Project #1: Sequential Detection

EEE 556
Spring 2002

Obligatory Statement on Intellectual Honesty

Your project report should represent your individual work. You may confer with class-mates about the project, but the derivations, programs, and writeup in the report that you turn in must be your own.

1 Introduction

In this project, you will become familiar with basic sequential detection methods and investigate their advantages and disadvantages relative to fixed length detection methods. Example 3.2 of *Fundamentals of Statistical Signal Processing: Volume II Detection Theory* shows that the performance of a detector can be improved by increasing the signal-to-noise ratio. Often, the noise power is fixed, and the detector performance can only be improved by increasing the signal energy $N\mu^2$, where μ is the amplitude of a transmitted signal. Power constraints frequently rule out the possibility of increasing μ; thus, we must increase the number of samples N. Increasing N typically means an increase in the time and resources necessary to reach a decision; in most applications, it is desireable to minimize the number of samples needed to make a decision.

One way to reduce the *average* number of samples necessary to make a decision is to use a sequential decision rule. In a sequential decision rule, samples are processed sequentially (one after another) to form a decision statistic which is compared to both an upper and a lower threshold. When the statistic crosses either threshold, processing stops and a decision is declared. Thus, the number of samples needed to reach a decision depends on the values of the samples. Usually, the average number of samples needed to obtain a desired level of performance is fewer than the number of samples necessary for a fixed sample length detector.

In this project, you will derive sequential decision rules for several problems. You will then evaluate these rules using Monte Carlo computer simulation, characterizing their performance in terms of designed and actual probabilities of detection and false alarm and in terms of expected number of samples needed to reach a decision.

The problems in Parts II and III of this project are based on real-life problems but are not very realistic. We need to know quite a bit more about detecting signals with unknown parameters before we can work more realistic problems; subsequent projects should have more realistic problems.

2 Theory of Sequential Detection

We consider a detection problem between $\mathcal{H}_0$ and $\mathcal{H}_1$. We assume that the observations $x[0], x[1], \ldots$ and that the likelihood functions of the observations under the two hypotheses are $p\left(x[0], x[1], \ldots, x[N-1]; \mathcal{H}_0\right)$ and $p\left(x[0], x[1], \ldots, x[N-1]; \mathcal{H}_1\right)$. A sequential test requires two thresholds A and B; we show below how these thresholds are computed to achieve desired values of P_D and P_{FA}. Given values of A and B, the sequential decision algorithm begins with $N = 1$ and a single observation $x[0]$, and makes decisions as follows:

- Choose $\mathcal{H}_0$ if

$$L\left(x[0], \ldots, x[N-1]\right) = \frac{p\left(x[0], \ldots, x[N-1]; \mathcal{H}_1\right)}{p\left(x[0], \ldots, x[N-1]; \mathcal{H}_0\right)} \leq B. \tag{1}$$

- Choose $\mathcal{H}_1$ if

$$L\left(x[0], \ldots, x[N-1]\right) = \frac{p\left(x[0], \ldots, x[N-1]; \mathcal{H}_1\right)}{p\left(x[0], \ldots, x[N-1]; \mathcal{H}_0\right)} \geq A. \tag{2}$$

- Otherwise, get another sample, increase N by one, and repeat the test.

Unlike the Neyman-Pearson test, with a sequential test you can specify both P_D and P_{FA} for a detector; the values of P_D and P_{FA} determine the average number of samples that will be needed to reach a decision.

We now derive (approximate) expressions for A and B in terms of the desired values P_D and P_{FA}. To obtain A, we consider the situation where the test terminates with N samples having chosen $\mathcal{H}_1$. In this case, Equation (2) holds. We assume that the amount by which $L\left(x[0], \ldots, x[N-1]\right)$ exceeds A in this test is negligibly small; this assumption is typically valid when a large number of samples is required to reach a decision. Under this assumption,

$$p\left(x[0], \ldots, x[N-1]; \mathcal{H}_1\right) = A p\left(x[0], \ldots, x[N-1]; \mathcal{H}_0\right)$$

Integrating both sides of this equality over R_1, we obtain

$$P_D = A P_{FA},$$

from which we can solve for A:

$$A = \frac{P_D}{P_{FA}}.$$

Similarly, to obtain B, we consider the situation where the test terminates with N samples and $\mathcal{H}_0$ is chosen. In this case, Equation (1) holds. We assume that the amount by which $L\left(x[0], \ldots, x[N-1]\right)$ is smaller than B is negligibly small:

$$p\left(x[0], \ldots, x[N-1]; \mathcal{H}_1\right) = Bp\left(x[0], \ldots, x[N-1]; \mathcal{H}_0\right)$$

Integrating both sides of this equality over R_0, we obtain

$$1 - P_D = B\left(1 - P_{FA}\right),$$

from which we can solve for B:

$$B = \frac{1 - P_D}{1 - P_{FA}}.$$

3 Part I

In Part I, you will implement a very simple sequential test to become familiar with its operation and simulation. The two hypotheses are

$$\begin{array}{ll} \mathcal{H}_0: & x[n] = w[n], \qquad n = 0, \ldots, N-1 \\ \mathcal{H}_1: & x[n] = \mu + w[n], \quad n = 0, \ldots, N-1 \end{array}$$

where μ is a known positive constant and $w[n]$ is a Gaussian white noise sequence with mean zero and known variance σ^2.

Do the following for Part I:

1. Derive the Neyman Pearson likelihood ratio test to choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ using a fixed number of samples N for a given P_{FA}. Compute and plot P_D as a function of N. Use the values of P_{FA}, μ, and σ^2 from item 4 below.

2. Derive the sequential test to choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ (you should get a threshold test). Graph the upper and lower thresholds as a function of N, and explain why they take the form that they do. Express the statistic used in the threshold test in a form that can be computed recursively (ie. the statistic for N observations can be computed from the Nth observation and the statistic for $N-1$ observations.)

3. Implement the sequential test in a computer simulation using the language/package of your choice. Your simulation should generate a data sequence $\{x[0], \ldots, x[N-1]\}$ given either $\mathcal{H}_0$ or $\mathcal{H}_1$ is true and apply the sequential test to this data sequence.

4. Use your computer simulation to determine the performance of the sequential test. In particular, set $P_{FA} = 0.1$, $P_D = 0.9$, $\mu = 1$, and $\sigma^2 = 4$, and determine by simulation the average number of samples needed to decide $\mathcal{H}_0$ or $\mathcal{H}_1$, the actual probability of corect detection, and the actual probability of false alarm. How does the average number of samples needed by the sequential test to reach a decision compare to the number of samples needed by the fixed length decision rule to obtain the same P_D and P_{FA}? How do you explain this result?

4 Part II

Now we will consider detection of a sinusoidal signal with known amplitude, frequency, and phase using a sequential detector. This is motivated by problems in radar and active sonar. However, in real world problems, you would not know the amplitude, phase, and frequency of the signal. In later projects, we will address detection of signals with unknown amplitude and phase.

The detection problem is the following:

$$\begin{array}{ll} \mathcal{H}_0: \;\; x[n] = w[n], & n = 0, \ldots, N-1 \\ \mathcal{H}_1: \;\; x[n] = s[n] + w[n], & n = 0, \ldots, N-1 \end{array}$$

where $w[n]$ is a Gaussian white noise sequence with mean zero and known variance σ^2 and $s[n]$ is a sinusoidal signal:

$$s[n] = \mu \cos\left(\frac{2\pi}{8} n\right)$$

Do the following for Part II:

1. Derive the Neyman Pearson likelihood ratio test to choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ using a fixed number of samples N for a given P_{FA}.

2. Derive the sequential test to choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ (you should get a threshold test). Express the statistic used in the threshold test in a form that can be computed recursively.

3. Implement the sequential test in a computer simulation using the language/package of your choice.

4. Use your computer simulation to determine the performance of the sequential test. In particular, set $P_{FA} = 0.1$, $P_D = 0.9$, $\mu = 1$, and $\sigma^2 = 4$, and determine by simulation the average number of samples needed to decide $\mathcal{H}_0$ or $\mathcal{H}_1$, the actual probability of corect detection, and the actual probability of false alarm. How does the average number of samples needed by the sequential test to reach a decision compare to the

number of samples needed by the fixed length decision rule to obtain the same P_D and P_{FA}?

5. How do your results in Part II differ from your results in Part I? How do you explain these differences?

5 Part III

Consider the following problem:

$$\begin{aligned} \mathcal{H}_0: &\quad x[n] = w[n], & n = 0, \ldots, N-1 \\ \mathcal{H}_1: &\quad x[n] = s[n] + w[n], & n = 0, \ldots, N-1 \end{aligned}$$

where $w[n]$ is Gaussian white noise with zero mean and known variance σ_N^2 and the signal $s[n]$ is a white Gaussian sequence with zero mean and known variance σ_S^2. This signal model might apply to passive underwater detection of ships, in which the signal from a ship in a given frequency band is due to the interaction of the hull with the water.

Do the following for Part III:

1. Derive the Neyman Pearson likelihood ratio test to choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ using a fixed number of samples N for a given P_{FA}. Compute and plot P_D as a function of N. (You may assume that N is large in this computation.)

2. Derive the sequential test to choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ (you should get a threshold test). Express the statistic used in the threshold test in a form that can be computed recursively.

3. Implement the sequential test in a computer simulation.

4. Use your computer simulation to determine the performance of the sequential test. In particular, set $P_{FA} = 0.1$, $P_D = 0.9$, $\sigma_N^2 = 1$, and $\sigma_S^2 = 1$, and determine by simulation the average number of samples needed to decide $\mathcal{H}_0$ or $\mathcal{H}_1$, the actual probability of corect detection, and the actual probability of false alarm. How does the average number of samples needed by the sequential test to reach a decision compare to the number of samples needed by the fixed length decision rule to obtain the same P_D and P_{FA}?

6 Report Guidelines

You will submit a report on your project. The following guidelines should be followed:

1. Your report should be neatly formatted; text should be typed, and equations may be neatly handwritten.

2. Your report should include all derivations and analyses that you do to complete the project.

3. Your report should include outlines of the programs that you write and descriptions of their function. Do not include the source code of your programs in the main body of the report; they should be in an appendix. (I want to be able to find out what you have done without looking through your source code.)

4. Your report should explain in detail all experiments that you perform with your computer programs.

5. Discuss and explain the observations that you make in the course of the project; include graphs and figures where appropriate and useful.