UNIFORMLY MOST POWERFUL CYCLIC PERMUTATION INVARIANT DETECTION FOR DISCRETE-TIME SIGNALS

F. C. Nicolls and G. de Jager
Department of Electrical Engineering, University of Cape Town
Rondebosch 7701, South Africa
Email: nicolls@dip.ee.uct.ac.za

Abstract

The uniformly most powerful invariant test is derived for the problem of detecting a signal with unknown location in a sequence of noise. This test has the property that for each possible signal location it has the greatest power of all tests which are invariant to cyclic permutations of the observations.

The test is compared to the generalised likelihood ratio test, which is more typically used for this detection problem. Monte-Carlo simulations are used to show that the powers of the two tests are comparable, and there is little loss involved in choosing whichever is simpler.

1 Introduction

The problem of detecting a known signal with unknown location in a sample of white Gaussian noise has been treated in the literature. A very simple approach is to use the magnitude of the Fourier transform of the observation as a cyclic-shift invariant statistic. However, because this statistic is not maximal, it is invariant to more than just translations. In fact, Hayes [2] has shown for the discrete case that even if the equivalence class includes circular shifts, time reversal, and change of sign of the sequence, the statistic is still not maximal. The result is suboptimality of any detector based on this quantity.

A more justifiable approach is to use the generalised likelihood ratio test (GLRT) formalism. Here the signal location is treated as an unknown parameter, which is estimated and used in the conventional likelihood ratio test (LRT). This gives rise to a convolution-style implementation where a window is moved across the data observation and a relevant statistic calculated for each position. A decision of target present is then made if any of these statistics exceed a predefined threshold. Under some fairly general conditions the GLRT is known to be asymptotically optimal [5, p. 240]. However, from a signal processing perspective this property is of questionable value since samples are seldom large. Additionally, there is no indication of just how many samples are required before the asymptotic distributions become appropriate.

In this paper, an invariance argument is used to derive the uniformly most powerful invariant (UMPI) test for detecting a signal in a sample of Gaussian noise. The signal is assumed to be known only to within an arbitrary cyclic permutation of its elements, which for discrete-time signals is the natural counterpart to unknown location (and shift invariance) in continuous time. The performance of this detector is then compared to the GLRT which, since it shares the same invariances, necessarily has lower power than the UMPI detector. It is shown that in at least some cases of practical interest the difference between the two detectors is negligible. This is a significant result since it validates the use of the simpler GLRT.

2 Problem formulation

It is assumed that N samples $x_1, \dots, x_N$ of data are observed. Under hypothesis H_0, these samples are independent and identically distributed as $N[0, \sigma^2]$ — a more general case is considered in Section 8. Under hypothesis H_1, some shifted version of the prototype target signal $s_1, \dots, s_N$ is added to the noise samples. Since for discrete-time observations it is natural to regard shifts as cyclic permutations of the elements, under H_1 the mean of the observations is some cyclic permutation of $s_1, \dots, s_N$.

The problem is most easily described in vector notation: letting $\mathbf{x} = (x_1, \dots, x_N)^T$, $\mathbf{n} = (n_1, \dots, n_N)^T$, and $\mathbf{s} = (s_1, \dots, s_N)^T$, the hypotheses are

$$H_0 : \mathbf{x} = \mathbf{n} \tag{1}$$

versus

$$H_1 : \mathbf{x} = \mathbf{P}^{\theta}\mathbf{s} + \mathbf{n}, \tag{2}$$

where $\mathbf{n} : N[\mathbf{0}, \sigma^2\mathbf{I}]$ and $\mathbf{P}$ is the cyclic permutation matrix

$$\mathbf{P} = \begin{pmatrix} 0 & \cdots & 0 & 1 \\ 1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 \end{pmatrix}. \tag{3}$$

It can be seen that premultiplying the column vector $\mathbf{x}$ by $\mathbf{P}$ cyclically permutes the elements one position downwards:

$$\begin{pmatrix} 0 & \cdots & 0 & 1 \\ 1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_{n-1} \\ x_n \end{pmatrix} = \begin{pmatrix} x_n \\ x_1 \\ \vdots \\ x_{n-1} \end{pmatrix}. \tag{4}$$

Thus, under H_1 the mean of $\mathbf{x}$ is some cyclic permutation of the signal vector $\mathbf{s}$, where the precise order of the permutation is unknown. θ is assumed to be an unknown deterministic quantity, and without loss of generality can be restricted to take on integer values from 0 to $N-1$.

3 GLRT solution to the problem

In this section the GLRT for the problem is derived. This is a typical test which is used in testing under composite hypotheses, where there are unknown (nuisance) parameters in the formulation (in this case the signal location). Effectively, maximum likelihood estimates are made of the unknown parameters under each hypothesis, and the resulting density functions used in a conventional likelihood ratio test.

The GLRT statistic is

$$\Lambda_{\text{GLRT}}(\mathbf{x}) = \ln \frac{\max_{\theta \in [0,N-1]} p_1(\mathbf{x}|\theta)}{p_0(\mathbf{x})}, \tag{5}$$

where $p_0(\mathbf{x})$ is the probability density function (pdf) of the observation $\mathbf{x}$ under H_0, and $p_1(\mathbf{x})$ the pdf under H_1. Assuming the distributions given in the previous section,

$$\begin{aligned} \Lambda_{\text{GLRT}}(\mathbf{x}) &= \max_{\theta \in [0,N-1]} -\frac{1}{2\sigma^2} \left\{ (\mathbf{x} - \mathbf{P}^\theta \mathbf{s})^T (\mathbf{x} - \mathbf{P}^\theta \mathbf{s}) - \mathbf{x}^T \mathbf{x} \right\} \\ &= \frac{1}{\sigma^2} \left\{ \max_{\theta \in [0,N-1]} (\mathbf{P}^\theta \mathbf{s})^T \mathbf{x} + \frac{1}{2} \mathbf{s}^T \mathbf{s} \right\}. \end{aligned} \tag{6}$$

The GLRT is to compare this statistic to a threshold, and decide H_1 when exceeded:

$$\begin{aligned} \frac{1}{\sigma^2} \left\{ \max_{\theta \in [0,N-1]} (\mathbf{P}^\theta \mathbf{s})^T \mathbf{x} + \frac{1}{2} \mathbf{s}^T \mathbf{s} \right\} &\underset{H_0}{\overset{H_1}{\gtrless}} \eta \\ \max_{\theta \in [0,N-1]} (\mathbf{P}^\theta \mathbf{s})^T \mathbf{x} &\underset{H_0}{\overset{H_1}{\gtrless}} \sigma^2 \eta - \frac{1}{2} \mathbf{s}^T \mathbf{s}. \end{aligned} \tag{7}$$

It can be seen that the final quantity on the left of this test is simply the maximum of the inner products between the observation $\mathbf{x}$ and all possible cyclic permutations of the signal $\mathbf{s}$. Since the null hypothesis is independent of the unknown parameter θ, the threshold can be chosen such that the test has a constant false alarm rate.

4 Invariance of the hypothesis testing problem

There is no uniformly most powerful (UMP) test for the hypothesis testing problem under consideration here. In the search for an optimal characterisation, it is therefore necessary to restrict the class of tests which are to be considered. For the case of detecting a signal with unknown location, it is natural to require that the hypothesis test be constrained such that the same decision be made for arbitrarily shifted versions of any given observation. The transformation group considered for the problem is therefore

$$\mathcal{G} = \{g(\mathbf{x}) | g(\mathbf{x}) = \mathbf{P}^k \mathbf{x}, k = 0, \ldots, N-1\}. \tag{8}$$

This places an equivalence on the observations $\{\mathbf{P}^0\mathbf{x}, \ldots, \mathbf{P}^{N-1}\mathbf{x}\}$, which is natural on account of the symmetry of the elements of the observations under each hypothesis. Thus the observations $(x_1, \ldots, x_{N-1}, x_N) \equiv (x_N, x_1, \ldots x_{N-1}) \equiv \cdots \equiv (x_2, \ldots, x_n, x_1)$ are all considered to be equivalent by the detector. By the way the hypothesis testing problem has been formulated, it cannot be said that enforcing this equivalence is restricting the form of the detector in any unreasonable way.

The testing problem can be seen to be invariant to the group $\mathcal{G}$. This can be established by considering the distributions of the observation $\mathbf{x}$ under each hypothesis: in both cases $\mathbf{x}$ is MVN with covariance matrix $\sigma^2\mathbf{I}$. However, under H_0 the mean is $\mathbf{0}$, and under H_1 it is one of the elements in the set $\{\mathbf{P}^0\mathbf{s}, \ldots, \mathbf{P}^{N-1}\mathbf{s}\}$. Consider now an element $g_k(\mathbf{x}) = \mathbf{P}^k\mathbf{x}$ of the group $\mathcal{G}$. Since this is a linear transformation of $\mathbf{x}$, the distribution of $\mathbf{y} = g_k(\mathbf{x})$ is $N[\mathbf{P}^k E\mathbf{x}, \sigma^2 \mathbf{P}^k (\mathbf{P}^k)^T]$, where $E\mathbf{x}$ is the expected value of $\mathbf{x}$. Noting now that $(\mathbf{P}^k)^T = \mathbf{P}^{N-k} = \mathbf{P}^{-k}$,

$$\mathbf{y} : N[\mathbf{P}^k E\mathbf{x}, \sigma^2 \mathbf{I}]. \tag{9}$$

Thus under H_0 the mean of the transformed vector $\mathbf{y}$ is $\mathbf{0}$, and under H_1 it is an element of the set $\{\mathbf{P}^k\mathbf{P}^0\mathbf{s}, \ldots, \mathbf{P}^k\mathbf{P}^{N-1}\mathbf{s}\} = \{\mathbf{P}^0\mathbf{s}, \ldots, \mathbf{P}^{N-1}\mathbf{s}\}$. The transformation g_k therefore preserves the form of the distribution and retains the partition of the parameter space under each hypothesis. Since this conclusion is valid for all elements $g \in \mathcal{G}$, the hypothesis testing problem is invariant-$\mathcal{G}$.

5 Maximal invariant statistic for the problem

Before being able to continue, it is necessary to find a maximal invariant statistic for the problem. This is a statistic which has the required invariances, but which also manages to retain all the useful information contained in the observation regarding the decision process. One such statistic can be obtained by defining $\mathbf{P}_{\max}(\mathbf{x})$ to be that function which cyclically permutes the elements of $\mathbf{x}$ until the element of $\mathbf{x}$ with the maximum value is in the first position. Note that for the distributions being considered here $\Pr\{x_i = x_j\} = 0$ for $i \neq j$, so the maximum element of $\mathbf{x}$ will be unique with probability 1.

The statistic $\mathbf{P}_{\max}(\mathbf{x})$ is clearly invariant to the group $\mathcal{G}$: since one of the elements of $\mathbf{x}$ is always maximum and elements of $\mathcal{G}$ simply permute the observation $\mathbf{x}$ cyclically, $\mathbf{P}_{\max}[g(\mathbf{x})] = \mathbf{P}_{\max}[\mathbf{x}]$ for all $g \in \mathcal{G}$. Additionally, for the same reasons, the condition $\mathbf{P}_{\max}[g(\mathbf{x}_1)] = \mathbf{P}_{\max}[g(\mathbf{x}_2)]$ means that $\mathbf{x}_1$ and $\mathbf{x}_2$ must be related to one another through a cyclic shift, so $\mathbf{x}_2 = g(\mathbf{x}_1)$ for some $g \in \mathcal{G}$. Thus the statistic $\mathbf{P}_{\max}(\mathbf{x})$ is maximal.

As explained by Lehmann [6] or Scharf [7], the significance of this result is that only functions of the maximal invariant statistic have to be considered when looking for a test which is invariant to $\mathcal{G}$.

6 Distribution of the maximal invariant statistic

The method described by Hogg and Craig [3, p. 142] in relation to order statistics provides a means of determining the distribution of the maximal invariant. Firstly it is reasserted that two elements of $\mathbf{x}$ are equal with probability zero, so the joint probability density of $\mathbf{x}$ can be defined to be zero at all points which have at least two of their coordinates equal. The set $\mathcal{A}$ where the probability density of $\mathbf{x}$ is nonzero can then be partitioned into N mutually disjoint sets:

$$\begin{aligned} \mathcal{A}_1 &= \{\mathbf{x} | x_1 = \max(x_1, \dots, x_N)\} \\ &\vdots \\ \mathcal{A}_N &= \{\mathbf{x} | x_N = \max(x_1, \dots, x_N)\}. \end{aligned} \tag{10}$$

Thus the set $\mathcal{A}_i$ is the set of all points in $\mathbb{R}^N$ which have no elements equal, and have x_i as their largest element.

Consider the function $\mathbf{y} = \mathbf{P}_{\max}(\mathbf{x})$. This defines a $1-1$ transformation of each of $\mathcal{A}_1, \dots, \mathcal{A}_N$ onto the same set $\mathcal{B}$, where it so happens that $\mathcal{B} = \mathcal{A}_1$. For points in $\mathcal{A}_i$, the transformation $\mathbf{y} = \mathbf{P}_{\max}(\mathbf{x})$ cyclically permutes the elements of $\mathbf{x}$ upwards by $i-1$ positions. Thus the inverse function is

$$\mathbf{x} = \mathbf{P}^{i-1}\mathbf{y} \tag{11}$$

which simply rotates them back downwards by the same amount.

Letting J_i be the determinant of the Jacobian of the inverse transformation corresponding to $\mathcal{A}_i$, it can be seen that

$$J_i = |\mathbf{P}^{i-1}|. \tag{12}$$

Now by the structure of $\mathbf{P}^{i-1}$, it is always possible to obtain an identity matrix by means of a number of row exchanges. Thus it must be the case that $J_i = +1$ or $J_i = -1$. Denoting the probability density of $\mathbf{x}$ by $f_x(\mathbf{x})$, the results of this section can be combined to yield the corresponding pdf $f_y(\mathbf{y})$ of $\mathbf{y} = \mathbf{P}_{\max}(\mathbf{x})$ as [3, p. 143]

$$f_y(\mathbf{y}) = \begin{cases} \sum_{k=0}^{N-1} f_x(\mathbf{P}^k \mathbf{y}) & y_1 = \max(y_1, \dots, y_N) \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

This expression can finally be used to find the distribution of the maximal invariant statistic under each hypothesis. Under H_0,

$$f_x(\mathbf{x}) = (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{x}^T\mathbf{x}}. \tag{14}$$

Therefore the distribution of $\mathbf{y} = \mathbf{P}_{\max}(\mathbf{x})$ is

$$f_y(\mathbf{y}) = \begin{cases} \sum_{k=0}^{N-1} (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{P}^k\mathbf{y})^T(\mathbf{P}^k\mathbf{y})} \\ \qquad\qquad y_1 = \max(y_1, \dots, y_N) \\ 0 \qquad\qquad \text{otherwise.} \end{cases} \tag{15}$$

Once again using the relation $(\mathbf{P}^k)^T = \mathbf{P}^{-k}$, this can be simplified to

$$f_y(\mathbf{y}) = \begin{cases} N(2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{y}} \\ \qquad\qquad y_1 = \max(y_1, \dots, y_N) \\ 0 \qquad\qquad \text{otherwise.} \end{cases} \tag{16}$$

When H_1 is in force, the situation is slightly more complex: now the mean of the observation takes some value in the set $\{\mathbf{P}^\theta\mathbf{s}, \theta = 0, \dots, N-1\}$. The probability density of $\mathbf{x}$ is therefore

$$f_x(\mathbf{x}) = (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{x}-\mathbf{P}^\theta\mathbf{s})^T(\mathbf{x}-\mathbf{P}^\theta\mathbf{s})}, \tag{17}$$

where θ is some integer in the range 0 to $N-1$. Substituting into the expression for $f_y(\mathbf{y})$ gives

$$f_y(\mathbf{y}) = \begin{cases} \sum_{k=0}^{N-1} (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{P}^k\mathbf{y}-\mathbf{P}^\theta\mathbf{s})^T(\mathbf{P}^k\mathbf{y}-\mathbf{P}^\theta\mathbf{s})} \\ \qquad\qquad y_1 = \max(y_1, \dots, y_N) \\ 0 \qquad\qquad \text{otherwise.} \end{cases} \tag{18}$$

The first case in this expression needs to be looked at in more detail: under the condition $y_1 = \max(y_1, \dots, y_N)$,

$$f_y(\mathbf{y}) = \sum_{k=0}^{N-1} (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{P}^{-k}\mathbf{P}^k\mathbf{y} - 2\mathbf{s}^T\mathbf{P}^{-\theta}\mathbf{P}^k\mathbf{y} + \mathbf{s}^T\mathbf{P}^{-\theta}\mathbf{P}^{\theta}\mathbf{s}}$$
$$= (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y}^T\mathbf{y}+\mathbf{s}^T\mathbf{s})} \sum_{k=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^{-\theta}\mathbf{P}^k\mathbf{y}}. \tag{19}$$

Using the fact that $\mathbf{P}^{-l} = \mathbf{P}^{N-l}$, the sum in this final expression can be written as

$$\begin{aligned}
\sum_{k=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^{-\theta}\mathbf{P}^k\mathbf{y}} &= \sum_{k=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^{k-\theta}\mathbf{y}} \\
&= \sum_{l=-\theta}^{-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} + \sum_{l=0}^{(N-1)-\theta} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} \\
&= \sum_{l=N-\theta}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} + \sum_{l=0}^{(N-\theta)-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} \\
&= \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}},
\end{aligned} \tag{20}$$

resulting in the final pdf for $\mathbf{y}$ under H_1 as

$$f_y(\mathbf{y}) = \begin{cases} (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y}^T\mathbf{y}+\mathbf{s}^T\mathbf{s})} \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} & \\ \qquad\qquad y_1 = \max(y_1, \dots, y_N) & \\ 0 \qquad \text{otherwise.} & \end{cases} \tag{21}$$

Under both cases the density of the maximal invariant is seen to be independent of the unknown parameter θ, as required.

7 Optimal invariant likelihood ratio test

Once the observation $\mathbf{x}$ has been mapped onto the corresponding maximal invariant statistic, a likelihood ratio test can be performed on this quantity. The likelihood ratio for the problem is

$$\begin{aligned}
l(y_1, \dots, y_N) &= \frac{(2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{y}} e^{-\frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s}} \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}}}{N(2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{y}}} \\
&= \frac{1}{N} e^{-\frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s}} \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}}.
\end{aligned} \tag{22}$$

The log-likelihood ratio is therefore

$$L(\mathbf{y}) = -\ln N - \frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s} + \ln \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}}. \tag{23}$$

The best invariant test is to compare this ratio to a threshold, and decide H_1 when exceeded:

$$\begin{aligned}
-\ln N - \frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s} + \ln \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} &\underset{H_0}{\overset{H_1}{\gtrless}} \eta \\
\ln \sum_{l=0}^{N-1} e^{\mathbf{s}^T\mathbf{P}^l\mathbf{y}} &\underset{H_0}{\overset{H_1}{\gtrless}} \eta + \ln N + \frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s}
\end{aligned} \tag{24}$$

The resulting test is uniformly most powerful out of all tests which share the same invariances. This means that no other test which is invariant to cyclic permutations of the observations can perform as well as this test, regardless of the value of the unknown parameter θ. Since the invariance is a perfectly reasonable one for the problem, it is fair to assert that this is the optimal test.

It is worth noting that the estimation of the parameter θ is explicit in the GLRT of equation 7. Thus the most likely location of the detected signal is also provided by the test. For the UMPI test, however, the dependence on the parameter is completely eliminated from the problem by the invariance condition. At no point does this test make use of an estimate of θ, either implicitly or explicitly.

8 Extension to correlated noise

The previous results were based on detection in a white noise environment. The results can be extended to the case where the noise has a known circulant covariance matrix. This is a special case of the general stationary condition, where the matrix has a Toeplitz structure. The constraint that the covariance matrix be circulant is required in order that the hypothesis testing problem remain invariant to cyclic permutations.

For this case, the hypotheses are as in equations 1 and 2, but now the distribution of the noise is $\mathbf{n} : N[\mathbf{0}, \mathbf{C}]$. Applying the (assumed invertible) whitening transformation $\mathbf{z} = \mathbf{C}^{-1/2}\mathbf{x}$ to the observed data, the hypotheses become $\mathbf{z} : N[\mathbf{0}, \mathbf{I}]$ under H_0 and $\mathbf{z} : N[\mathbf{C}^{-1/2}\mathbf{P}^{\theta}\mathbf{s}, \mathbf{I}]$ under H_1. Now, if $\mathbf{C}$ is circulant then $\mathbf{C}^{-1/2}$ is also circulant, so $\mathbf{C}^{-1/2} = \mathbf{P}^{\theta}\mathbf{C}^{-1/2}\mathbf{P}^{-\theta}$. Writing this as $\mathbf{C}^{-1/2}\mathbf{P}^{\theta} = \mathbf{P}^{\theta}\mathbf{C}^{-1/2}$, the distribution under H_1 is $\mathbf{z} : N[\mathbf{P}^{\theta}\mathbf{C}^{-1/2}\mathbf{s}, \mathbf{I}]$. This can be recognised as the problem of invariant detection of the modified signal $\mathbf{C}^{-1/2}\mathbf{s}$ in white noise. The test given in the previous section can therefore be used in this modified problem, and is once again UMPI.

Finally, it is noted that the components of a random vector with a circulant covariance matrix can be diagonalised by means of the discrete Fourier transform (DFT). This can provide a fast method of calculating the required test statistic.

9 Distributions of the test statistics

Summarising the results of the previous sections, the GLRT is

$$t_{\mathrm{GLRT}}(\mathbf{x}) = \max_{\theta \in [0,N-1]} (\mathbf{P}^{\theta}\mathbf{s})^T \mathbf{x} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_{\mathrm{GLRT}}, \qquad (25)$$

and the UMPI test is

$$t_{\mathrm{UMPI}}(\mathbf{x}) = \ln \sum_{l=0}^{N-1} e^{\mathbf{s}^T \mathbf{P}^l \mathbf{y}} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_{\mathrm{UMPI}}. \qquad (26)$$

The thresholds η_{GLRT} and η_{UMPI} are constants which are chosen to yield the desired false alarm rate.

The components of the sum in the UMPI statistic are statistically dependent upon one another, and each have a log-normal distribution. The UMPI statistic is therefore given by the logarithm of the sum of dependent lognormal variates. Sums of dependent and independent lognormal random variables have been discussed at some length in the literature, and it is well-known that no closed form exists for their distribution [1]. It is common therefore to approximate the sum by yet another lognormal distribution (which is appropriate particularly in the tails [4]), and perform identification by means of moment matching[8]. The effectiveness of this approach seems reasonable for some commonly occurring situations discussed in the references.

In order not to get engrossed in the details of analytical approximations, the results in this paper will be based instead on Monte-Carlo simulations. The distributions of the test statistics can be very easily obtained by applying the tests to a large sample of observations from any specific configuration of signal and noise. The signals used for testing purposes are shown in Figure 1. These signals are normalised

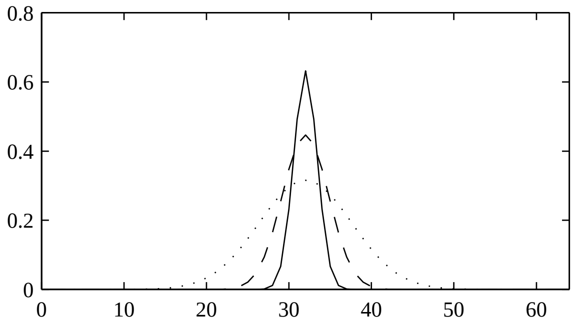

Figure 1: Target signals used in Monte-Carlo simulations. The signals are Gaussians centred on 32 with standard deviations of 2, 4, and 8.

to have unit energy. Various scalings of these signals are considered, corresponding to signals with energies of 2, 4, and 8. In all cases, the additive noise is comprised of 32 uncorrelated samples of zero-mean Gaussian noise with unit variance.

Figure 2 shows distributions for each of the test statistics for the case of the signals rescaled to have an energy of 8. The results correspond to the signals shown in Figure 1, and in each case the distribution to the right corresponds to H_1.

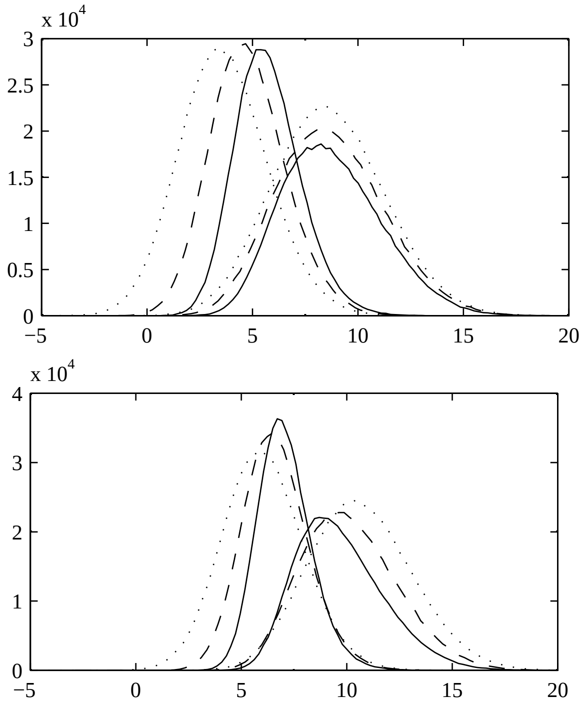

Figure 2: Distributions of the GLRT statistic (top) and UMPI statistic (bottom) under both H_0 and H_1 for each of the signals tested. The signals were rescaled to have an energy of 8.

10 Comparison of GLRT and UMPI test powers

In order to estimate the receiver operating characteristics (ROC) of the tests, 500000 samples of noise for nine cases of interest were generated. These comprised the three signals in Figure 1, scaled to have energies of 2, 4, and 8. The results are shown in Figure 3, each plot demonstrating the results for one specific energy. The curves relate to the signals plotted with line styles corresponding to those used in Figure 1. The ROC curves for both the GLRT and the UMPI tests are plotted in the same style, with the UMPI ROC always being the upper one. (This is necessarily the case since the UMPI test is uniformly more powerful than any other invariant test.)

It is clear from the plots that the difference in performance between the tests is marginal for the cases which were analysed. In addition, the differences seem to become even less

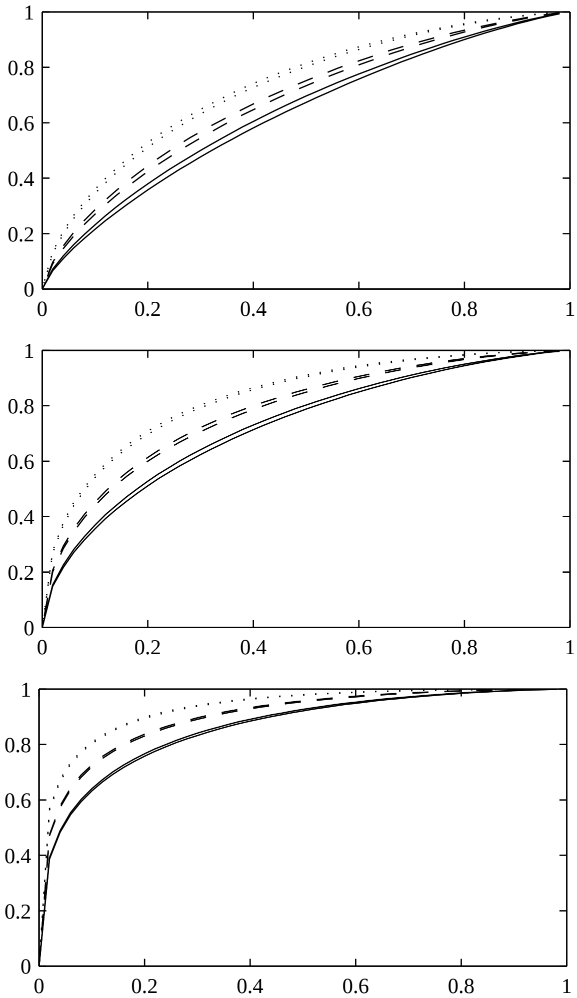

Figure 3: ROC curves for detection of the test signals scaled to have energy of 2 (top), 4 (middle), and 8 (bottom).

pronounced when the energy of the signal to be detected becomes high.

11 Discussion and Conclusions

In this paper a test for detecting a signal with unknown location in white noise is derived, which is uniformly most powerful in the class of all tests which are invariant to cyclic permutations of the observations. It is demonstrated for some specific cases of interest that the performance of this test is not significantly better than that of the GLRT, which is a suboptimal but more common solution to the problem.

The importance of these results are twofold. Firstly, insofar as the invariances are sensible for the problem, the UMPI test cannot be improved upon by any other test which shares the same invariances. Thus it provides an ideal baseline against which the performance of competing tests can be assessed. Secondly, the fact that the power of the two tests is comparable means that detectors based on the conventional GLRT are not significantly suboptimal. Thus it provides evidence that the asymptotic approach of the GLRT to the UMPI test occurs quite quickly with increasing observation length.

Invariance to cyclic permutations is not always strictly appropriate for all unknown signal location problems. Some problems are not inherently cyclic, especially those which result from discretisation of continuous-time problems. A UMP test cannot be expected for these situations: invariance is essentially a symmetry condition, and can necessarily only be applied to situations which exhibit the required symmetries. Nonetheless, the results given in this paper provide insight into the nature of the detection process, and lend credibility to the GLRT.

References

[1] Abu-Dayya, A. A. and Beaulieu, N. C., "Outage probabilities in the presence of correlated lognormal interferers," *IEEE Transactions on Vehicular Technology*, vol. 43, pp. 164–173, Feb. 1994.

[2] Hayes, M. H., "The reconstruction of a multidimensional sequence from the phase or magnitude of its Fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-30, pp. 140–154, Apr. 1982.

[3] Hogg, R. V. and Craig, A. T., *Introduction to Mathematical Statistics*. Macmillan, third ed., 1970.

[4] Janos, W. A., "Tail of the distribution of sums of lognormal variates," *IEEE Transactions on Information Theory*, vol. IT-16, pp. 299–302, May 1970.

[5] Kendall, M. and Stuart, A., *The Advanced Theory of Statistics*, vol. 2. Charles Griffin & Company, fourth ed., 1979.

[6] Lehmann, E. L., *Testing Statistical Hypotheses*. Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, second ed., 1986.

[7] Scharf, L. L., *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison-Weseley Series in Electrical and Computer Engineering: Digital Signal Processing, Addison-Wesley, 1991.

[8] Schwartz, S. C. and Yeh, Y. S., "On the distribution function and moments of power sums with log-normal components," *Bell Systems Technical Journal*, vol. 61, pp. 1441–1462, Sept. 1982.