

November 1996

LIDS-P-2369

Research Supported By:

Air Force Office of Scientific Research
(grants F49620-95-1-0083 F49620-96-1-0028)

Defense Advanced Research Projects Agency
(grant F49620-93-1-0604)

**A Multiscale Hypothesis Testing Approach to Anomaly Detection and
Localization From Noisy Tomographic Data**

Austin B. Frakt, W.Clem Karl, Alan S. Willsky

A Multiscale Hypothesis Testing Approach to Anomaly Detection and Localization From Noisy Tomographic Data

Austin B. Frakt, W. Clem Karl, Alan S. Willsky

Direct correspondence to Austin B. Frakt, 77 Massachusetts Ave., Room 35-425, Cambridge, MA 02139. Telephone: 617-253-4874. Fax: 617-258-8553. Email: frakt@mit.edu.

Austin B. Frakt and Alan S. Willsky are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology. W. Clem Karl is with the Department of Electrical, Computer, and Systems Engineering, Boston University.

This work is supported by the Air Force Office of Scientific Research (grants F49620-95-1-0083 and F49620-96-1-0028, the latter awarded to Boston University), the Defense Advanced Research Projects Agency (grant F49620-93-1-0604). Austin B. Frakt is supported by a National Defense Science and Engineering Graduate Fellowship awarded by the Defense Advanced Research Projects Agency.

Abstract

In this paper we investigate the problems of anomaly detection and localization from noisy tomographic data. These are characteristic of a class of problems which cannot be optimally solved because they involve hypothesis testing over hypothesis spaces with extremely large cardinality. Our multiscale hypothesis testing approach addresses the key issues associated with this class of problems. A multiscale hypothesis test is a hierarchical sequence of composite hypothesis tests which discards large portions of the hypothesis space with minimal computational burden and zooms in on the likely true hypothesis. For the anomaly detection and localization problems, hypothesis zooming corresponds to spatial zooming— anomalies are successively localized to finer and finer spatial scales. The key challenges we address include how to hierarchically divide a large hypothesis space and how to process the data at each stage of the hierarchy to decide which parts of the hypothesis space deserve more attention. To answer the former we draw on [1, 6–9]. For the latter, we pose and solve a non-linear optimization problem for a decision statistic which maximally disambiguates composite hypotheses. With no more computational complexity, our optimized statistic shows substantial improvement over conventional approaches. We provide examples which demonstrate this and which quantify how much performance is sacrificed by the use of a sub-optimal method as compared to that achievable if the optimal approach were computationally feasible.

Keywords

anomaly detection, tomography, hypothesis zooming, hypothesis testing, composite hypothesis testing, non-linear optimization, quadratic programming

I. INTRODUCTION

In this paper we present a new approach to hierarchical, multiresolution anomaly detection and localization from noisy tomographic data. This problem is of interest in several areas of active research. For example, it arises in non-destructive evaluation as well as adaptive tomographic reconstruction from limited data in which one attempts to use degrees of freedom frugally by localizing areas of interest to be imaged at finer resolutions. In addition, the anomaly detection

and localization problems raise issues that surface in many other contexts. Thus, a second objective of this paper is to provide some insight into these more general issues.

In particular, a fundamental characteristic of a large class of signal and image analysis problems is that they involve hypothesis testing over hypothesis spaces of extremely large cardinality—so large that enumeration of all hypotheses and exhaustive comparison is computationally infeasible. This characteristic is present in the specific problems considered in this paper since the enumeration of all possible anomalies leads to a very large hypothesis space. Large hypothesis spaces also arise in other applications including detection of regions of interest for automatic target recognition from wide-area imagery as well as model-based object recognition.

Such problems thus invite efficient hierarchical approaches. Our approach achieves great efficiency by discarding large parts of the hypothesis space with minimal computational burden and zooming in on a smaller part of the hypothesis space to be scrutinized more extensively and, perhaps, exhaustively. We call this hypothesis zooming methodology a *multiscale hypothesis test (MSHT)*. As depicted in Figure 1, a MSHT is a scale-recursive sequence of composite hypothesis tests which increasingly disambiguates hypotheses at finer scales and zooms in on one element of the global set of hypotheses, $\mathcal{H} \triangleq \{H_m\}_{m=0}^{M-1}$. At the top (or coarsest) scale, the MSHT divides \mathcal{H} into a family of, possibly overlapping, subsets and then chooses one of these subsets, discarding all the others. Similarly, at subsequent scales, the remaining hypotheses are divided into subsets, all but one of which are discarded. Note that the statistical decision problem at each scale is a composite hypothesis testing problem in which one composite hypothesis¹, $\tilde{H}_n \subset \mathcal{H}$, is selected for finer scale investigation.

There are several questions that must be addressed in formulating and solving problems in

¹We use a tilde ($\tilde{\cdot}$) to indicate composite hypotheses and quantities and objects associated with them to distinguish these from individual (non-composite) hypotheses and their associated similar objects.

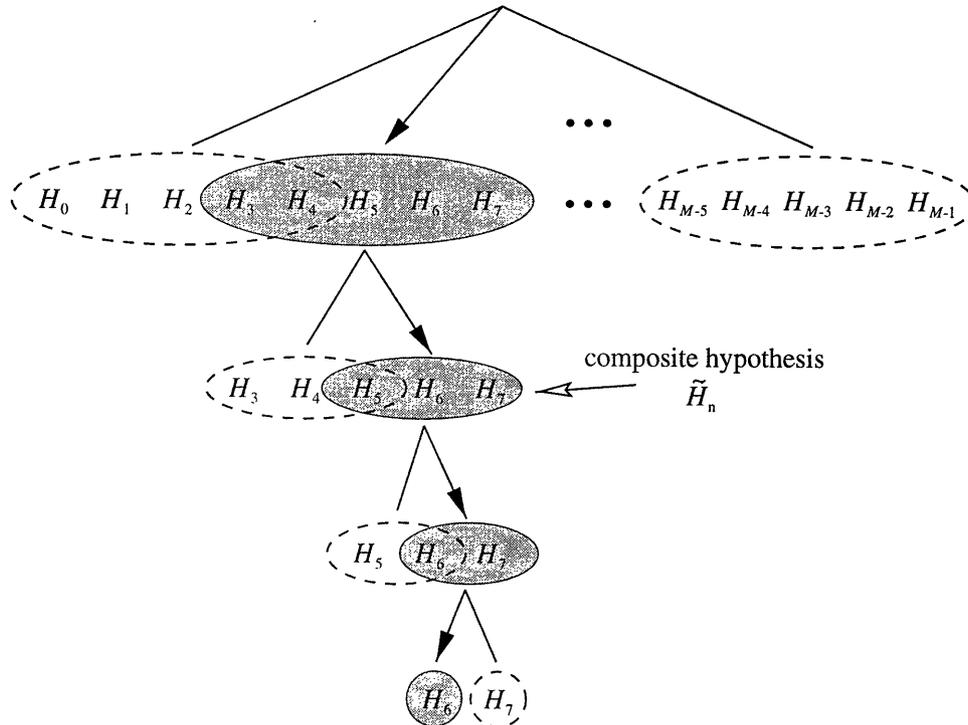


Fig. 1. A MSHT. Composite hypotheses indicated by dashed lines are discarded, shaded ones retained.

the manner we have described. The first is: how should a large hypothesis space be divided hierarchically so that hypothesis zooming can proceed? For anomaly detection in imaging (e.g., tomographic) problems there is a natural choice—grouping hypotheses spatially, i.e., subdividing the region of interest at a succession of resolutions so that anomalies can be localized to finer and finer spatial scales. This is not an original idea, and, in particular, it has been used in [8, 9] for the detection and localization of anomalies in inverse scattering data and is standard in the image segmentation community [1, 6, 7]. We also adopt this method for dividing our hypothesis space by defining composite hypothesis \tilde{H}_n as containing all hypotheses, H_m , which correspond to anomalies with support within a region of the image domain which we denote \tilde{R}_n . Figure 2 illustrates this at the coarsest scale at which the regions \tilde{R}_n are $s_{\text{hyp}} \times s_{\text{hyp}}$ in size. Composite hypotheses at finer scales have smaller cardinality and, in the anomaly detection and localization problems, naturally correspond to regions of smaller area so that hypothesis zooming and spatial zooming coincide.

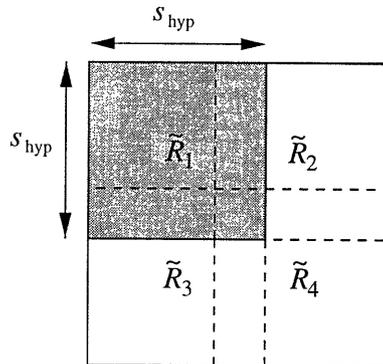


Fig. 2. Coarse-scale subdivisions. Subdivision \tilde{R}_n corresponds to \tilde{H}_n . \tilde{R}_1 is shaded and with a solid border.

The second important question is: once composite hypotheses are defined, how should the data be processed at each stage in the hierarchy in order to determine which composite hypotheses to discard? We need to address this question in light of the fact that computational complexity is a key concern. In particular, if computational load were not an issue, then we could solve the problem optimally using a generalized likelihood ratio test (GLRT) [13] as follows. At a coarse level in the hierarchy, we would enumerate all hypotheses, H_m , process the data to choose among these optimally, and then choose the composite hypothesis, \tilde{H}_n , that contains this optimal choice. Since an underlying assumption for problems of interest here is that complete enumeration and comparison is unacceptably complex, we limit our consideration to the class of decision rules that have acceptable computational complexity. By no means do we provide a complete solution to the problem of finding the best decision rule under a complexity constraint since ranking the relative complexities of different decision rules is difficult and the characterization of all decision rules with complexity less than some specified level is both ill-posed and prohibitively complex. What we do provide, however, is a solution within a specific class of decision rules that includes what can be thought of as the “natural” choices. In the process, we highlight issues that are of importance in any composite hypothesis testing problem.

For anomaly detection and localization in imaging applications there is an obvious first choice for

the decision statistic and decision rule to be used at a particular stage in the hypothesis zooming procedure. Specifically, for each composite hypothesis, \tilde{H}_n , at that stage, compute a likelihood statistic for the hypothesis that an anomaly has support over the *entire* region \tilde{R}_n . The decision rule is simple: keep the composite hypothesis \tilde{H}_n corresponding to the largest likelihood statistic.²

Approaches of this type have been proposed and used in the literature [8–11]. A principle objective of this paper is to demonstrate that we can achieve much better performance with the same level of complexity by choosing the decision statistics in a significantly different manner. To understand this, note that the statistic mentioned in the preceding paragraph corresponds to a hypothesis that, in some sense, “covers” all of the hypotheses that comprise a composite hypothesis \tilde{H}_n . Thus, this statistic sacrifices sensitivity to each individual hypothesis in \tilde{H}_n in order to achieve *some* sensitivity to *all* of the hypothesis in \tilde{H}_n . In contrast, the statistics used in the computationally intractable but optimal GLRT have, as we shall see, *significantly* greater sensitivity to each of hypotheses in a composite hypothesis. Our aim is to find statistics which approach the sensitivity of those used in the GLRT but with the same computational simplicity as that required for the “natural” decision rule described previously.

In this paper we limit ourselves to a class of decision rules based on the computation of a single linear function of the measured data for each composite hypothesis. However, we design these linear statistics to maximize their utility for the job for which they will be used. In particular, since each such statistic will be used to decide whether or not to discard its associated composite hypothesis, we aim to design a statistic that maintains strong sensitivity to the hypotheses that comprise the corresponding composite hypothesis and has minimal sensitivity to all other hypotheses. Using this philosophy, we are led to an interesting optimization problem, the solution of which is a decision

²As an aside we note that for problems in which the data are in a different domain from the image (e.g., tomographic problems), the calculation of the likelihood statistic is most easily done directly in the data domain rather than by first forming an image and then calculating the statistic [12].

statistic that is quite different from and superior to the natural choices. In particular, these optimized statistics do a far better job of increasing separation between composite hypotheses.

In this paper we develop the MSHT methodology and optimized statistic design for the tomographic anomaly detection and localization problems and also present results to quantify how well our methods work. Of particular interest to us are quantifying (a) how much performance is sacrificed by using a sub-optimal method as compared to that achievable if exhaustive hypothesis enumeration and comparison were computationally feasible; and (b) how much better our approach is compared to one based on the natural choice of decision statistics. In addition, as we have stated, one of our objectives in this paper is to lay out what we believe to be the key issues in large hypothesis testing problems more generally, and in the conclusion to this paper we provide a discussion of some of the issues and questions which we believe our work brings into sharper focus.

This paper is organized as follows. In Section II we outline all modeling assumptions and set up the problems considered in this paper. We discuss two types of decision statistics in Section III—a conventional likelihood statistic and our optimized statistic. Section IV includes pseudo-code for our MSHT algorithm and considers the computational complexity of this and the optimal approach. Examples are provided in Section V. Closing remarks and a discussion of some complements and extensions are found in Section VI.

II. PROBLEM FORMULATION

A. Observation Equation

We model tomographic data collection with the equation

$$g_i = \int_{\mathbb{R}^2} f(x, y) S_i(x, y) dx dy + \nu_i \quad (1)$$

where g_i is the i -th datum, $f(x, y)$ is a real function of two spatial variables representing the object, ν_i is the i -th sample measurement noise. The function $S_i(x, y)$ is one over the i -th strip and zero

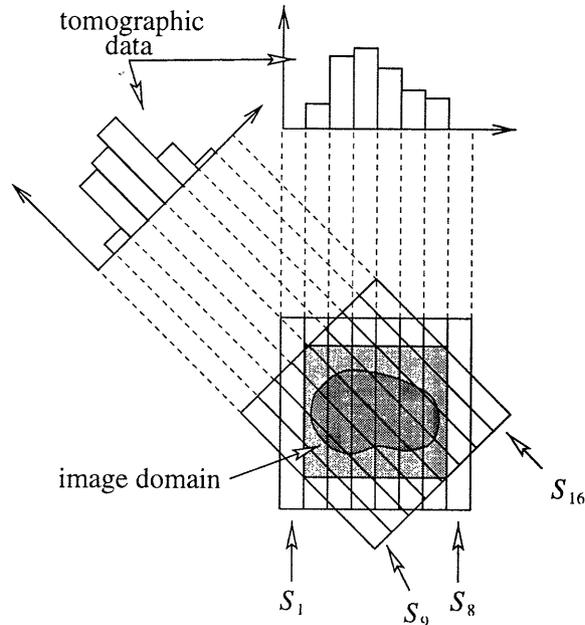


Fig. 3. Tomographic projections with strip indicator functions $S_i(x, y)$.

elsewhere. Data acquisition with these functions is illustrated in Figure 3. We let N_ϕ be the number of projections (angular positions) and we assume these positions are equally spaced in the interval $[0, \pi)$. The number of equally-spaced samples per projection angle is N_s . Finally, for computational purposes, the object, $f(x, y)$, is discretized in a rectangular pixel basis so that

$$f(x, y) = \sum_{j=1}^{N_p} f_j p_j(x, y) \quad (2)$$

where $p_j(x, y)$ is one over the j -th pixel and zero elsewhere and there are N_p pixels corresponding to a $N \times N$ field.

Combining equations (1) and (2) we find that

$$\mathbf{g} = \mathbf{T}\mathbf{f} + \boldsymbol{\nu} \quad (3)$$

where \mathbf{g} , \mathbf{f} , and $\boldsymbol{\nu}$ are vectors containing the data values, field pixel values, and noise values, respectively. The field vector, \mathbf{f} , is lexicographically ordered and \mathbf{g} and $\boldsymbol{\nu}$ are formed by stacking the values obtained at each projection. For example, $\mathbf{g} = [\mathbf{g}_1^T \ \mathbf{g}_2^T \ \dots \ \mathbf{g}_{N_\phi}^T]^T$ where \mathbf{g}_k contains the data values at the k -th projection angle. The components of the matrix \mathbf{T} are given by,

$$[\mathbf{T}]_{ij} = \int_{\mathbb{R}^2} S_i(x, y) p_j(x, y) dx dy, \quad i = 1, \dots, N_\phi N_s, \quad j = 1, \dots, N_p. \quad (4)$$

Equation (3), coupled with whatever *a priori* knowledge we have about ν and \mathbf{f} , represents our observational model. The tomographic projection matrix, \mathbf{T} , captures a discrete representation of the strip integrals. The application of \mathbf{T} to \mathbf{f} is called the projection of \mathbf{f} .

B. Models

Our notion of an anomaly is a localized region of the image domain which differs statistically from our prior set of expectations of the image. Therefore, we statistically characterize an anomaly-free background and define a parameterized class of anomalies. We model the field, \mathbf{f} , as a superposition of a field, \mathbf{f}_a , which contains, at most, a single anomaly from this class and an anomaly-free background field, \mathbf{f}_b . That is,

$$\mathbf{f} = \mathbf{f}_a + \mathbf{f}_b. \quad (5)$$

The anomaly field and background field are statistically independent.

The class of anomalies we consider is parameterized by an intensity, c , a size, s , and a position, (i, j) as follows. The anomaly field, \mathbf{f}_a , is zero everywhere except over a square patch where it is constant. Our notation is

$$\mathbf{f}_a = c\mathbf{b}(s, i, j) \quad (6)$$

where $c \geq 0$ and $\mathbf{b}(s, i, j)$ is the lexicographically ordered indicator vector associated with an $N \times N$ field which is zero everywhere except over the $s \times s$ support area with upper left corner at pixel (i, j) where $\mathbf{b}(s, i, j)$ takes the value one. The size, s , and location, (i, j) , of are unknown. We assume knowledge, however, of the minimum possible size, s_{\min} , and the maximum possible size, s_{\max} , the anomaly can be where $1 \leq s_{\min} \leq s \leq s_{\max} \ll N$.

The background field, \mathbf{f}_b , is a zero-mean, wide sense stationary (WSS), Gaussian random field

with known covariance Λ . We consider two types of background covariance statistics in our examples in Section V—white and fractal. The fractal background has a power spectral density of the form $1/\omega^2$ where ω is frequency. Additional structural details are found in [3, 4]. We consider a fractal background because fractal fields accurately model a wide range of natural textures [14].

The additive measurement noise, ν , is assumed to be a zero-mean, white, WSS, Gaussian random vector with intensity λ and is statistically independent of the background and anomaly fields. Therefore, the data are jointly Gaussian³:

$$\mathbf{g} \sim \mathcal{N}(c\mathbf{T}\mathbf{b}(s, i, j), \Lambda_{\mathbf{g}}) \quad (7)$$

where

$$\Lambda_{\mathbf{g}} = \mathbf{T}\Lambda\mathbf{T}^T + \lambda\mathbf{I}. \quad (8)$$

Notice that, even when the background is white (Λ is diagonal) the data are correlated due to the structure of the tomographic projection matrix, \mathbf{T} .

Figure 4 illustrates an example of the kind of anomaly and background field which are considered in this paper. Projections of \mathbf{f} are also shown with and without the addition of noise. Figures (a) and (b) are views of the image domain while Figures (c) and (d) are views of the data domain.

C. Problem statement

The anomaly detection problem is to determine whether or not \mathbf{f}_a is identically zero. The anomaly localization problem is to determine the values of the size $s_{\min} \leq s \leq s_{\max}$ and location (i, j) of the anomaly if indeed one is present. The optimal solution to these problems includes one hypothesis, H_m , for each possible anomaly size, s_m , and each possible location of the anomaly's upper left hand corner, (i_m, j_m) , where

³The notation $\mathbf{x} \sim \mathcal{N}(\mathbf{m}, \mathbf{P})$ means that \mathbf{x} is a Gaussian random vector with mean \mathbf{m} and covariance \mathbf{P} .

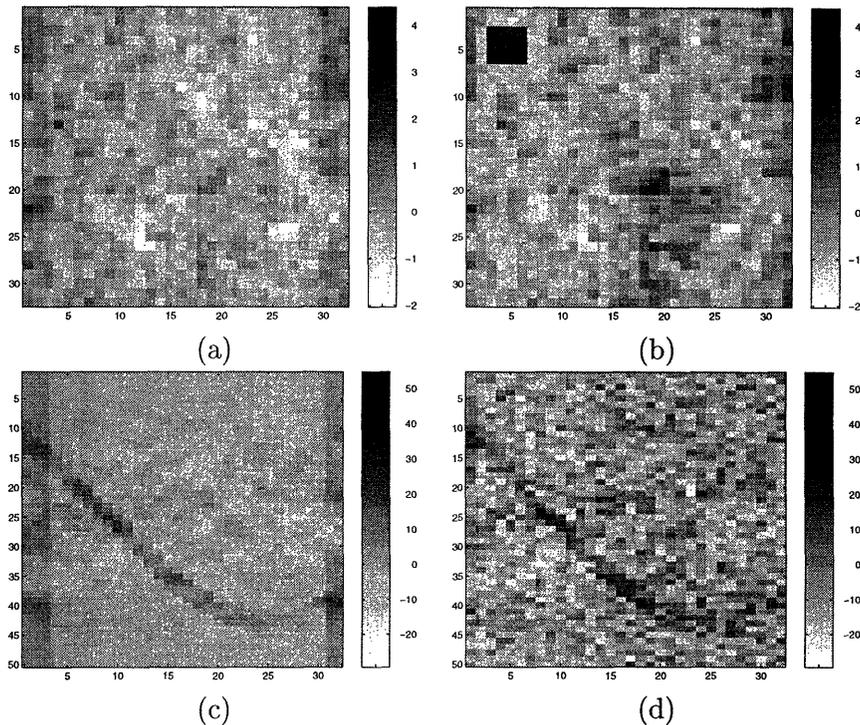


Fig. 4. Figure (a) illustrates a fractal background, \mathbf{f}_b . Figure (b) is a superposition of the background shown in (a) and an anomaly near the upper left corner. Figure (c) illustrates the projection of the anomaly plus background field. The horizontal axis is the projection number ($N_\phi = 32$). The vertical axis is the sample number ($N_s = 50$). In figure (d), measurement noise has been added to the projections.

$$s_m \in \{s_{\min}, s_{\min} + 1, \dots, s_{\max} - 1, s_{\max}\}, \quad (9)$$

$$(i_m, j_m) \in \{1, 2, \dots, N - s_m + 1\}^2. \quad (10)$$

The optimal decision statistics are likelihood ratios and the optimal decision rule is a likelihood ratio test. While the optimal test is straightforward, it is computationally infeasible for all but trivial-sized problems. Therefore, we propose the MSHT as an efficient and effective alternative.

D. Composite Hypothesis Structure

A MSHT has two main high level characteristics: the form of the composite hypotheses and the form of the statistics used to decide which composite hypotheses to discard. In this section, we define the form of the composite hypotheses. We defer discussion of decision statistics to Section III. For clarity of presentation and notational simplicity, we specify in detail only the coarsest scale

composite hypothesis test of the MSHT. The processing at other scales follows by analogy.

Figure 2 provides an interpretation of the composite hypotheses. Composite hypothesis \tilde{H}_n contains hypotheses which correspond to anomalies with support entirely within the $s_{\text{hyp}} \times s_{\text{hyp}}$ region denoted as \tilde{R}_n . For example, all hypotheses associated with anomalies with support *entirely* within the shaded region labeled \tilde{R}_1 belong to \tilde{H}_1 . We associate each composite hypothesis with an indicator function $\tilde{\mathbf{b}}_n \triangleq \mathbf{b}(s_{\text{hyp}}, i_n, j_n)$ which is one over the region \tilde{R}_n and zero elsewhere. The composite hypothesis regions overlap by at least $s_{\text{max}} - 1$ pixels so that each possible anomaly lies entirely within at least one region. This ensures that $\mathcal{H} = \tilde{H}_1 \cup \tilde{H}_2 \cup \tilde{H}_3 \cup \tilde{H}_4$.

III. DECISION STATISTICS

In this section we specify the form of the decision statistics for the coarsest scale composite hypothesis test of a MSHT. The statistics at subsequent scales are easily understood by analogy. We will discuss two types of decision statistic. The first type, discussed in Section III-A, is a coarse-scale likelihood statistic of the form used in [8, 9] for problems similar to the ones addressed here. As discussed in Section I, while this statistic is natural and intuitive, it sacrifices considerable sensitivity to achieve computational simplicity. We use this statistic as a benchmark against which to compare what we call an *optimized statistic* which we discuss in Section III-B.

A. Coarse-Scale Likelihood Statistics

As discussed in the introduction, the coarse-scale likelihood statistic associated with composite hypothesis \tilde{H}_n is the log-likelihood ratio to discriminate between a single hypothesis for a coarse-scale anomaly with support over the entire region \tilde{R}_n and the hypothesis that no anomaly exists. To derive the coarse-scale likelihood statistic, we associate with each \tilde{H}_n a coarse-scale anomaly: $\mathbf{f}_a = c\tilde{\mathbf{b}}_n$ where $\tilde{\mathbf{b}}_n$, as defined in Section II-D, is the indicator function for \tilde{R}_n .

From equation (7) and standard results in hypothesis testing [13], we find that the log-likelihood

for each of these four hypotheses is given by an affine operation on the observed data, \mathbf{g} , namely,

$$\tilde{\ell}_n(\mathbf{g}) = (\mathbf{T}\tilde{\mathbf{b}}_n)^T \Lambda_{\mathbf{g}}^{-1} \mathbf{g} - \frac{1}{2} (\mathbf{T}\tilde{\mathbf{b}}_n)^T \Lambda_{\mathbf{g}}^{-1} \mathbf{T}\tilde{\mathbf{b}}_n, \quad n = 1, 2, 3, 4 \quad (11)$$

and the resulting decision rule consists of choosing the largest of these four values.

Note that at this coarsest scale in the MSHT, we have that

$$(\mathbf{T}\tilde{\mathbf{b}}_i)^T \Lambda_{\mathbf{g}}^{-1} \mathbf{T}\tilde{\mathbf{b}}_i = (\mathbf{T}\tilde{\mathbf{b}}_j)^T \Lambda_{\mathbf{g}}^{-1} \mathbf{T}\tilde{\mathbf{b}}_j, \quad i, j = 1, 2, 3, 4. \quad (12)$$

This follows from the symmetry of the composite hypothesis regions, the fact that we have a complete set of data, and the wide sense stationarity of \mathbf{g} . Consequently, at this level we can drop the second term in equation (11) as it has no influence on the decision rule. Note, however, that relations such as equation (12) need not hold at subsequent scales since the composite hypothesis regions do not have the requisite symmetry. Despite this, we have found that such relations are good approximations at all scales.

B. Optimized Statistics

In this section we design a statistic which is no more computationally complex than the coarse-scale likelihood ratio but is much more sensitive and discriminating. The statistic we design is linear in the observed data and has the form

$$\tilde{\ell}_i(\mathbf{g}) = \mathbf{a}_i \mathbf{g}, \quad i = 1, 2, 3, 4. \quad (13)$$

Roughly, our objective in designing such a statistic is to choose the vector \mathbf{a}_i to force $\tilde{\ell}_i$ to be significantly larger, on average, when \tilde{H}_i is true than when it is false. That is, we would like $\tilde{\ell}_i$ to have a large separation between its mean when any $H_j \in \tilde{H}_i$ is true and its mean when any $H_k \notin \tilde{H}_i$ is true. Since doubling the magnitude of \mathbf{a}_i will double this difference in mean values, we normalize this difference by the standard deviation of $\tilde{\ell}_i$.

More precisely, we define

$$m_{ij} \triangleq E[\tilde{\ell}_i | H_j] = \mathbf{c}\mathbf{a}_i^T \mathbf{T}\mathbf{b}_j \quad (14)$$

$$\sigma_i^2 \triangleq \text{var}[\tilde{\ell}_i | H_j] = \mathbf{a}_i^T \mathbf{\Lambda}_g \mathbf{a}_i \quad (15)$$

where we have introduced the shorthand notation $\mathbf{b}_j \triangleq \mathbf{b}(s_j, i_j, k_j)$ for the anomaly indicator function associated with hypothesis H_j . Notice that the conditional mean is linear in \mathbf{a}_i while the conditional variance is quadratic in \mathbf{a}_i . Also note that the conditional variance is independent of j .

The criterion we adopt is to maximize the worst-case (i.e., smallest) normalized difference between m_{ij} for $H_j \in \tilde{H}_i$ and m_{ik} for $H_k \notin \tilde{H}_i$ where σ_i is used for normalization. That is, we choose $\hat{\mathbf{a}}_i$ as the solution to the optimization problem

$$\hat{\mathbf{a}}_i = \arg \max_{\mathbf{a}} \min_{(j,k) \in \mathcal{A}_i} \frac{m_{ij} - m_{ik}}{\sigma_i} \quad (16)$$

where $\mathcal{A}_i \triangleq \{(j,k) | H_j \in \tilde{H}_i \text{ and } H_k \notin \tilde{H}_i\}$. Substituting in the definitions of equations (14) and (15), the optimization problem reduces to

$$\hat{\mathbf{a}}_i = \arg \max_{\mathbf{a}} \min_{(j,k) \in \mathcal{A}_i} \frac{\mathbf{a}^T \mathbf{T}\mathbf{b}_j - \mathbf{a}^T \mathbf{T}\mathbf{b}_k}{\sqrt{\mathbf{a}^T \mathbf{\Lambda}_g \mathbf{a}}}. \quad (17)$$

As is shown in Appendix A, using Lagrange duality theory this problem can be reformulated as the quadratic program

$$\hat{\mathbf{y}} = \arg \min_{\mathbf{y}} \mathbf{y}^T \mathbf{Q} \mathbf{\Lambda}_g^{-1} \mathbf{Q}^T \mathbf{y} \quad (18)$$

$$\text{subject to } \begin{cases} \mathbf{e}^T \mathbf{y} = 1 \\ \mathbf{y} \geq 0 \end{cases} \quad (19)$$

where

$$\mathbf{Q} \triangleq \begin{bmatrix} (\mathbf{T}\mathbf{b}_{j_1} - \mathbf{T}\mathbf{b}_{k_1})^T \\ (\mathbf{T}\mathbf{b}_{j_2} - \mathbf{T}\mathbf{b}_{k_2})^T \\ \vdots \end{bmatrix} \quad \text{and } \mathbf{e} \triangleq [1 \ 1 \ 1 \ \dots \ 1]^T \quad (20)$$

and all pairs $(j_n, k_n) \in \mathcal{A}_i$. This quadratic program may be solved using, for example, methods described in [5]. Since the solution depends only on the data covariance matrix and the structure

of the hypothesis space, but not on the data itself, solving the quadratic program is an off-line procedure. Once $\hat{\mathbf{y}}$ is found, $\hat{\mathbf{a}}_i$ is given by

$$\hat{\mathbf{a}}_i = \frac{\Lambda_{\mathbf{g}}^{-1} \mathbf{Q}^T \hat{\mathbf{y}}}{\sqrt{\hat{\mathbf{y}}^T \mathbf{Q} \Lambda_{\mathbf{g}}^{-1} \mathbf{Q}^T \hat{\mathbf{y}}}}. \quad (21)$$

Note that in optimizing over all linear functions of the tomographic data we implicitly consider the coarse-scale likelihood statistic as well as all linear functions of any linear reconstruction of the image. Therefore, pre-processing the data with any linear reconstruction routine (like convolution back-projection) cannot result in a better statistic.

IV. MSHT ALGORITHM AND COMPUTATIONAL COMPLEXITY

A. The MSHT Algorithm

Given the hierarchical hypothesis space decomposition described in Section II-D and either choice of decision statistic discussed in Section III, the MSHT algorithm detects and localizes an anomaly in the coarse-to-fine manner described in Section I. At each scale, other than the finest, one of four subdivisions of the remaining part of the image domain is selected for finer-scale investigation. The scale-recursion terminates at a scale at which the optimal test is feasible and at which the regions \tilde{R}_n are no smaller than $s_{\max} \times s_{\max}$. Then, having localized the anomaly to an area significantly smaller than the entire image domain, the optimal test is performed which includes only hypotheses associated with anomalies in that area. Finally, the statistic value associated with the selected hypothesis is compared with a threshold. If it is larger than the threshold, the associated hypothesis is selected, otherwise it is declared that no anomaly exists.

The following pseudo-code summarizes the algorithm. The inputs to the algorithm are the region to investigate (initialized to the entire image domain) and the scale number (initialized to one). The output is a hypothesis in \mathcal{H} or the decision that no anomaly exists.

MSHT Pseudo-Code

$$\hat{H} = MSHT(\tilde{\ell}, \tilde{R}, scale)$$

Step 1. *If scale is finest possible, perform optimal test on anomalies in \tilde{R} . Call the selected hypothesis H and the associated statistic value ℓ .*

$$\hat{H} = \begin{cases} H & \text{if } \ell > \eta \\ \text{"no anomaly"} & \text{otherwise} \end{cases} .$$

Step 2. *Otherwise, subdivide the region \tilde{R} into four overlapping squares where the amount of overlap is at least $s_{\max} - 1$. Denote these squares \tilde{R}_n for $n = 1, 2, 3, 4$.*

Step 3. *For each subdivision \tilde{R}_n compute the statistic $\tilde{\ell}_n$.*

Step 4. *Let k be such that $\tilde{\ell}_k > \tilde{\ell}_n, \forall k \neq n$. Call $\hat{H} = MSHT(\tilde{\ell}_k, \tilde{R}_k, scale + 1)$.*

B. Computational Complexity

Our primary motivation for applying a MSHT to the anomaly detection and localization problems is that the optimal hypothesis test is too computationally costly. The MSHT formulates fewer hypotheses than the optimal test and is therefore more efficient. In this section we quantify this claim by calculating the computational complexity of the optimal algorithm and the MSHT algorithm. To do so, we will compute the number of operations per hypothesis and the number of hypotheses formulated in each algorithm.

Both the optimal test and the MSHT formulate linear statistics which require $O(N_\phi N_s)$ operations (adds and multiplies). This result follows from the fact that a linear statistic is an inner product between two length $N_\phi N_s$ vectors. Since each hypothesis requires a constant amount of work, the overall complexity of either algorithm scales with the number of hypotheses formulated. Hence we take the number of hypotheses formulated as a measure of algorithmic complexity.

First consider the optimal test. Suppose the linear size of the square field is N and that we know the minimum and maximum possible size of the anomaly: s_{\min} and s_{\max} , respectively. The number of hypotheses formulated is $M_{\text{opt}} = \sum_{r=s_{\min}}^{s_{\max}} (N - r + 1)^2 = O(N^2)$.

Now consider the MSHT algorithm. Its computational complexity is a function of, among other parameters, the scale at which the hypothesis zoom terminates and exhaustive enumeration of all

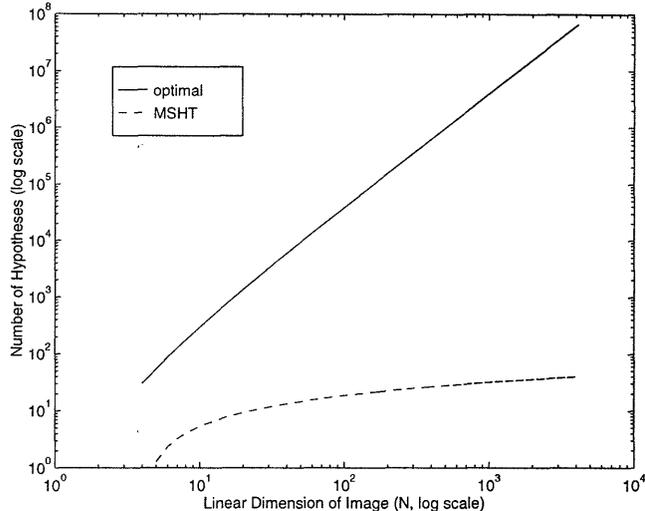


Fig. 5. The complexity of the optimal algorithm (top curve), and a MSHT (bottom curve). The vertical-axis is the number of hypotheses (log scale), the horizontal-axis is the linear dimension of the field (log scale). Here $s_{\max} = 4$ and $s_{\min} = 1$.

remaining hypotheses is conducted (Step 1 of the above pseudo-code) and the amount by which the regions \tilde{R}_n overlap. We may neglect the effect of overlapping in our order-of-magnitude calculation since the overlapping is on the order of s_{\max} and $s_{\max} \ll N$. Therefore, the number of hypotheses formulated is $M_{\text{MSHT}} = O(\log N) + C$ where the constant, C , accounts for the number of hypotheses formulated in the finest-scale exhaustive enumeration step. If the finest-scale exhaustive enumeration step is conducted at a scale at which the \tilde{R}_n are on the order of $s_{\max} \times s_{\max}$ then C is negligible since $s_{\max} \ll N$.

Figure 5 displays the number of hypotheses for the optimal and MSHT algorithms as a function of N for $s_{\min} = 1$ and $s_{\max} = 4$. For the case illustrated, it is assumed that the exhaustive enumeration step of the MSHT is conducted at a scale at which the \tilde{R}_n are $s_{\max} \times s_{\max}$. As can be seen, the difference in the number of hypotheses considered is quite large. For example, for a 512×512 image (i.e., $N = 512$), the MSHT considers about 50 hypotheses while the optimal algorithm considers just over 10^6 hypotheses.

V. EXAMPLES

In this section we present several types of examples. First, in Section V-A, we introduce a means of directly comparing the sensitivity of a coarse-scale likelihood statistic with that of an optimized statistic. In Section V-B we investigate the performance of the first, coarsest-stage of a MSHT algorithm in several cases of differing data quality for problems of sufficiently small size so that comparison to the optimal test is feasible. We focus on the coarsest scale because it is the scale at which the composite hypotheses are largest in cardinality and, therefore, most difficult to disambiguate [9]. That is, it is at this stage that the MSHT should have its greatest difficulty and thus is most in need of maximally sensitive statistics.

In Section V-C, we conclude with examples illustrating the performance of a “full” MSHT algorithm (one which continues past the coarsest scale and successively localizes the anomaly to the finest possible scale). This full algorithm includes three scales, the two coarsest scales are composite hypothesis tests. By the third scale, the anomaly has been sufficiently localized so that full enumeration and comparison of the remaining hypotheses is feasible. Hence, an optimal test on the remaining hypotheses is performed at the third scale of the full algorithm.

The data in all our examples are simulated based on the models presented in Section II-B and we include examples corresponding to two different field sizes. The examples in Sections V-A and V-B correspond to a 16×16 field while the examples in Section V-C correspond to a 32×32 field. The parameters of the problem addressed in each section are detailed in Table I. Throughout this section, a “false alarm” occurs when no anomaly exists but a “no anomaly” decision is not returned by the algorithm. A “detection” occurs when an anomaly exists and the true composite hypothesis (Section V-B) or hypothesis (Section V-C) is selected.

Before proceeding to examples, we present the definitions of *signal-to-noise ratio* (SNR) and *anomaly-to-background ratio* (ABR):

<i>Section</i>	<i>Field Size</i>	<i>Background Field</i>	N_ϕ	N_s	s_{\min}	s_{\max}	$s_{\text{hyp}} \text{ scale } 1$	$s_{\text{hyp}} \text{ scale } 2$
V-A	16×16	fractal & white	16	24	4	4	10	N/A
V-B	16×16	fractal	16	24	3	4	10	N/A
V-C	32×32	fractal	32	50	4	4	18	11

TABLE I
DETAILS RELATING TO THE EXAMPLES PRESENTED IN SECTION V.

$$\text{SNR} \triangleq \frac{\text{trace}(\mathbf{T}\mathbf{\Lambda}\mathbf{T}^T)}{\text{trace}(\lambda\mathbf{I})} \quad (22)$$

$$\text{ABR} \triangleq \frac{\mathbf{f}_a^T \mathbf{f}_a}{\text{trace}(\mathbf{\Lambda})}. \quad (23)$$

The SNR measures the relative power between the projected background and the additive noise, while the ABR measures the relative power between the anomaly and background fields.

A. Direct Comparison of Statistics

In this section, we illustrate in a direct way the superiority of the optimized statistics over the coarse-scale likelihood statistics. To make the comparison, recall the definitions of m_{ij} and σ_i^2 . The former is the mean value of statistic $\tilde{\ell}_i$ (associated with \tilde{R}_i which is depicted in Figure 2) conditioned on hypothesis H_j . The latter is the associated variance. We define the standard-deviation-normalized conditional mean as

$$\mu_{ij} \triangleq \frac{m_{ij}}{\sigma_i}. \quad (24)$$

The value of μ_{ij} indicates how sensitive $\tilde{\ell}_i$ is to hypothesis H_j .

Figures 6 and 7 illustrate values of μ_{1j} for the case of a white and fractal background respectively at $\text{SNR} = 1$. The two-dimensional bar at position (m, n) in the plots of these figures corresponds to the hypothesis that $\mathbf{f}_a = \mathbf{b}(4, m, n)$. The shaded regions of these two-dimensional bar charts are the areas we wish to be large; that is, these correspond to values of μ_{1j} for hypotheses, H_j ,

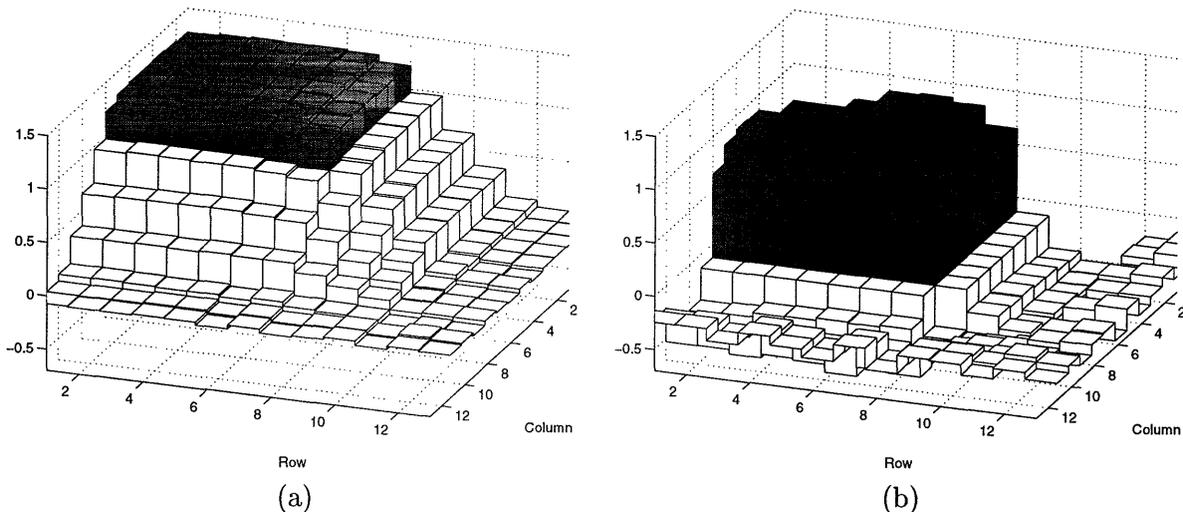


Fig. 6. μ_{1j} with a white background at SNR=1. Figure (a) is for the coarse-scale likelihood statistic and figure (b) is for the optimized statistic. Pixel (m, n) corresponds to $\mathbf{f}_a = \mathbf{b}(4, m, n)$.

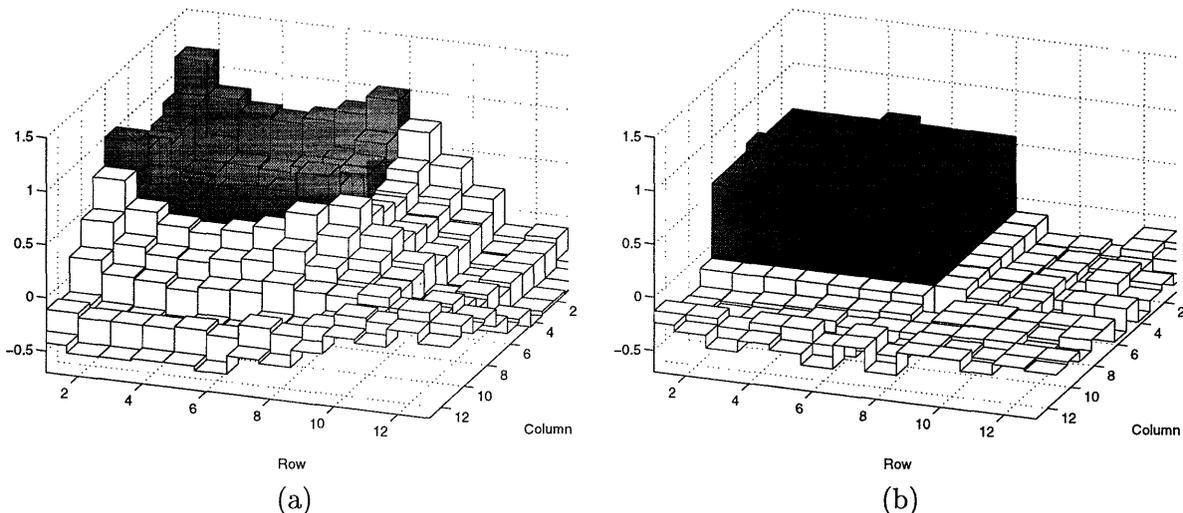


Fig. 7. μ_{1j} with a fractal background at SNR=1. Figure (a) is for the coarse-scale likelihood statistic and figure (b) is for the optimized statistic. Pixel (m, n) corresponds to $\mathbf{f}_a = \mathbf{b}(4, m, n)$.

that comprise the composite hypothesis \tilde{H}_1 . The unshaded portions of these plots correspond to $H_j \notin \tilde{H}_1$, i.e., these are values of μ_{1j} we would like to be significantly smaller. In both figures, plot (a) corresponds to the coarse-scale likelihood statistic and plot (b) to the optimized statistic.

The shape of both plots in Figure 6 exhibit precisely the type of behavior we want. The value of μ_{1j} is relatively high for $H_j \in \tilde{H}_1$ (shaded region) and relatively low for $H_j \notin \tilde{H}_1$ (unshaded region). Notice, however, that for the optimized statistic (Figure 6(b)) there is a more abrupt

transition between the shaded and unshaded regions as compared to the coarse-scale likelihood statistic (Figure 6(a)). Note that at this coarsest stage (and at any stage prior to the final one), our objective is to *discard* hypotheses and *not* to make an absolute decision about the presence or absence of an anomaly (this is considered at the final stage). As a consequence, it is the sharpness of the transition and the *relative* (not absolute) sizes of the μ_{1j} between the shaded and unshaded regions that are of importance. Thus, while it might appear that sensitivity is lost in Figure 6(b) because the values in the shaded regions are somewhat lower than in Figure 6(a), this is not the case. Indeed, because of the sharper transition in Figure 6(b), the optimized statistic does a significantly better job at disambiguating composite hypotheses as we will verify in the next section.

While Figure 6 shows that there is some enhancement using the optimized statistic for the case of a white background, this point becomes much stronger if we consider a correlated background field as we do in Figure 7. In particular, in Figure 7 the background has a fractal covariance structure. In the case of Figure 7(a), we see that $\tilde{\ell}_1$ is sensitive to the wrong hypotheses. That is, there exist $H_k \notin \tilde{H}_1$ and $H_j \in \tilde{H}_1$ for which $\mu_{1k} > \mu_{1j}$. Moreover, there is no clear, sharp transition between the shaded and unshaded regions in Figure 7(a). Therefore, in the case of the fractal background, the coarse-scale likelihood statistic is ineffective. Comparing Figure 7(a) with Figure 7(b) we see that the optimized statistic is significantly better than the coarse-scale likelihood statistic—it is sensitive to the correct hypotheses and is highly discriminating between those hypotheses which belong to \tilde{H}_1 and those which do not. Though we shall not provide the corresponding figures for the three other coarse-scale composite hypotheses ($\tilde{H}_2, \tilde{H}_3, \tilde{H}_4$), the same conclusions apply.

B. Coarse-Scale Performance

In this section, we compare the performance of three algorithms for coarse-scale anomaly detection and localization. One algorithm is the GLRT. The other two are similar to the coarsest scale

<i>Algorithm</i>	<i>Number of Hypotheses Formulated</i>
GLRT	365
coarse-scale MSHT with coarse-scale likelihood statistics	4
coarse-scale MSHT with optimized statistics	4

TABLE II

THE NUMBER OF HYPOTHESES FORMULATED BY THE ALGORITHMS CONSIDERED IN SECTION V-B.

of a MSHT. For these, statistics associated with four coarse-scale regions (shown in Figure 2) are computed. The largest is compared to a threshold. If it is above the threshold, its associated region is declared as containing the anomaly. Otherwise, a “no anomaly” decision is returned.

We have two objectives. The first is to contrast the performance of two coarse-scale MSHT algorithms (as just described), one using coarse-scale likelihoods as decision statistics and the other optimized statistics. In particular, we will demonstrate the significant performance enhancement obtained using optimized statistics. The second objective is to contrast the coarse-scale MSHT with optimized statistics to the optimal GLRT. To do this, of course, we must examine a problem of sufficiently small size so that the GLRT is computationally feasible.

In Figure 8 we illustrate receiver operator characteristic (ROC)⁴ curves at different ABRs for the three methods: GLRT (top curve in each plot), coarse-scale MSHT algorithm with optimized statistics (middle curve in each plot), coarse-scale MSHT algorithm with coarse-scale likelihood statistics (bottom curve in each plot). The anomaly considered is $\mathbf{f}_a = c\mathbf{b}(4, 6, 6)$ where c varies with the ABR. The background is fractal and $\text{SNR} = 1$. Table II indicates the number of hypotheses formulated for each of the three methods considered here.

Figure 8 unquestionably illustrates the superiority of the optimized statistic over the coarse-scale likelihood statistic. At all ABRs, the coarse-scale MSHT performance using the optimized statistic outperforms that using the coarse-scale likelihood statistic by a wide margin. Indeed, in

⁴Note that an ROC curve for an M -ary hypothesis test, unlike that for a binary hypothesis test, need not go through the point $(\text{Pr}(\text{detection}), \text{Pr}(\text{false alarm})) = (1, 1)$ because more is required for a detection in the M -ary hypothesis case. It is not enough that the null hypothesis is not chosen, the *correct* alternative hypothesis must be selected.

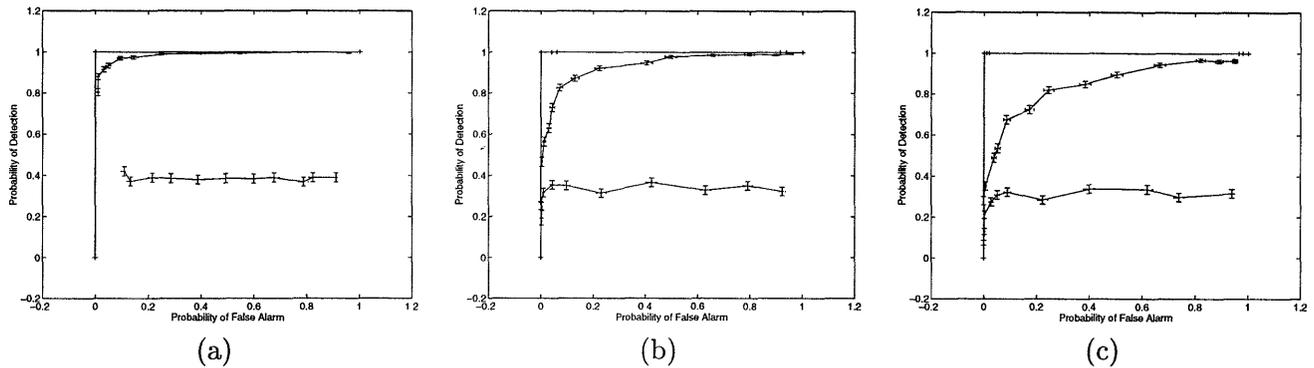


Fig. 8. Comparison of three algorithms at different ABRs. In each plot, ROCs for the GLRT (top curve), coarse-scale MSHT with optimized statistics (middle curve), coarse-scale MSHT with coarse-scale likelihood statistics (bottom curve) are shown. Five-hundred Monte Carlo runs were conducted per data point and error bars are drawn plus and minus one standard deviation. The background is fractal and SNR = 1. Figure (a): ABR = 9.1. Figure (b): ABR = 5.1. Figure (c): ABR = 3.5.

all but Figure 8(a), the ROCs for the coarse-scale MSHT using coarse-scale likelihood statistics do not stray far from a probability of detection of 0.2, the performance level which *blind guessing* achieves⁵. In fact, even in Figure 8(a), at an ABR of 9.1, the coarse-scale likelihood statistics are not much better than blind guessing. These experimental results support the analytical ones of the previous section—the coarse-scale likelihood statistic has far lower discriminating power and sensitivity compared to the optimized statistic.

Comparing the performance of the coarse-scale MSHT with optimized statistics to the optimal GLRT provides us with a quantification of the performance loss due to the use of a computationally simpler decision rule. There are two pieces of information to extract from such results. The first is that it defines a range of ABRs for which there is only minimal loss of performance as compared to the optimal test. For example, as Figure 8(a) indicates, the performance loss in using the coarse-scale MSHT with linear statistics at ABRs of about 9 or larger is quite small.

This comparison also identifies a range of ABRs over which the constraint of using a single linear statistic for each composite hypothesis, \tilde{H}_n , is too severe, resulting in too great a loss of performance relative to the optimal test. For example, the performance in Figure 8(c) and even in Figure 8(b)

⁵Recall that there are five choices: “no anomaly” and the choice that an anomaly exists in region \tilde{R}_n for $n = 1, 2, 3, 4$.

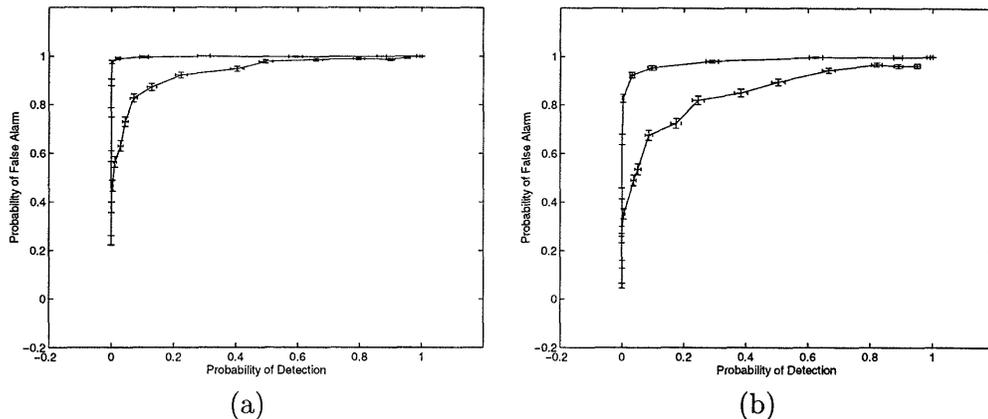


Fig. 9. ROCs for the coarse-scale MSHT. For the top curves in each plot, four linear statistics were computed per \tilde{H}_n . For the bottom curves, one linear statistic was computed. Five-hundred Monte Carlo runs were conducted per data point and error bars are drawn plus and minus one standard deviation. The background is fractal, SNR = 1. Figure (a): ABR = 5.1. Figure (b): ABR = 3.5.

suggest that at this range of ABRs we need to consider more complex decision rules to achieve near-optimal performance. The GLRT and the optimization procedure we have described suggest a method for doing this. Specifically, the GLRT involves the calculation of *many* linear statistics for each \tilde{H}_n —one statistic matched to each individual hypothesis. In contrast, the coarse-scale MSHT algorithm using optimized statistics as we have described it, uses a single linear statistic. An obvious generalization is to use *several*—say, four—optimized statistics for each \tilde{H}_n where each statistic is sensitive to a different subset of \tilde{H}_n . The resulting decision rule would then compute all 16 statistics and choose the \tilde{H}_n corresponding to the largest one. Figure 9 depicts the result of applying such a rule for the same cases shown in Figure 8(b) and (c). We see from Figure 9 that this rule, which has four times the complexity of the rule using a single statistic per \tilde{H}_n , has significantly better (and in the cases shown nearly perfect) performance.

C. Full Algorithm Performance

In this section we illustrate an example for a full MSHT algorithm. In contrast to the problems considered in the previous section, the one considered here is larger⁶. The field is 32×32 and the

⁶This larger size prohibits the computation of Monte Carlo results based on the optimal test.

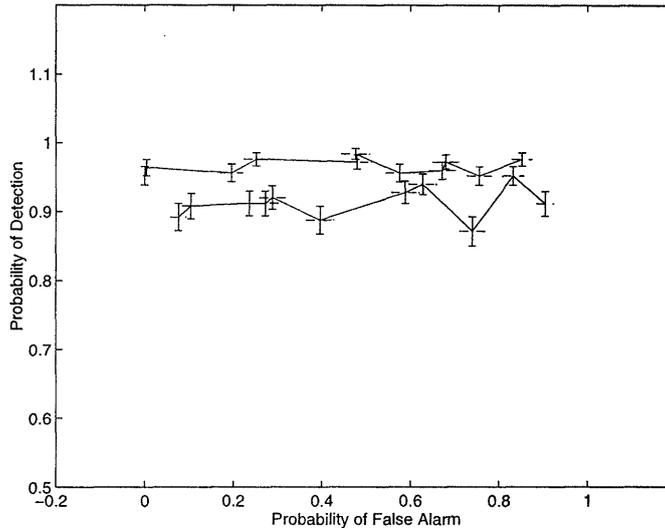


Fig. 10. ROCs for the MSHT at two ABRs. Top: ABR = 5.1. Bottom: ABR = 3.5. SNR = 1, $\mathbf{f}_a = \mathbf{cb}(4, 3, 3)$ where c varies with the ABR. The background is fractal, the number of Monte Carlo runs is 250 and error bars are drawn plus and minus one standard deviation.

anomaly considered is $\mathbf{f}_a = \mathbf{cb}(4, 3, 3)$ where c varies with the ABR. One optimized statistic is used for each composite hypothesis in the tests at the first two scales of the algorithm. After the second scale, the anomaly has been localized to a 11×11 region which is small enough that the optimal test can be performed over the hypotheses corresponding to anomalies with support in that region. Further details are provided in Table I. The background is fractal with SNR = 1.

Figure 10 illustrates ROCs at two ABRs. The top curve corresponds to ABR = 5.1 and the bottom to ABR = 3.5. At these ABRs we see that performance is quite good, indicating that the MSHT is indeed an effective way to navigate a large hypothesis space to find the true hypothesis. At lower ABRs, however, where performance is significantly below the level we see here, additional statistics per composite hypothesis can be used to increase performance levels, as discussed in the previous section.

VI. DISCUSSION AND CONCLUSION

We have presented the multiscale hypothesis test with optimized statistics as a new approach to the anomaly detection and localization problems from noisy tomographic projections. We have

shown that, in certain data-quality regimes, this hierarchical, hypothesis zooming method can achieve good performance with great efficiency. The key to achieving high performance with low computational complexity is the design of highly selective statistics. We developed and solved an optimization problem for such statistics and, in several ways, quantified their superiority over conventional statistics.

While we have developed the MSHT framework in the context of anomaly detection and localization from tomographic data, we have touched on the fundamental issues relevant to a broad class of problems—those involving large hypothesis spaces. The key obstacles in dealing with large hypothesis spaces include how to organize the space for hypothesis zooming and how to process the data for efficient decision-making. The MSHT framework as applied to the anomaly detection and localization problems addresses these challenges and provides a guiding philosophy for solutions to similar large-hypothesis-space problems.

In our development of the MSHT framework we have imposed certain conditions, the relaxation of which suggest ways to achieve additional performance gain using the methodology we have described. For example, as we described in Section V-B, it is possible to generalize our approach by designing several statistics per composite hypothesis. With this extension, our approach provides a set of MSHT algorithms ranging from the simplest (using one statistic per composite hypothesis), to the fully optimal GLRT (using as many statistics as there are individual hypotheses). This then provides a systematic framework for identifying the minimally complex MSHT as a function of data quality and performance level.

Further extensions of our approach are also possible. For example, in defining the MSHT we have held the hierarchical hypothesis space decomposition fixed. For the problems of interest in this paper, the choice we have made (namely, defining \tilde{H}_n to correspond to anomalies in region \tilde{R}_n) is natural. But, in more general problems, hypothesis space decomposition needs to be considered

jointly with the design of decision statistics. In general, this is a prohibitively complex combinatorial problem. However, we expect that our framework can lead to feasible iterative algorithms for joint statistic design and hypothesis space decomposition. An iterative procedure might start with an initial composite hypothesis (call it \tilde{H}) and its associated optimized statistic. Then, using the type of information presented in Figures 6 and 7, it can be determined which hypotheses are not placed in the correct composite hypothesis and ought to be redistributed. For example, the next suggested composite hypothesis would not include hypotheses which are in \tilde{H} but have a relatively low average statistic value and would include hypotheses which are not in \tilde{H} but have a relatively large average statistic value. Iterating this process—i.e., redesigning the statistic to discriminate this new group of hypotheses from all others, identifying those that seem misplaced and redistributing them—provides a systematic procedure for using our methodology to build composite hypotheses that are maximally distinguished by an algorithm using optimized statistics.

APPENDIX

I. QUADRATIC PROGRAMMING FORMULATION FOR OPTIMIZED STATISTICS

In this appendix we show how to formulate the optimization problem for statistics posed in Section III-B as a quadratic programming problem. To do so we shall employ Lagrange duality theory which is a standard technique for recasting constrained optimization problems [2]. The optimization problem we consider in Section III-B is

$$\hat{\mathbf{a}}_i = \arg \max_{\mathbf{a}} \min_{(j,k) \in \mathcal{A}_i} \frac{\mathbf{a}^T \mathbf{T} \mathbf{b}_j - \mathbf{a}^T \mathbf{T} \mathbf{b}_k}{\sqrt{\mathbf{a}^T \boldsymbol{\Lambda}_g \mathbf{a}}}. \quad (25)$$

It is sufficient to consider vectors \mathbf{a} for which $\mathbf{a}^T \boldsymbol{\Lambda}_g \mathbf{a} = 1$. Making a few additional notational changes, we rewrite the problem as

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} z \quad (26)$$

$$\text{subject to } \begin{cases} \mathbf{P}\mathbf{x} - z\mathbf{e} \geq 0 \\ 1 - \mathbf{x}^T\mathbf{x} \geq 0 \end{cases} \quad (27)$$

where

$$\mathbf{x} \triangleq \Lambda_{\mathbf{g}}^{1/2}\mathbf{a}, \quad \mathbf{P} \triangleq \mathbf{Q}\Lambda_{\mathbf{g}}^{-1/2} \quad (28)$$

and \mathbf{Q} and \mathbf{e} are as defined in Section III-B. This is the *primal problem*.

Let us call the optimal cost to the primal problem \hat{z} . For simplicity, we assume that $\hat{\mathbf{x}}$ exists and that $\hat{z} < \infty$. Introducing Lagrange multipliers \mathbf{y} and μ , we define the Lagrangian cost function as

$$L(z, \mathbf{x}, \mu, \mathbf{y}) \triangleq z + \mathbf{y}^T(\mathbf{P}\mathbf{x} - z\mathbf{e}) + \mu(1 - \mathbf{x}^T\mathbf{x}). \quad (29)$$

Our aim is to find values for the Lagrange multipliers such that maximizing L is the same as solving the primal problem. Toward this end we define $J(\mu, \mathbf{y}) \triangleq \max_{z, \mathbf{x}} L(z, \mathbf{x}, \mu, \mathbf{y})$. The function J is the maximum of the Lagrangian cost as a function of the Lagrange multipliers. It is straightforward to show that in searching for the $\hat{\mathbf{y}}$ and $\hat{\mu}$ we must consider only non-negative values. It is also clear that the weak duality relationship, $J(\mu, \mathbf{y}) \geq \hat{z}$, holds for all values of \mathbf{y} and μ . The dual problem attempts to find the smallest upper bound. In our case it is not hard to show strong duality, i.e., that the smallest upper bound is, in fact, tight ($\min_{\mu, \mathbf{y}} J(\mu, \mathbf{y}) = \hat{z}$).

The *dual problem*, therefore, is

$$[\hat{\mu} \hat{\mathbf{y}}^T] = \arg \min_{\mu, \mathbf{y}} J(\mu, \mathbf{y}) \quad (30)$$

$$\text{subject to } \begin{cases} \mu \geq 0 \\ \mathbf{y} \geq 0 \end{cases} \quad (31)$$

All that remains is to put the dual problem into a more useful form. To begin doing so, recall that J is the maximum of L over all z and \mathbf{x} . A necessary condition at the maximum of L is that the gradient of L is zero. Setting the partial derivative of L with respect to z and the gradient of L with respect to \mathbf{x} to zero yields the conditions $\mathbf{y}^T\mathbf{e} = 1$ and $\hat{\mathbf{x}} = \frac{1}{2\mu}\mathbf{P}^T\mathbf{y}$. Plugging these conditions

back into L yields $J(\mu, \mathbf{y}) = \mu + \frac{\mathbf{y}^T \mathbf{P} \mathbf{P}^T \mathbf{y}}{4\mu}$.

Having found a workable expression for J , the dual problem is to minimize it. Using the fact that μ cannot be zero (or else $\hat{\mathbf{x}}$ is unbounded) a necessary condition for μ at the minimum is $\hat{\mu} = \frac{1}{2} \sqrt{\mathbf{y}^T \mathbf{P} \mathbf{P}^T \mathbf{y}}$. Having found the optimal μ we plug this into J to get $J(\hat{\mu}, \mathbf{y}) = \sqrt{\mathbf{y}^T \mathbf{P} \mathbf{P}^T \mathbf{y}}$.

Putting all this together, the dual problem is

$$\hat{\mathbf{y}} = \arg \min_{\mathbf{y}} \mathbf{y}^T \mathbf{P} \mathbf{P}^T \mathbf{y} \quad (32)$$

$$\text{subject to } \begin{cases} \mathbf{e}^T \mathbf{y} = 1 \\ \mathbf{y} \geq 0 \end{cases} . \quad (33)$$

Recalling the fact that $\mathbf{P} \triangleq \mathbf{Q} \mathbf{\Lambda}_g^{-1/2}$ and $\mathbf{x} \triangleq \mathbf{\Lambda}_g^{1/2} \mathbf{a}$, we may rewrite the dual problem and the optimal primal solution as shown in Section III-B.

REFERENCES

- [1] J-M. Beaulieu and M. Goldberg. Hierarchy in picture segmentation: A stepwise optimization approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):150–163, February 1989.
- [2] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, MA, 1995.
- [3] M. Bhatia. *Wavelet Transform-Based Multi-Resolution Techniques For Tomographic Reconstruction and Detection*. PhD thesis, Massachusetts Institute of Technology, August 1994.
- [4] A. B. Frakt. Multiscale hypothesis testing with application to anomaly characterization from tomographic projections. Master's thesis, Massachusetts Institute of Technology, May 1996.
- [5] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, 1981.
- [6] S. L. Horowitz and T. Pavlidis. Picture segmentation by a tree traversal algorithm. *Journal of the ACM*, 23(2):368–388, April 1976.
- [7] J. Le Moigne and J. C. Tilton. Refining image segmentation by integration of edge and region data. *IEEE Transactions on Geoscience and Remote Sensing*, 33(3):605–615, May 1995.
- [8] E. L. Miller and A. S. Willsky. A multiscale, decision-theoretic algorithm for anomaly detection in images based upon scattered radiation. In *First International Conference on Image Processing*, Austin, Texas, November 1994.

- [9] E. L. Miller and A. S. Willsky. Multiscale, statistical anomaly detection analysis and algorithms for linearized inverse scattering problems, Jan 1997. To appear in *Multidimensional Systems and Signal Processing*.
- [10] L. M. Novak, G. J. Halversen, G. J. Owirka, and M. Hiett. Effects of polarization and resolution on the performance of a SAR automatic target recognition system. *Massachusetts Institute of Technology Lincoln Laboratory Journal*, 8(1), Spring–Summer 1995.
- [11] J. B. Poline and B. M. Mazoyer. Analysis of individual brain activation maps using hierarchical description and multiscale detection. *IEEE Transactions on Medical Imaging*, 13(4):702–710, December 1994.
- [12] D. J. Rossi and A. S. Willsky. Reconstruction from projections based on detection and estimation of objects—parts I and II: Performance analysis and robustness analysis. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-32(4):886–906, 1984.
- [13] H. L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I*. John Wiley and Sons, Inc., 1968.
- [14] B. J. West and A. L. Goldberger. Physiology in fractal dimensions. *American Scientist*, 75:354–365, 1987.