H_1(x))}{p(\mathbf{y}|H_0)} = 2 \log \frac{p_{\mathbf{e}}(\mathbf{y} - \mathbf{H}\hat{x}^{ML})}{p_{\mathbf{e}}(\mathbf{y})} = \dots$$
$$= \mathbf{y}^T \mathbf{R}^{-T/2} \underbrace{\left(\mathbf{R}^{-T/2}\mathbf{H}(\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})^{-1}\mathbf{H}^T\mathbf{R}^{-1/2}\right)}_{\Pi} \mathbf{R}^{-1/2}\mathbf{y}$$

Test with alternate linear model (2/2)

One can show that the test statistic is χ^2 distributed

$$T(\mathbf{y}) \sim \begin{cases} \chi^2_{\mathrm{dim}(x)} & \text{under } H_0, \\ \chi^2_{\mathrm{dim}(x)}\big((x^o)^T \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} x^0\big) & \text{under } H_1, \end{cases}$$

Threshold selection:

- Chose h from the $\chi^2_{\mathrm{dim}(x)}$ distribution to get a pre-selected false alarm rate.
- Compute the power from the non-central χ^2 distribution.

Note:

Why not always apply the simplest possible test without any alternate model? It always applies? Because the GLR test gives much better power!

GLR Optimality

Recall that the ML estimate is efficient, so asymptotically it is unbiased and attains the CRLB lower bound. That is, asymptotically ML cannot be beaten.

GLRT Optimality

The GLR test is the uniformly most powerful one, meaning that asymptotically it gives the best power of all tests

First example, revisited

Two sensors with good angle but poor range resolution, as studied in Chapter 2. Detect from a measurement pair if a target is present. Key quantities for the test:

$$y_1 = x + e_1, \quad \mathrm{Cov}(e_1) = R_1$$
$$y_2 = x + e_2, \quad \mathrm{Cov}(e_2) = R_2$$

$$\mathbf{y} = \mathbf{H}x + \mathbf{e}, \quad \mathrm{Cov}(\mathbf{e}) = \mathbf{R} \quad \mathbf{H} = \begin{pmatrix} I \\ I \end{pmatrix}$$

$$\mathbf{R} = \begin{pmatrix} R_1 & 0 \\ 0 & R_2 \end{pmatrix}$$

$$T(\mathbf{y}) = \mathbf{y}^T \mathbf{R}^{-T/2} \Pi \mathbf{R}^{-1/2} \mathbf{y},$$

$$\Pi = \mathbf{R}^{-T/2}\mathbf{H}(\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})^{-1}\mathbf{H}\mathbf{R}^{-1/2} = \begin{bmatrix} 22.8 & 22.2 & 5.0 & 0.0 \\ 22.2 & 22.8 & 0.0 & 5.0 \\ 5.0 & 0.0 & 22.8 & -22.2 \\ -0.0 & 5.0 & -22.2 & 22.8 \end{bmatrix}$$

Numerical simulation

Threshold for the test:

```
h=erfinv(chi2dist(2),0.999)
h =
10.1405
```

Note: for no model (standard statistical test), $\Pi = I$. Both methods perform perfect $P_D = 1$.

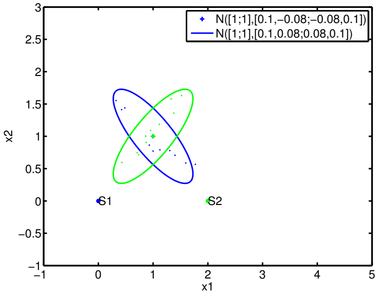

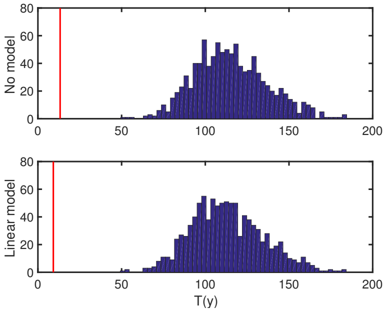

Sensor network example 1/2

It is quite easy to define a sensor network in the Signal and System Lab. The following code sets up a random configuration with 5 TOA sensors, makes an illustration and simulates data.

```
ny=5; nx=2;
s=exsensor('toa',ny,1);      % Default network
s.pe=0.2*eye(ny);            % Set noise level
plot(s)
y=simulate(s);
T=y.y*inv(cov(s.pe))*y.y'; % Test statistic
```

We can analyse the power from the two tests with no alternate model where $T \sim \chi^2_4$ and a linear alternate model where $T \sim \chi^2_2$

```
[b,level]=detect(chi2dist(ny),y.y*inv(cov(s.pe))*y.y')
level =
    0.9626
[b,level]=detect(s,y)
level =
    0.9985
```

Note that the different degrees of freedom affects the power!

Sensor network example

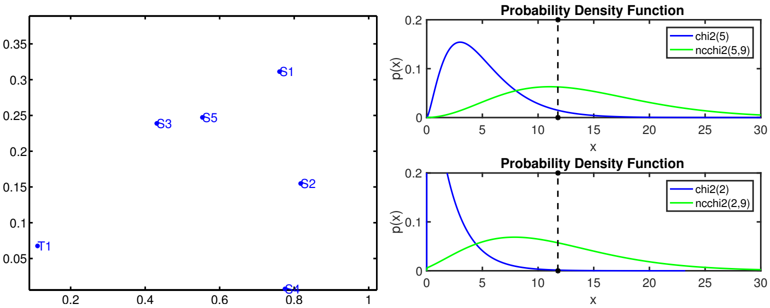

Note:

The model based approach decreases the degrees of freedom in the χ^2 distribution under H_0 from $\dim(\mathbf{y}) = 4$ to $\dim(x) = 2$. The threshold can thus be significantly decreased for the same false alarm rate!

Summary

Two fundamentally different approaches for detecting if a model is present or not:

1. Test if $\mathbf{y} = \mathbf{e}$ or if it deviates from the assumed noise distribution. The test statistic is $T(\mathbf{y}) = \mathbf{y}^T \mathbf{R}^{-1} \mathbf{y} \sim \chi^2_{\dim(\mathbf{y})}$
2. Test if $\mathbf{y} = \mathbf{e}$ or if $\mathbf{y} = \mathbf{H}x + \mathbf{e}$. The generalized likelihood ratio (GLR) test statistic is $T(\mathbf{y}) = \mathbf{y}^T \mathbf{R}^{-1/2} \Pi \mathbf{R}^{-1/2} \mathbf{y} \sim \chi^2_{\dim(x)}$

The latter gives a much smaller threshold and better test (higher power).

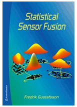

Chapter 5