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

F. Nicolls and G. de Jager

Department of Electrical Engineering, University of Cape Town
Rondebosch 7701, South Africa

ABSTRACT

The uniformly most powerful invariant (UMPI) test is derived for detecting a target with unknown location in a noise sequence. This test has the property that for each possible target 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 (GLRT), which is commonly used as a solution to this detection problem. Monte-Carlo simulations show that the powers of the two tests are comparable, thereby justifying near-optimality of the GLRT.

1. INTRODUCTION

The problem of detecting a known target with unknown location in a sample of noise often arises in signal and image processing. A common solution involves the use of a generalised likelihood ratio testing (GLRT) formalism, where the target location is treated as an unknown parameter. An estimate of the location is made as part of the detection procedure, and the value obtained is used in a conventional likelihood ratio test (LRT). The resulting test has no optimality properties.

In this paper, an invariance argument is used to derive the uniformly most powerful invariant (UMPI) test for the specific case where shifts are defined to be circular. The target is therefore assumed to be known only to within an arbitrary cyclic permutation of its elements. For discrete-time signals is the natural counterpart to unknown location (and shift invariance) in continuous time. The performance of this detector is compared to the GLRT, which, since it shares the same invariances, necessarily has lower power than the UMPI detector. It is shown in at least some cases of practical interest that the difference between the UMPI test and the GLRT is negligible. This is a significant result, since it indicates that for the cases investigated the GLRT is near-optimal. It is also demonstrated that this conclusion could not have been reached by comparing the GLRT to the ideal matched filter, which assumes the location of the target known.

2. PROBLEM FORMULATION

It is assumed that N samples $x_1, \ldots, x_N$ 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 $s_1, \ldots, s_N$ 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 $s_1, \ldots, s_N$.

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

The GLRT is commonly used for composite hypothesis testing problems: maximum likelihood estimates are made of the unknown parameters under each hypothesis, and the resulting density functions used in a conventional likelihood ratio test. For the case of unknown target location, the GLRT statistic is

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

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,

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

The GLRT compares this statistic to a threshold, with H_1 being decided when exceeded:

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

The 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 target $\mathbf{s}$. 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

Since no uniformly most powerful (UMP) test exists for the problem outlined, in the search for an optimal characterisation it is necessary to restrict the class of permissible tests. For the problem considered, 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 for the problem is therefore

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

This places an equivalence on the observations $\{\mathbf{P}^0\mathbf{x}, \dots, \mathbf{P}^{N-1}\mathbf{x}\}$, which is natural on account of the symmetry of the elements under each hypothesis. Thus the observations $(x_1, \dots, x_{N-1}, x_N) \equiv (x_N, x_1, \dots x_{N-1}) \equiv \dots \equiv (x_2, \dots, x_n, x_1)$ are all considered to be equivalent by the detector. Enforcing this equivalence is in no way restricting the form of the test in any unreasonable way.

The testing problem is invariant to the group $\mathcal{G}$: under both hypothesis $\mathbf{x}$ 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 $\{\mathbf{P}^0\mathbf{s}, \dots, \mathbf{P}^{N-1}\mathbf{s}\}$. Consider now an element $\mathbf{g}_k(\mathbf{x}) = \mathbf{P}^k\mathbf{x}$ of the group $\mathcal{G}$. Since this is a linear transformation of $\mathbf{x}$, the distribution of $\mathbf{y} = \mathbf{g}_k(\mathbf{x})$ 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 $\{\mathbf{P}^k\mathbf{P}^0\mathbf{s}, \dots, \mathbf{P}^k\mathbf{P}^{N-1}\mathbf{s}\} = \{\mathbf{P}^0\mathbf{s}, \dots, \mathbf{P}^{N-1}\mathbf{s}\}$. The transformation $\mathbf{g}_k$ 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 $\mathbf{g} \in \mathcal{G}$, the hypothesis testing problem is invariant-$\mathcal{G}$.

5. MAXIMAL INVARIANT STATISTIC FOR THE PROBLEM

To continue, a maximal invariant statistic for the problem is required. 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, $\mathbf{P}_{\max}[\mathbf{g}(\mathbf{x})] = \mathbf{P}_{\max}[\mathbf{x}]$ for all $\mathbf{g} \in \mathcal{G}$. Additionally, for the same reasons, the condition $\mathbf{P}_{\max}[\mathbf{g}(\mathbf{x}_1)] = \mathbf{P}_{\max}[\mathbf{g}(\mathbf{x}_2)]$ means that $\mathbf{x}_1$ and $\mathbf{x}_2$ must be related to one another through a cyclic shift, so $\mathbf{x}_2 = \mathbf{g}(\mathbf{x}_1)$ for some $\mathbf{g} \in \mathcal{G}$. Thus the statistic $\mathbf{P}_{\max}(\mathbf{x})$ is maximal.

As explained by Lehmann [1] or Scharf [2], 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. Since two elements of $\mathbf{x}$ are equal with probability zero, 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 $\mathcal{A}_i$ 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 $\mathbf{x} = \mathbf{P}^{i-1}\mathbf{y}$, 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

$$J_i = |\mathbf{P}^{i-1}|. \tag{11}$$

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]

$$f_y(\mathbf{y}) = \begin{cases} \sum_{k=0}^{N-1} f_x(\mathbf{P}^k\mathbf{y}) & y_1 = \max(y_1, \dots, y_N) \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

The distribution of the maximal invariant statistic can now be found under each hypothesis. Under H_0,

$$f_x(\mathbf{x}) = (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{x}^T\mathbf{x}}. \tag{13}$$

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

$$f_y(\mathbf{y}) = \begin{cases} N(2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{y}} & \\ & y_1 = \max(y_1, \dots, y_N) \\ 0 & \text{otherwise.} \end{cases} \tag{14}$$

where use has been made of the relation $(\mathbf{P}^k)^T = \mathbf{P}^{-k}$. When H_1 is in force, 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

$$f_x(\mathbf{x}) = (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{x}-\mathbf{P}^\theta\mathbf{s})^T(\mathbf{x}-\mathbf{P}^\theta\mathbf{s})}, \tag{15}$$

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

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

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

$$f_y(\mathbf{y}) = (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y}^T\mathbf{y}+\mathbf{s}^T\mathbf{s})} \sum_{k=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^{-\theta}\mathbf{P}^k\mathbf{y}}. \quad (17)$$

Using the fact that $\mathbf{P}^{-l} = \mathbf{P}^{N-l}$ and noting that the sum in this expression involves the inner product between $\mathbf{s}$ and all cyclic permutations of $\mathbf{y}$, the range of summation can be modified:

$$\sum_{k=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^{-\theta}\mathbf{P}^k\mathbf{y}} = \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}}. \quad (18)$$

This yields the final pdf for $\mathbf{y}$ under H_1 as

$$f_y(\mathbf{y}) = \begin{cases} (2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y}^T\mathbf{y}+\mathbf{s}^T\mathbf{s})} \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}} \\ \qquad\qquad y_1 = \max(y_1, \dots, y_N) \\ 0 \qquad\qquad \text{otherwise.} \end{cases} \quad (19)$$

Under both hypotheses the density of the maximal invariant is 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

$$\begin{aligned} l(\mathbf{y}) &= \frac{(2\pi\sigma^2)^{-N/2} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{y}} e^{-\frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s}} \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}}}{N(2\pi\sigma^2)^{-\frac{N}{2}} e^{-\frac{1}{2\sigma^2}\mathbf{y}^T\mathbf{y}}} \\ &= \frac{1}{N} e^{-\frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s}} \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}}. \end{aligned} \quad (20)$$

The log-likelihood ratio is therefore

$$L(\mathbf{y}) = -\ln N - \frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s} + \ln \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}}. \quad (21)$$

The best invariant test involves comparing this ratio to a threshold, and deciding H_1 when exceeded:

$$\ln \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}} \underset{H_0}{\overset{H_1}{\gtrless}} \eta + \ln N + \frac{1}{2\sigma^2}\mathbf{s}^T\mathbf{s}. \quad (22)$$

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

Again noting that the summation in the expression for this test goes over terms involving inner products between all cyclic permutations of $\mathbf{y}$ with $\mathbf{s}$, it is evident that

$$\sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}} = \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{x}} \quad (23)$$

Thus the test in Equation 22 can be written in terms of the original data observation $\mathbf{x}$.

An estimate 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 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 is Toeplitz. The constraint that the covariance matrix be circulant is required to ensure invariance of the hypothesis testing problem to cyclic permutations.

Suppose that the hypotheses are as in equations 1 and 2, but with the noise distributed as $\mathbf{n} : N[\mathbf{0}, \mathbf{C}]$. 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 $\mathbf{C}^{-1/2} = \mathbf{P}^\theta\mathbf{C}^{-1/2}\mathbf{P}^{-\theta}$. Thus 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. COMPARISON OF GLRT AND UMPI TEST POWERS

Summarising the results of the previous sections, the GLRT is

$$t_{\text{GLRT}}(\mathbf{x}) = \max_{\theta \in [0, N-1]} (\mathbf{P}^\theta\mathbf{s})^T\mathbf{x} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_{\text{GLRT}}, \quad (24)$$

and the UMPI test is

$$t_{\text{UMPI}}(\mathbf{x}) = \ln \sum_{l=0}^{N-1} e^{\frac{1}{\sigma^2}\mathbf{s}^T\mathbf{P}^l\mathbf{y}} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_{\text{UMPI}}. \quad (25)$$

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 test are statistically dependent upon one another, and each have a lognormal 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 [4]. Maxima of correlated normal variates are similarly intractable.

To avoid the details of analytical approximations, Monte-Carlo methods are used to estimate the receiver operating characteristics (ROCs) for each of the tests. Targets used for testing purposes are shown in Figure 1. Three scalings of each of these targets are considered, corresponding to energies of 2, 4, and 8. In all cases, the additive noise is uncorrelated zero-mean Gaussian noise with

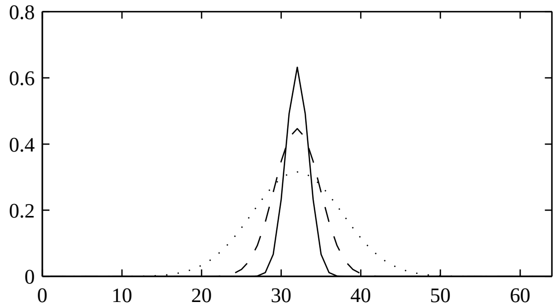

Fig. 1. Targets used in Monte-Carlo simulations. The signals are Gaussians centered on 32 with standard deviations of 2, 4, and 8.

unit variance. The test properties are estimated from a total of 50000 simulated noise and target plus noise samples in each case.

Results are shown in Figure 2, each plot demonstrating ROCs for one specific energy. The line styles used in these plots correspond to those used in Figure 1 for each of the targets. 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. It is seen that the performance of the GLRT is almost as good as that of the optimal UMPI test. This indicates near-optimality of the GLRT solution.

Also shown in the plots in Figure 2 are the ROCs for the ideal or clairvoyant matched filter, which assumes prior knowledge of the target location. This represents an upper bound on the performance of any test. Since matched filter performance depends only on the energy of the target being detected, a single curve can characterise the test properties for all three targets at any constant energy. It is observed that the performance of the GLRT is substantially worse than that of the matched filter, particularly for the more localised targets. In the absence of the UMPI test results, this could raise some doubt as to the effectiveness of the GLRT for the problem. However, the UMPI test results indicate that the decrease in performance is caused by the target location in fact being unknown, rather than by the GLRT being significantly suboptimal.

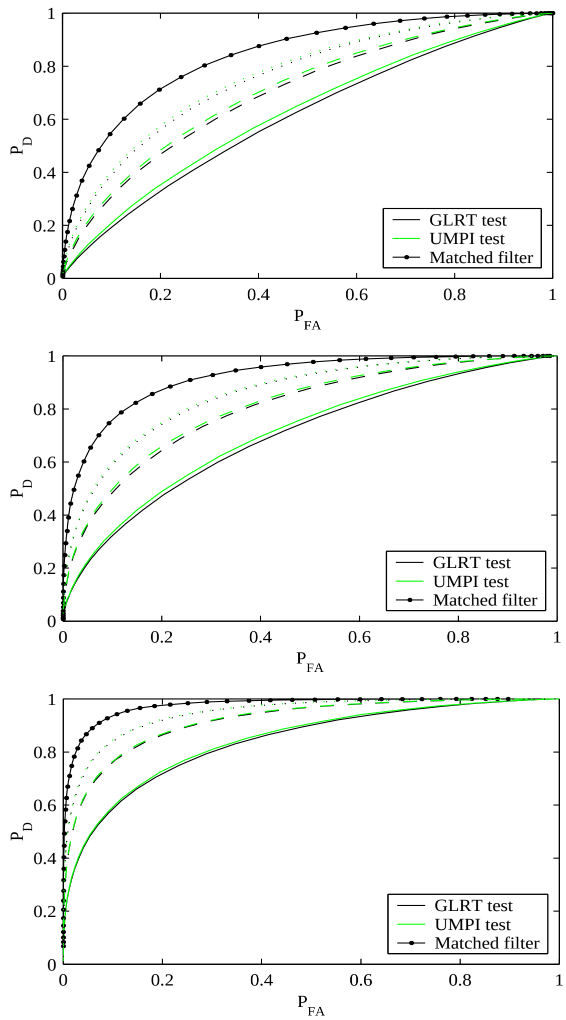

Fig. 2. ROC curves for detection of the test targets scaled to have energy of 2 (top), 4 (middle), and 8 (bottom).

10. DISCUSSION AND CONCLUSIONS

In this paper a test for detecting a target 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. The test can be extended to the correlated noise case, as long as the covariance matrix is circulant and known. 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. This provides a measure of justification for the use of the GLRT in problems of detecting targets with unknown location in noise.

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.

11. REFERENCES

[1] Erich Leo Lehmann, *Testing Statistical Hypotheses*, Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, second edition, 1986.

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

[3] Robert V. Hogg and Allen T. Craig, *Introduction to Mathematical Statistics*, Macmillan, third edition, 1970.

[4] S. C. Schwartz and Y. S. Yeh, "On the distribution function and moments of power sums with log-normal components," *Bell Systems Technical Journal*, vol. 61, no. 7, pp. 1441–1462, Sept. 1982.